

University of Cambridge
Department of History and Philosophy of Science

**Radical Pluralism,
Ontological Underdetermination,
and the Role of Values in Species Classification**

Stijn Conix

Queens' College

A dissertation submitted for the degree of
Doctor of Philosophy, December 2017

Declaration of Length

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

It does not exceed the prescribed word limit of 80,000 words, including footnotes, references and appendices but excluding the bibliography.

Abstract

The main claim of this thesis is that value-judgments should play a profound role in the construction and evaluation of species classifications. The arguments for this claim will be presented over the course of five chapters. These are divided into two main parts; part one, which consists of the two first chapters, presents an argument for a radical form of species pluralism; part two, which comprises the remaining chapters, discusses the implications of radical species pluralism for the role of values in species classification.

The content of the five chapters is as follows. Chapter 1 starts with a discussion of the theoretical assumptions concerning species and natural kinds that form the broad framework within which the arguments of the thesis are placed. The aim of this chapter is to introduce a set of relatively uncontroversial assumptions that frame the rest of the thesis. On the basis of these assumptions, chapter 2 presents an argument for radical species pluralism. The chapter substantiates this argument with a broad range of examples, and compares this position to other forms of species pluralism. Chapter 3 returns to the main interest of the thesis, namely, the role of values in species classification. It introduces the notion of values and presents an argument for the value-ladenness of taxonomy on the basis of the considerations in the first two chapters. It then sketches three important views on values in science in the literature. Chapter 4 argues that the case presented in chapter 3 provides strong support for one of these views, called the 'Aims View', and against two other prominent views, called the 'Epistemic Priority View' and the 'Value-Free Ideal'. The resulting view, in line with the Aims View, is that value-judgments should play a particularly substantial role in species classification. Chapter 5 then considers the popular assumption that these value-judgments in taxonomy commonly take the shape of generally accepted classificatory norms, and argues that this assumption is not tenable. Finally, a brief concluding chapter points at some implications of the claims and arguments in this thesis.

Acknowledgments

I would like to thank my supervisor, Tim Lewens, for his support and patient feedback on many previous drafts of the thesis. I am also grateful to my advisor, Nick Jardine, for his guidance and detailed comments on all chapters of the thesis. I also want to thank Thomas Reydon and Jacob Stegenga, who were the examiners for this thesis and gave very helpful and constructive feedback.

The AHRC and the department of History and Philosophy of Science provided funding for this PhD, for which I am very thankful.

I would also like to thank my family for their continued support. Finally, I want to thank my wife, Lin, for the quantum leap in my work and for her incredible support throughout my PhD.

Contents

DECLARATION OF LENGTH	1
ABSTRACT	2
ACKNOWLEDGMENTS	3
CONTENTS	4
INTRODUCTION: TAXONOMIC INFLATION	7
1. Splitting and Lumping: Taxonomic Inflation	8
2. Taxonomic Inflation and Values	11
3. The Research Question	17
3.1. Classification and species ontology	18
3.2. Principled, descriptive, and normative	18
3.3. The role of values	20
4. Overview of the Thesis	21
4.1. Chapter overview	21
4.2. Originality of the thesis	23
CHAPTER 1: SPECIES AND NATURAL KINDS	25
1. Species Concepts	26
1.1. The genealogical approach	27
1.2. The process approach	29
1.3. The similarity approach	32
2. Natural Kinds and Individuals	33
3. Natural Kinds	36
3.1. Essentialism	38
3.2. Causal kinds	41
3.3. Bare property cluster kinds	46
4. Summary	49
CHAPTER 2: RADICAL SPECIES PLURALISM	50
1. Moderate Species Pluralism	52
2. Radical Pluralism and Causal Kinds	56
2.1. Selection of mechanisms	57
2.2. Individuating mechanisms	58

3. Radical Pluralism and Species	61
3.1. The genealogical approach and pluralism	61
3.2. The process approach and pluralism	70
3.3. The Evolutionary Species Concept and pluralism	72
4. Different Kinds of Pluralism	74
5. Summary	79
CHAPTER 3: THE ARGUMENT FROM ONTOLOGICAL UNDERDETERMINATION	80
1. Values and Epistemic Standards	80
2. The Argument from Ontological Underdetermination	84
3. Three Objections to Ontological Underdetermination	87
4. Three Views on Values in Science	91
4.1. The Value-Free Ideal	92
4.2. The Epistemic Priority View	97
4.3. The Aims View	104
4.4. A continuum of views	106
5. Summary	109
CHAPTER 4: AGAINST THE WEDGE STRATEGY	111
1. The First Wedge: Internal and External Stages	112
2. Problems for the Internal/External Distinction	114
2.1. Selection of research topics	115
2.2. Ethical constraints	116
2.3. Dissemination of results	117
2.4. Failure of the internal/external distinction	118
3. The Second Wedge: Epistemic and Non-Epistemic Reasons	120
4. The Entanglement of Epistemic and Non-Epistemic Reasons	126
4.1. The non-epistemic in the epistemic	127
4.2. Radical pluralism and ontological underdetermination	129
5. Objection: A weakened Epistemic Priority View	132
6. A Positive Argument from Ontological Underdetermination	136
CHAPTER 5: CLASSIFICATORY NORMS AND NORMATIVE UNDERDETERMINATION	139
1. Classificatory Norms and Radical Pluralism	140
2. Classificatory Norms and the Aims View	144
3. Assumptions of the Classificatory Norms View	148

4. Normative Underdetermination	150
5. Combining Normative and Ontological underdetermination	158
CONCLUDING REMARKS	161
1. The Epistemology of Species Delimitation	161
2. Values in Science	163
3. Values and Natural Kinds	164
BIBLIOGRAPHY	166

Introduction: Taxonomic Inflation

In 2007, *The Economist* published an editorial which stated that 'Linnaeus's system is being subtly debauched by over-eager taxonomists, trying to help conservation' (Leaders, 2007, p. 1). According to this editorial, species numbers of charismatic taxa like primates are rising strongly, but not simply due to the discovery of new, legitimate species. Instead, the editorial claims, many of the newly accepted species are groups to which taxonomists assign species status merely in the hope that this will benefit the groups' chances of receiving special conservation status. In other words, the claim is that taxonomists recognize certain groups of organisms as species not because the biological data suggest this, but because of their preferences, interests and value-judgments. Karl and Bowen refer to the highly endangered black turtle (*Chelonia agassizii*) as an example of this practice. They write that 'conservationists cling to dubious taxonomy for the sake of protecting wildlife within existing legal frameworks' (Karl & Bowen, 1999, p. 966).

In the comments section of the online version of the editorial, several taxonomists reacted strongly to the claims in *The Economist*. Stepan (2007), a taxonomist, wrote that '[s]ystematists (taxonomists) are scientists and generally do not nor should they consider the "value" of taxa in drawing their conclusions'. Ruedas and Brown (2007) wrote that 'we (taxonomists) labor in obscurity, we labor in penury, we labor with love, but our objectivity should not be questioned'. In short, these taxonomists claim that, generally, taxonomy is not influenced by such value-laden considerations, and that this would be highly inappropriate if it were the case. The aim of this thesis is to delve deeper into this issue, and investigate the role of values in biological species classification. More precisely, three broad questions will guide this thesis; first, do values play a role in species classification? Second, do we want values to play a role in species classification? And finally, if values influence species classification, what precisely is their role?

In this chapter I introduce these questions by discussing the case of 'taxonomic inflation', a recent incarnation of the debate between splitters and lumpers, and subject of the *Economist* editorial referred to above. Section 1 introduces taxonomic inflation, and gives a brief

overview of the scientific debate regarding this notion. Section 2 discusses the role values play in this debate. This raises the question about the appropriate role of values in species classification. Section 3 connects this to the research question of this thesis. Finally, section 4 gives a brief overview and summary of the five chapters of this thesis, and points out which of its claims and arguments are novel.

1. Splitting and Lumping: Taxonomic Inflation

An old and recurring debate in taxonomy occurs between so-called ‘splitters’ and ‘lumpers’. The prominent evolutionary theorist George Gaylord Simpson (1945, p. 23) jokingly described this debate as follows:

Splitters make very small units – their opponents say that if they can tell two animals apart, they place them in different genera, and if they cannot tell them apart, they place them in different species. Lumpers make large units - their opponents say that if a carnivore is neither a dog nor a bear they call it a cat.

This opposition between recognizing fewer, more inclusive categories, and smaller, more precise categories in biological taxonomy is most prominently associated with debates concerning species delimitation. Both splitter and lumper taxonomic cultures recur regularly throughout the history of biological species classification. A classic instance of splitting is Merriam’s proposal (1918) to recognize eighty-two different species of the north-American grizzly bear (today recognized as a single species). A classic instance of lumping occurred in zoological taxonomy of the early 20th century, when it was discovered that many groups formerly recognized as distinct species intergrade at places where they overlap. Instead of recognising these intergrading groups as distinct species, they were grouped into large polytypic species containing multiple subspecies or geographical varieties (see Mallet, 2004; Poulton, 1903).

This debate is not merely a historical curiosity from the time classification was based on morphological similarity. The controversy concerning lumping and splitting now occurs between different forms of evolutionary classification. One recent example of this is the debate about ‘taxonomic inflation’. While biodiversity on Earth is decreasing rapidly, recorded species numbers are increasing at a rate far higher than ever before. Some authors argue that this is simply due to continued scientific effort and new molecular methods which allow us to

discern boundaries between species that were previously not visible. However, others argue that the rapid increase in species numbers is due to a particular taxonomic culture of splitting that has been increasingly popular since the 1980s, and they call the increase due to this splitting-culture 'taxonomic inflation'.

More specifically, taxonomic inflation refers to the strong increase in species numbers not due to new data or discoveries, but due to a change in the species concept or the criteria for species delimitation. While the term has been used occasionally throughout the history of taxonomy, it became particularly prevalent after a controversial paper by Isaac, Mallet and Mace (2004). In that paper, the authors claim that a significant proportion of the increase in species numbers since the 1980s is caused by a general switch by taxonomists from the Biological Species Concept (BSC) to the Phylogenetic Species Concept (PSC). The BSC defines a group of organisms as a species if they interbreed and are reproductively isolated from other such groups. The PSC, on the other hand, defines species as groups of organisms that are diagnosable, i.e. that share at least one inherited trait not shared by organisms outside that group.¹ As interbreeding groups of organisms often consist of multiple smaller diagnosable groups, the BSC generally picks out larger units than the PSC. Consequently, many groups recognised as subspecies by the BSC count as species under the PSC. Thus, Isaac et al. (2004) claim that the switch from the BSC to the PSC leads to splitting existing species into smaller groups, and consequently to a substantial increase in the total number of recognized species. They use the term 'inflation' for this phenomenon because the increase in species numbers does not reflect newly discovered variation, but rather the splitting of this variation into smaller chunks. As a consequence of this splitting, each unit represents less 'diversity-value'.

It is important to see these claims against the historical background of technological and methodological changes in taxonomy. Isaac et al. (2004) claim that the switch from the BSC to the PSC happened in the course of the 1980s. At around the same time, taxonomy changed drastically with the introduction of molecular methods. Until the late 1970s, species classification was mostly based on morphological characters. While morphological description is still required for accepting new species, the reliance on morphological data is now largely replaced by the use of various forms of molecular data such as allozyme data, chromosome

¹ There are several distinct species concepts that are often referred to as the PSC. The PSC referred to in this chapter is the diagnosability-based PSC, first introduced by Cracraft (1983). I come back to the PSC, the BSC and other species concepts in chapters 1 and 2.

numbers, and, most importantly, DNA data. The switch to molecular data uncovered many morphologically cryptic groups, i.e. groups that are morphologically similar to sister groups and can only be distinguished on the basis of molecular data. Such cryptic variation is common because small groups of organisms can rapidly become diagnosable on a molecular level through drift or strong selection pressures. Such diagnosability on the molecular level often arises long before the groups become morphologically diagnosable. For this reason, molecular data is more likely to uncover many small diagnosable groups than morphological data. This is important because the PSC delimits species on the basis of diagnosability. Combining the PSC with molecular data thus leads to even more splitting and an even larger increase in species than the switch to the PSC alone.

Isaac et al.'s (2004) paper started a vigorous scientific debate. Various authors argue that excessive splitting is indeed a significant problem in contemporary taxonomy. For example, the number of Lemur species has increased from thirty-six to over a hundred in twenty-five years. Tattersall (2007, p. 21) claims that only forty-eight of these Lemur species are valid, while the rest of the increase is due to the adoption of the PSC. Heller et al. (2013) claim that a recent revision of bovid species by Groves and Grubb (2011) is strongly distorted by taxonomic inflation. They claim that the increase from 143 species to 279 is mainly due to the use of diagnosability as the sole criterion for species delimitation. In line with this, Zachos et al. (2013; see also Zachos, 2014) claim that many newly recognized mammal species, including a tiger species, European red deer species and ten klipspringer species, are the result of adopting the PSC and do not constitute newly discovered diversity.

Taxonomists have argued against these claims concerning taxonomic inflation in two distinct ways. The first form of argument is exemplified by Sangster's (2009, 2014) studies on avian taxonomy, which investigate the criteria for species delimitation used since 1980. His results show that even though there was a very strong increase in species numbers, there was no switch from the use of reproductive isolation (BSC) to diagnosability (PSC) as the main criterion. Consequently, Sangster argues, Isaac et al.'s claims about taxonomic inflation do not apply to avian taxonomy. Köhler et al. (2005) come to a similar conclusion in a worldwide study on the number of recognized amphibian species, which increased by 25% between 1992 and 2004. Their research shows that the strong increase is mainly due to exploration of the species-rich tropical areas. Almost all newly discovered species had not been collected before,

and only very few species were a result of raising subspecies to the species-level. Thus, these papers suggest that if taxonomic inflation occurs at all, it only affects some parts of taxonomy.

The second form of argument against taxonomic inflation holds that the PSC and diagnosability form a better approach to species delimitation than the BSC and reproductive isolation. This is how several authors (Cotterill, Taylor, Gippoliti, Bishop, & Groves, 2014; Groves, 2013) respond to Heller et al.'s (2013) claim about taxonomic inflation in bovid species. They argue that the use of diagnosability for species delimitation is preferable to using reproductive isolation because the latter is hard to operationalize. Instead of defining taxonomic inflation as a problem for taxonomy, they raise the opposite problem of 'taxonomic conservatism' (Cotterill et al., 2014, p. 1) and 'inertial bias' (Bernardo, 2011) in the work of those who argue against this splitting culture. It is interesting to note that this second counterargument does not question any of the empirical claims about taxonomic inflation, but simply evaluates the phenomenon differently: instead of inflation, the increase in species numbers is scientific progress. In the next section, I turn to the role of such value-laden judgments in the debate about taxonomic inflation.

2. Taxonomic Inflation and Values

The scientific debate on taxonomic inflation has received wide attention. This is not surprising as the impact of taxonomic inflation is not limited to the field of taxonomy, but extends to all uses of species classification, species lists and species numbers. Faurby et al. (2016), for example, argue that taxonomic inflation impacts macro-evolutionary research such as analyses of lineage diversification, and Isaac et al. (2004) point out that species are important units in macroecological research. For example, the interpretation of macroecological patterns such as the increase of species ranges at higher latitudes, known as 'Rapaport's rule', crucially depends on the species classifications that are used. If these are inconsistent because of inflation, the outcome of ecological research is affected too.

Taxonomic inflation is also likely to affect conservation biology, because conservation planning is often dependent on measures derived from species numbers, such as endemism, species richness and species abundance (Fleishman, Noss, & Noon, 2006). More precisely, if taxonomic inflation only occurs in some taxa and not others (as we saw indicated by the works of e.g. Sangster (2009) and Köhler et al. (2005) cited above), it could make species lists

inconsistent and consequently species numbers meaningless. In a similar way, taxonomic inflation is also likely to affect studies that investigate changes in biodiversity over time. Taxonomic inflation leads to an increase in species richness, species abundance, endemism, extinctions and threatened species, even if no new species actually went extinct or came into existence over the time of measurement. Finally, it is also worth noting that a switch from the BSC to the PSC and consequent changes in species classification change the areas recognized as biodiversity hotspots. This has significant consequences because these hotspots play an important role in area-based conservation planning (Mittermeier, Turner, Larsen, Brooks, & Gascon, 2011; Zachos & Habel, 2011).

The effect of taxonomic inflation on biodiversity conservation is particularly interesting because it shows that the phenomenon can have profound consequences outside science. Species that were not threatened before inflation might receive special conservation status after being split into several smaller groups. For example, revision of the genus *Sorbus* by Meyer et al. (2005) led to the recognition of many new species which subsequently received special conservation status from local protection authorities.² More generally, changes in the criteria for species delimitation and related changes in species numbers may make conservation policies and laws ineffective (see Mace, 2004). Similarly, conservation targets expressed in species numbers may become impossible to meet or are rendered meaningless due to taxonomic inflation.

It is clear then, that the debate about taxonomic inflation is not merely an entirely theoretical controversy. Taxonomic inflation has important consequences in various branches of science and their applications, and for biodiversity conservation. In this way, the scientific controversy about taxonomic inflation is entangled with value-judgments related to these consequences. This entanglement of values and taxonomy was also emphasized in the *Economist* editorial on taxonomic inflation, cited at the start of this chapter. This editorial (Leaders, 2007) claimed that an important motivation behind the recognition of small species lies in how conservation efforts often depend on the rank of a group. More precisely, it wrote that taxonomists 'debauch' taxonomy by recognising certain groups of organisms as species in the hope of improving the conservation chances of those groups. This also explains, according to the

² For more examples and an overview of the effects of taxonomic revisions on conservation status, see Morrison III et al. (2009).

author, why taxonomic inflation is particularly strong in charismatic species, like primates and other large mammals; classification in those charismatic groups is not only guided by the data, but also by the interest and appraisal of the taxonomists.

While one may think that such a substantial role for values in taxonomy is uncommon, there are a number of cases that confirm the practice described by *The Economist*. For example, recent research on the classification of Australian herpetofauna has been troubled by ill-supported and radical taxonomic revisions. In several controversial self-published articles, Rob Wells and Ross Wellington (1984, 1985) recognized almost 500 new species, changing the nomenclature and taxonomic arrangement of nearly all amphibians and reptiles in New Zealand and Australia. Many of the new species were subspecies raised to species status without new research or evidence. Even if some of the revisions were later substantiated by further research, the work of Wells and Wellington is generally considered ‘unscientific’ (Williams, Wüster, & Fry, 2006, p. 925): they provide no or very little evidence or justification for the changes they propose, they describe species that they have not investigated, and they name new species in bizarre ways (such as ‘Vaderscincus’, after Darth Vader). Because the publications followed the appropriate rules for the recognition of new species, the International Code of Zoological Nomenclature refused to suppress the names despite the pleas of 150 professional herpetologists trying to avoid the taxonomic chaos that these changes caused. When asked for the reasons for publishing these taxonomic revisions, Wells and Wellington state that they wanted to revitalize the debate about the classification of Australian herpetofauna by forcing other taxonomists to disprove or affirm their revisions. Other authors (Kaiser et al., 2013; Williams et al., 2006) claim that name-bagging, taxonomic terrorism (i.e. intentionally confusing existing classifications by incorrect revisions) and scooping taxonomic revisions from other taxonomists were important motivations.

This chaos in Australasian reptile taxonomy has been perpetuated by the work of amateur taxonomist Raymond Hoser (Hoser, 1989, 1991; see Naish, 2013; and Williams et al., 2006 for discussion). Hoser proposes extensive taxonomic changes through papers published either in self-edited journals or those without peer review processes. Like Wells and Wellington, Hoser’s proposals strongly inflate species numbers, often by raising subspecies to species status without providing evidence or justification, or relying on minimal evidence such as emails from museum staff (Wüster, Bush, Keogh, O’Shea, & Shine, 2001). Williams et al. (2006)

even claim that in many cases Hoser refers to non-existent DNA evidence. Hoser was himself responsible for 76% of the taxonomic revisions in Australian snakes between 2002 and 2013, all in single-authored, self-published papers (Kaiser et al., 2013). Unsurprisingly, several professional taxonomists have pointed out that Hoser's work is problematic in many ways, and have argued that name-bagging and even a personal grudge against professional taxonomists are among the main motivations for his radical taxonomic revisions. However, Hoser's revisions follow the rules of the ICZN, and therefore cannot be suppressed easily without proper investigation of the revisions.

It is generally accepted that Hoser, Wells and Wellington's decisions concerning species classification were not based on empirical evidence, but on their personal values and preferences: name-bagging, scooping taxonomic revisions, personal grudges and revitalizing taxonomic work (Kaiser et al., 2013; Naish, 2013; Williams et al., 2006; Wüster et al., 2001). Since values clearly play a significant role in species classification in such corrupted science, one might make the inference that as long as values are kept at bay, science will not be corrupted. This is important, as the practices of these Australian taxonomists are not the norm in taxonomy. Indeed, most taxonomic revisions differ from Hoser's, Wells's and Wellington's cases in that they are published in peer reviewed journals and are well supported by empirical evidence. As section 1 of this chapter has shown, there is an extensive scientific debate about taxonomic inflation, published in peer-reviewed journals such as *TRENDS in Ecology and Evolution* and *Systematic Biology*.

While it might appear then that values only enter corrupted science, this is not the case: closer consideration of the debates on taxonomic inflation shows that value-judgments play a role in this scientific debate, and, more generally, in scientific decisions about splitting or lumping species. As discussed earlier, taxonomic inflation has potentially severe implications for biodiversity conservation. Thus, conservation values are entangled with the debate through the potential consequences of splitting and lumping. It is also worth noting that the term 'inflation' already expresses a value-judgment, and is countered by Bernardo (2011) with the equally value-laden term 'taxonomic conservatism'. Another example of value considerations becoming entangled with the scientific debate is the alleged role of species charisma. Isaac et al. (2004) argue that taxonomic inflation is strongest in species that humans find charismatic, like primates and other large mammals. Other values recurrently mentioned in taxonomic

literature are pragmatic or epistemic, such as the ease of operationalization of a species concept or the precision of criteria for species delimitation (Cotterill et al., 2014).

Two examples from the literature on taxonomic inflation provide particularly interesting illustrations of the role that value-judgments play in decisions about splitting and lumping. In a series of articles, Frankham and co-authors (Frankham, 2015; Frankham et al., 2012, 2014) discuss the effect of habitat fragmentation on species. The habitats of many species have been fragmented due to human impact on the environment. As a result, many smaller groups that occasionally exchanged genes before fragmentation are now reproductively isolated. In such cases, the BSC and PSC lead to different species classifications. The PSC recognizes groups as species if they form a diagnosable unit. The existence of such PSC species is promoted by habitat fragmentation, because small, isolated populations can rapidly become diagnosable through drift. The BSC, on the other hand, is more likely to recognize such fragmented populations as a single species, because reproductive isolation usually evolves slowly and, consequently, interbreeding would have been likely had their habitat not been fragmented.

Frankham et al. (2012, 2014) argue that in such cases the BSC should be preferred over the PSC because using the latter would be detrimental to the conservation of these groups. Because of the founder-effect and their small size, these isolated groups are often at risk of inbreeding depression and subsequent extinction. The best way to save these groups from inbreeding depression is by restoring gene-flow between the different isolated populations. For this to happen, however, the different groups would have to be part of a single conservation unit instead of several separate conservation units. Thus, conservation of these groups would be more efficient and have better chances if the different groups were recognized as a single interbreeding species.

Conservation values play a crucial role in these arguments. The reason we should prefer the BSC over the PSC, according to Frankham et al. (2012), is that the former is better for biodiversity conservation. Thus, they argue that conservation values should determine whether we split or lump the isolated groups. This does not mean that empirical considerations do not play a role in their arguments. Indeed, they provide empirical support for the claim that the PSC is detrimental to species conservation, and they provide support for the claim that the BSC would not suffer from these problems. However, these empirical

arguments can be distinguished from the normative claim that we should lump groups rather than split them to increase the likelihood of these groups surviving.

The second example concerns classification of *Anopheles*, a species-rich mosquito genus containing hundreds of species. In a contribution aptly titled ‘what are species and why does it matter’, Attenborough (2015, pp. 140–146) argues that effective prevention and combat of malaria require species classification that is as specific and detailed as possible, effectively consisting of PSC species rather than BSC species. More precisely, Attenborough points out that there are many morphologically cryptic groups with fixed genetic differences. While these groups cannot be readily distinguished except by molecular data, the variables affecting malaria transmission often differ between these groups. Thus, effective combatting of the transmission of malaria requires distinguishing between these groups. He argues that species classifications that overlook these differences risk leading to interventions that likewise overlook these differences. Even though Attenborough also points to the epistemic virtues of the PSC, he emphasizes that using the PSC is important ‘to improve human health in the tropical Western Pacific, sub-Saharan Africa and other places still greatly afflicted by this scourge’ (Attenborough, 2015, p. 147). Thus, he explicitly invokes social or moral values as a criterion to choose the PSC, and consequently to split rather than lump.

The point of these examples is to show that value-judgments do not only play a role in ‘bad’ taxonomy, like that of the three Australian taxonomists and the phenomenon discussed in *The Economist*. Frankham et al. (2012) and Attenborough (2015), two well-respected biologists, both explicitly use value-judgments as reasons for accepting one group of organisms as a species rather than another. Even though there are significant differences between these two examples and that of the Australian taxonomists, it is easy to see that they are significantly similar too. In all these cases, decisions about splitting or lumping are made on the basis of value-judgments. In the case of *The Economist* and the Australian taxonomists, the value-judgments are related to the charisma of the groups, their conservation value, and various personal preferences of the taxonomists. In Frankham et al.’s and Attenborough’s papers, decisions about splitting and lumping are made on the basis of conservation values and reasons concerning human health respectively.

3. The Research Question

Section one of this chapter gave a brief overview of the scientific debate about taxonomic inflation. It showed that various empirical and methodological disagreements between taxonomists lead to different interpretations of the recent increase in species numbers, and ultimately to different classifications of organisms into species. The second section showed that values are part of this controversy concerning splitting and lumping in various ways. Some of these value-influences, for example the role of name-bagging and personal grudges in the classifications of Wells and Wellington, seem undesirable because they affect science in a negative way. The cases of population fragmentation and *Anopheles* taxonomy, on the other hand, may suggest that it is sometimes desirable that values play a role in taxonomy. Indeed, species classifications are constructed for their use in various practical domains; it would seem unwise not to take into account these uses when constructing the classifications.

More generally, these examples suggest that there are two conflicting intuitive attitudes with regard to the role of values in species classification. The first is nicely illustrated by Ruedas', Brown's and Stepan's reactions to the article in *The Economist* (referred to above), which emphasise the objectivity of taxonomy to counter claims about value-ladenness. Similarly, taxonomists also frequently refer to the importance of objective and evidence-based taxonomy in their arguments against Wells, Wellington and Hoser. Most tellingly, Williams ends his discussion of Wells' and Wellington's revisions arguing against their contributions because 'conviction is no substitute for evidence' (Williams et al., 2006, p. 926).

The general gist of these reactions is the same: value-laden considerations such as charisma, personal preferences or the conservation value of species should not influence the decision of whether a particular group is a species. The motivation for these reactions is the belief that values somehow pose a threat to the objectivity of species classifications. The aim of taxonomy is to discover, classify and name groupings in the organic world; thus, classifications should depend only on what the world is like, and not on the value-judgments of the taxonomists who make the classifications. If value-judgments play a profound role, as in the case of the Australian taxonomists, species classifications risk reflecting these values rather than the world.

The second intuitive attitude is directly opposed to the first. While it is obvious that some value-influences are damaging, there are also many value-influences that seem harmless and even desirable for taxonomy. The most obvious examples here come from epistemic values such as empirical adequacy and internal consistency, or pragmatic values such as operationalizability and simplicity. The mosquito and inbreeding-depression cases, both discussed at the end of section two, also provide convincing examples: if splitting and lumping have severe consequences for malaria-prevention and biodiversity conservation, should taxonomists not take these considerations into account when constructing classifications? More generally, one might argue that classifications are man-made representations or even tools that are constructed for both scientific and practical purposes. Accordingly, it seems sensible to claim that value-judgments should play a role in constructing these classifications. These opposing but plausible intuitions concerning the role of value-judgments in taxonomy raise the question of what constitutes an appropriate role for values in species classification. More precisely, the question is whether values should influence taxonomy, and if so, what kind of role they should play. The next three subsections further clarify the scope of these questions.

3.1. Classification and species ontology

The aim of this thesis is to provide a better understanding of the role that values play in species classifications as they are constructed by biologists. Following this line of thought, the investigation will focus on the classifications, and not the things that are classified and represented by these classifications. This does not mean, however, that I will not discuss species ontology. In order to understand the role of values in scientific species classifications, it is important to understand what scientists mean by 'species', and how they go about classifying them. For this reason, the first two chapters of the thesis attempt to make explicit the view on species ontology that is implied by current taxonomic practice and theory. I then take this view on species ontology as a starting point, and investigate its implications for the role of values in constructing species classifications.

3.2. Principled, descriptive, and normative

This question about the role of values in species classification can be interpreted in both a descriptive and a normative sense. In the descriptive sense, the question simply is whether

species classifications are value-laden. In the normative sense, the question is whether species classifications should be value-laden. This thesis aims to address both these questions.

There is a sense in which the answer to the descriptive question is trivially affirmative. The examples of the Australian taxonomists show that values played a role in the construction of at least some species classifications. More generally, one could argue that it is unavoidable that sometimes scientists' decisions are, willingly or unwillingly, influenced by their personal value-judgments. As Rudner (1953, p. 1) points out, 'the scientist cannot escape his quite human self – he is a "mass of predilections" and these predilections must inevitably influence all of his activities not excepting his scientific ones'. While such trivial value-ladenness is interesting and can have severe consequences (as the case of Hoser and Wells suggests), I will not consider it further in this thesis. I do this for the same reasons that I call this value-ladenness 'trivial'. First, this kind of value-ladenness does not follow from anything particular about species classification. Rather, it is a consequence of general characteristics of science and scientists, and is in no way different from the way values enter any other scientific field. Second, and in relation to this, trivial value-ladenness is accidental rather than necessary. It is caused by careless or fraudulent scientists, and may be avoided by processes such as peer review. It does not affect taxonomy in a systematic way, but on a case-by-case basis. Thus, it is easy to think that, if we are careful enough and install appropriate procedures, such value-ladenness can be avoided.

Instead, one might ask whether there is anything particular about species classification that gives us reason to think that values play a significant role, and whether this role of values is an unavoidable rather than accidental characteristic of species classifications. An affirmative answer to these questions would have philosophically interesting implications on the way we think about species classifications. In this thesis, I will defend such an affirmative answer. I argue that species classifications are inevitably value-laden as a direct consequence of the scientific view on species ontology. Given that this value-ladenness follows directly from the assumed ontology of species, it cannot be avoided by more careful research or processes such

as peer review. Thus, this value-ladenness is both characteristic of species classifications and inevitable.³

The question about the role of values in species classifications can also be interpreted in a normative sense: *should* values play a role in classifying species? It might seem that this question is already partially answered by the descriptive question: if species classifications are unavoidably value-laden, it does not make sense to claim that they should be value-free. However, the two positions are more independent than this suggests. One can coherently claim that species classifications are unavoidably value-laden, but that we should aim to limit the influence of these values as much as possible. Thus, the inevitable value-ladenness of species classifications is compatible with the ideal of value-freedom. The normative question then is whether taxonomists should aim for value-freedom (regardless of whether this is completely possible), or whether it is desirable that values play a role in species classifications. Another way of putting this question is whether we should limit the role of values in constructing species classifications as much as possible, or whether values have a positive role to play. My answer to the normative question will support the latter option, i.e. that good species classification requires that values play a profound role.

3.3. The role of values

The aim of this thesis is not merely to show that values do and should influence species classification. I also want to clarify what role they play and how profound this role is. Very broadly put, values can be taken to play a superficial role if they are involved in some decisions in the ‘external stages’ of science, such as the choice of research topics, without substantially affecting the research process and outcomes of taxonomy. On the other hand, values play a profound role if they are substantially involved in scientific decisions throughout the research process, including decisions in the ‘internal stages’, such as the interpretation of evidence and the acceptance of hypotheses.

The argument for the value-ladenness of taxonomy in this thesis will focus on one particular way in which values enter taxonomy, namely, through what I will call ontologically underdetermined decisions. Consequently, the evaluation of the role of values in taxonomy

³ This is not to say that there are no other scientific representations that are value-laden for similar reasons as the ones discussed in this thesis. However, I will focus only on species classification, and do not commit to the application of my arguments outside this domain.

will also focus only on values that enter in this particular way. However, this does not limit the scope of the arguments. While ontological underdetermination is only one of many ways in which values enter taxonomy, I argue that it provides good reasons to think that values do and should play a substantial role in both the internal and external stages. This will lead to the main conclusion of this thesis, namely, that values should play a profound role in the classification of species. Other sources of value-ladenness may further strengthen this conclusion; however, the arguments in this thesis are not dependent on them.

4. Overview of the Thesis

Before I set out these arguments, it will be helpful to give a brief overview of the structure and contents of this thesis, and to point out how and where its claims are original. I do the former in subsection 4.1, and the latter in subsection 4.2.

4.1. Chapter overview

As discussed above, the main aim of this thesis is to get a better understanding of the role that values should play in species classification. This question naturally leads to a division of the thesis into two main parts. The first part (chapters 1 and 2) investigates the principles that guide scientific species classification. The conclusion of this first part will be that these principles are such that there are innumerable many equally good ways of classifying organisms into species. The second part (chapters 3, 4 and 5) then considers the consequences of this radical species pluralism for the role that values should play in species classification. The conclusion of this second part will be that values should play a profound role in classifying organisms into species. I will develop these arguments over the course of five chapters. The main claims of these chapters are as follows.

Chapter 1 sets the stage for the arguments of the remaining chapters. It introduces and clarifies the notions of 'species' and 'natural kinds', and points out which parts of the extensive debates on these two notions are relevant to the arguments in this thesis. More precisely, the aim of this chapter is to present the principles on which taxonomists rely for individuating species, and to connect this on a more abstract level to the individuation of natural kinds. Chapter 2 then argues that if we assume these principles for individuating kinds and species, it follows that there are innumerable different ways of carving up the organic world into species. I call this position 'radical species pluralism'.

Chapter 3 returns to the main interest of the thesis, namely, the role of values in species classification. It first introduces and substantiates an argument for the value-ladenness of taxonomy that follows directly from radical species pluralism. This argument, which I call the ‘argument from ontological underdetermination’, shows that values are needed in taxonomy to choose between multiple ontologically good classifications. The second part of the chapter then introduces three competing views for thinking about and regulating this role for values in taxonomy. These views all acknowledge that taxonomy should be value-laden, but disagree about the profoundness of this value-ladenness. The first, called the ‘Value-Free Ideal’, holds that values have no role to play in the internal stages of science. The second, called the ‘Epistemic Priority View’, holds that values can only influence science if they do not interfere with epistemic considerations. The third, called the ‘Aims View’, holds that values and epistemic considerations should stand on equal footing in both the internal and external stages.

The fourth chapter investigates which of these three views on values in science best captures the role of values that enter through ontological underdetermination. I show that the Value-Free Ideal and the Epistemic Priority View, both of which imply a relatively superficial role for values in taxonomy, rely on assumptions that are incompatible with the argument from ontological underdetermination. Consequently, I argue that the Aims View, which implies a particularly profound role for values in both the internal and external stages of taxonomy, best captures the role that values should play in taxonomy.

Chapter 5 then considers a popular view in the literature on natural kinds that is significantly similar to the view defended in chapter 4. This popular view holds that values influence taxonomy in the shape of generally accepted classificatory norms that determine which features of the world to focus on for classification. After presenting this view, the chapter argues that generally accepted classificatory norms cannot fully capture the way values influence taxonomy due to ontological underdetermination. It is argued that in addition to selecting which features of the world to focus on, values also play a role in shaping the species that are recognised. This suggests a profound role for values in line with the conclusions of chapter 4.

The concluding remarks, finally, sum up three implications of the profound value-ladenness of taxonomy, and point to further research questions suggested by the arguments in this thesis.

4.2. Originality of the thesis

The claim that there are multiple good ways of classifying organisms into species is of course not new. After convincing arguments by Dupré (1981, 1993, 1999), Kitcher (1984b) and Ereshefsky (1998, 2001, 2010a) among others, species pluralism has arguably become the consensus view among philosophers of biology. Similarly, the claim that values play a substantial role in science is accepted by most if not all philosophers and scientists (Anderson, 1995a; Douglas, 2009; Kitcher, 2001; Kourany, 2010; Longino, 1990). In addition, both species pluralism and the value-ladenness of science are the subjects of longstanding debates with practically innumerable contributions. It is therefore worth briefly indicating what the arguments and claims in this thesis contribute to these debates, and in what respect these contributions are original. These points will be discussed in more detail over the course of the thesis.

The variant of ontological species pluralism that I attribute to taxonomy can be distinguished from other pluralist positions by its radical character.⁴ Most significantly, it is more radical than Ereshefsky's (2001) and Kitcher's (1984b) quite moderate forms of species pluralism, and even more radical than Dupré's (1981, 1993) promiscuous realism and Stanford's (1995) anti-realism. Positions similar to radical species pluralism have been proposed in closely related debates about populations (Stegenga, 2016), evolutionary groups (Barker & Velasco, 2013), and, more generally, natural kinds (Craver, 2009). Chapter 2 points towards these positions, and discusses in detail the case of radical species pluralism.

It should be noted that the focus of my arguments is on the implications of species pluralism for the role of values in species classification, and not on species pluralism itself. Thus, the main original contributions of this thesis can be found in the last three chapters, and concern the role of values in species classification. In line with an increasing consensus in the philosophy of science, I claim that values play a substantial role in taxonomy. There are three main ways in which this claim takes an original shape in this thesis.

First, while most debates about values in science focus on theory choice, this thesis looks at the role of values in scientific classification. Despite the *prima facie* obvious entanglement of classification with value-judgments, surprisingly little philosophical attention has been paid to

⁴ I come back to the difference between my radical pluralism and these other variants in chapter 2, section 4.

this subject in the debates on values in science. While many authors hint at the connection between values and classification (Anderson, 2004; Dupré, 2007a; Ludwig, 2015; Stanford, 1995), very few have explored it in detail. This is unfortunate, as scientific classification is interestingly different from the other kinds of scientific outcomes that debates about values in science usually focus on. By investigating the case of species classification, this thesis aims to make a step in overcoming this relative gap in scholarship.

Second, my argument for the value-ladenness of taxonomy, which I call the argument from ontological underdetermination, is different from the arguments of epistemological underdetermination that have received the lion's share of attention in the philosophical literature on values in science (Douglas, 2009; Kitcher, 2001; Longino, 1990; Steel, 2010). Instead of citing uncertainty as the source of values in science, I refer to pluralism. Several others, particularly in the context of species pluralism, have hinted at similar arguments (Anderson, 2004; Dupré, 1993; Ereshefsky, 2001; Kitcher, 1984b; Ludwig, 2015; Stanford, 1995). However, these authors rely on moderate rather than radical pluralism. As I will argue in chapter 4, the difference between moderate and radical pluralism has significant implications for the role of values that follows from this pluralism. Moreover, these similar arguments rarely extend beyond brief remarks or hints, and the role for values that they support is never developed in much detail. In addition to proposing a new and detailed form of this argument, chapters 3, 4 and 5 also extensively discuss how profound this particular role of values is, and how it relates to other ways in which science is value-laden.

Finally, this thesis connects this argument from ontological underdetermination to a particularly profound and substantial role for values in science. More particularly, I argue that the argument from ontological underdetermination supports a role for values on equal footing with epistemic considerations in all stages of research. This is noteworthy because arguments similar to that one presented here are usually associated with a relatively superficial role for values in science. Moreover, the profound role for values defended in this thesis goes against what is arguably the most popular view on values in science, namely, the Epistemic Priority View (Douglas, 2009; Steel, 2010; Steel & Whyte, 2012), and provides support for the more controversial Aims View (M. J. Brown, 2013; de Melo-Martín & Intemann, 2016; Kourany, 2010).

Chapter 1: Species and Natural Kinds

This thesis investigates the role of values in the construction of species classifications. Thus, it is important to clarify what is meant by the notions species and classification. This chapter provides such clarification by giving a brief overview of recent philosophical work on the species concept and natural kinds. These two notions are, of course, the subject of old and ongoing debates with innumerable contributions and distinct positions. In this chapter, I do not aim to add new arguments or insights to these debates. Instead, the aim is to set the stage for the arguments in the following chapters. To do this, I point out which parts of these debates are irrelevant to the questions of this thesis, and only focus on the relevant parts of the debates; in doing so, I provide a broad overview of the main philosophical positions that are plausible given the purposes of this thesis without committing to any particular one of these. This will allow me to show in later chapters that my arguments apply over a wide range of current views on species and natural kinds. It should be noted that the discussion of these views in this chapter remains on a very general level. Apart from obvious reasons of space, this is because more detail is not required for answering the leading questions of this thesis. Where necessary for my arguments, chapter 2 will add more detail to the discussion in this chapter.

The structure of the chapter is as follows. Section 1 clarifies the notion of species by introducing three main approaches to species classification used in taxonomy, namely, the genealogical approach, the process approach, and the similarity approach. Sections 2 and 3 then elaborate on the relation between these classifications and the world. Section 2 argues that while we can think of this relation both in terms of natural kinds and individuals, it is for the purposes of this thesis most convenient to adopt the language of natural kinds. Section 3 elaborates on the notion of natural kinds; it presents three views on natural kinds, namely essentialism, a causal view, and a bare property cluster view, and argues that only the latter two are suitable for the purposes of this thesis. Section 4 ends the chapter with a brief summary.

1. Species Concepts

There is abundant literature on what species are, and more than thirty different species concepts have been proposed over the past decades (Coyne & Orr, 2010; Ereshefsky, 2001; Mallet, 2001; Mayden, 1999). I will not discuss all of these here. Instead, I classify these concepts into three general approaches. Given that I am interested in the role values play in delimiting species, this taxonomy of approaches is based on what is considered to individuate groups of organisms as species. These three approaches are the genealogical approach (section 1.1), which individuates species on the basis of genealogical history; the process approach (section 1.2), which individuates species on the basis of evolutionary processes responsible for evolutionary cohesion between organisms; and the similarity approach (section 1.3), which individuates species on the basis of similarity between organisms. All groups recognised as species by taxonomists are individuated on the basis of one of these approaches, or some combination of them.

Before I discuss these approaches in more detail, two remarks are in place. First, the general description of these approaches does not do justice to the sophistication of the various species concepts they cover, nor does it touch on all issues relevant for evaluating them. However, for the purposes of this thesis, the taxonomy of approaches used here provides sufficient detail. It is not my aim to defend one particular view of species ontology, or even to defend the plausibility of all these views; rather, I merely want to give an overview of the kinds of groups taxonomists recognise as species. Second, the taxonomy of approaches given here is merely pragmatic, and some species concepts fit in more than one category. Hence, other taxonomies of species concepts are possible too; examples can be found in Hull (1997), Mallet (2001), Baum (2009) and Reydon (2005). These different classifications emphasize different aspects of species concepts, and the one used here simply offers the most convenient exposition for the purposes of this thesis. Thus, this taxonomy should not be taken to have any metaphysical implications. More precisely, by adopting this taxonomy I do not commit to all these species concepts as being part of a single category (see Ereshefsky, 1998) or to all species as being the same kind of entity (see Reydon, 2005).

1.1. The genealogical approach

The genealogical approach defines species on the basis of the unique common genealogical history of the members of the species. To be part of *Homo sapiens*, for example, is to be part of the genealogical nexus that connects all other organisms that are part of this species to a common ancestor or group of common ancestors. As the tree of life represents genealogical relations, one can also say that the genealogical approach defines species as chunks of the tree of life.

This broad characterization of the genealogical approach requires at least three further specifications. First, there are relations of ancestry and descent between entities on different levels of organization. The tree of life can be seen as a tree of species, a tree of genes, a tree of organisms, and a tree of populations. Assuming that these are all biologically important levels of organization, it is not *prima facie* clear which of these histories should act as the basis for the genealogical approach to classification. However, two of these options can be eliminated immediately. As thinking of the tree of life as a tree of species assumes that we already have a way of identifying species, it is not helpful if we want to use the tree of life as a guide to species delimitation. A second viable simplification is that genealogical relations between different populations can be reduced to the genealogical relations between organisms or genes. However, this still leaves us with two levels of organization, the organism-level and the gene-level, which cannot straightforwardly be reduced to each other (Haber, 2012; Velasco, 2010).

Second, genealogical approaches need to specify how to understand the notion of unique common history as a criterion for grouping organisms. This notion is often cashed out in terms of monophyly (e.g. Mishler & Brandon, 1987; Mishler & Donoghue, 1982; Mishler & Theriot, 2000). Roughly speaking, a group is monophyletic if it contains an ancestral population and all and only the descendants of this population. However, while monophyly is suitable to describe taxa above the level of species, it is unclear how it applies to groups on the species level. This is because the genealogical relations of sexually reproducing organisms do not form strictly nested, non-overlapping groups (Velasco, 2009).⁵ I am, for example, part of the monophyletic

⁵ Another problem for applying monophyly to groups on the species level is that it makes it impossible for a species to be ancestral to another species. As soon as a new species branches off from the ancestral species, the latter is no longer monophyletic (Reydon, 2006; Rieppel, 2010; Sober, 2000). To apply the notion of monophyly

group consisting of my maternal grandmother and all and only her descendants, but also of the overlapping monophyletic group consisting of my paternal grandmother and all and only her descendants. Consequently, monophyly cannot serve as a basis of classification that organizes organisms into non-overlapping basic-level groups. In response to this worry, some authors have argued that unique common history should be understood in terms of exclusivity (Baum, 2009; Baum & Shaw, 1995; Velasco, 2009). An exclusive group consists of organisms that are more closely related to each other than to any organisms outside that group. This approach can be used both for sexually and asexually reproducing organisms, as both form non-overlapping nested exclusive groups. However, clarifying unique common history in terms of exclusivity is not entirely without problems either, because it can still be interpreted in various ways. Most importantly, exclusivity can be defined in terms of genetic relatedness and in terms of a most recent common ancestor not shared by organisms outside that group (Velasco, 2009, 2010).

The final issue that requires clarification concerns the ranking of exclusive or monophyletic groups as species (e.g. Ereshefsky, 1992, pp. 673–674; Mishler & Donoghue, 1982, p. 499; Velasco, 2009, p. 484). As the tree of life has a branching structure down to the species level, groups that are monophyletic or exclusive can be found on many nested levels. These levels are represented in Linnaean classification as different ranks, such as genus, family or order. It follows that genealogical approaches to species classification need a further criterion for deciding which genealogical groups should be recognized as species. Generally, genealogical approaches locate species on the boundary between the reticulate relations that characterize sexually interbreeding populations, and the non-reticulate relations on higher ranks. Thus, species are the least inclusive monophyletic or exclusive groups. However, there is no consensus about how precisely this is to be interpreted. Many authors rely on criteria for ranking that are not genealogical, such as diagnosability (Cracraft, 1983) or ‘group size, gap size, geological age, ecological and geographical criteria, degree of intersterility, tradition and possibly others’ (Mishler & Donoghue, 1982, p. 499). Baum (2009, p. 74) holds that ranking is inherently a subjective matter, requiring ‘semisubjective’ decisions. Mishler (1999), finally,

to species then, we would have to drop the requirement that monophyletic groups contain *all* the descendants of an ancestor.

believes that this subjectivity means that the species rank should be abolished, and classification of organisms should only be in terms of rank-free monophyletic groups.

Over the past decades, different genealogical approaches to species classification have been proposed (Baum, 2009; Baum & Donoghue, 1995; Baum & Shaw, 1995; Cracraft, 1983; Eldridge & Cracraft, 1980; Kornet, 1993; Kornet & McAllister, 2005; Mishler & Brandon, 1987; Nixon & Wheeler, 1990; Queiroz & Donoghue, 1988; Ridley, 1989). The differences between these views lie largely in the position they take with respect to the three issues discussed above, and in the way the approach is operationalized in taxonomic practice. I will not go into the differences between these accounts or their relative strengths and weaknesses here, but will discuss some of the relevant issues in the arguments for species pluralism in chapter 2.

1.2. The process approach

The process approach recognizes species on the basis of the evolutionary processes that create and maintain cohesion between the organisms of a species. David Baum (2009, p. 74) calls such species ‘functional’ because their cohesion allows them to act as single, unified entities in evolutionary processes. For this reason, they are also often called ‘units of evolution’ (e.g. Dupré, 2001; Reydon, 2005). These units of evolution play a role in processes like speciation, extinction, and possibly even species selection, and therefore also figure in evolutionary explanations of these phenomena.

There are a number of species concepts that adopt a process-definition of species, but at the same time emphasize that these species form evolutionary lineages that are extended over time and interconnected by relations of ancestry and descent. Instead of focusing on species as synchronously living groups of organisms that actively participate in evolutionary processes as a unit, they focus on the pattern that is the passive result of these processes over time. Thus, one could argue that these are genealogically defined groups. However, while these species concepts regard monophyly or exclusivity as important signals of the evolutionary unity of a group of organisms, these properties are not taken to define these groups as species.⁶ What defines such species at any point in time is their functional cohesion. Thus, these evolutionary lineage-based species concepts are, in the first place, part of the process

⁶ In that sense, these lineage-based species concepts (such as the Evolutionary Species Concept) need not be limited to lineages based on genetic inheritance. In line with Kendig (2014, p. 169), these concepts could in theory also include epigenetic, ecological, cultural and behavioural inheritance systems.

approach. The main examples of concepts that emphasize this temporal dimensionality of process-species are various versions of the Evolutionary Species Concept (Mayden, 1999; Simpson, 1945; Wiley, 1978), and closely related concepts such as De Queiroz' (1999, 2005a, 2005b, 2007) General Lineage Concept.

Regardless of whether they are defined as active participants in evolutionary processes or as the lineages resulting from these processes, the process-approach defines species on the basis of cohesion between the organisms of which the species consists. This notion of cohesion requires further clarification. Because different organisms of a species are usually materially discrete, they are not literally cohesive in the sense of being directly connected.⁷ However, the brief discussion in the previous paragraph does suggest a sense in which process-based species are metaphorically cohesive. The crucial feature of the cohesion of process-based species is that it allows them to act as a unit in evolutionary processes. In other words, they are cohesive in the sense that they are affected by the same evolutionary pressures in the same way, and respond to these pressures in a similar way. Thus, process-based species are cohesive 'with respect to evolutionary pressures' (Barker, 2007, p. 655). Another popular way of stating this is to say that the organisms of a process-based species share a unique evolutionary trajectory. This formulation ties nicely into the temporal dimensionality of some process-species as discussed in the previous paragraph.

There are many mechanisms and processes that contribute to the evolutionary cohesion that causes the organisms of a species to share an evolutionary trajectory, including gene-flow, developmental homeostasis and natural selection. Closer consideration of these mechanisms suggests that there are two general forms that this cohesion can take. Barker and Velasco (2010) and Barker and Wilson (2010) call these two forms of evolutionary cohesion 'integrative cohesion' and 'response cohesion', and similar distinctions can be found in Ereshefsky (1991), Stegenga (2016) and Mishler and Brandon (1987).

A group of organisms has integrative cohesion if there are strong causal connections between them that promote or maintain integration into the whole. Consequently, integrative cohesion usually requires spatial contiguity that allows for direct interaction between the organisms of a species. This does not mean that in order to be integratively cohesive, all

⁷ There are of course exceptions to this, such as slime molds.

organisms of a species need to be connected by direct causal interaction. As Stegenga (2016) points out, such a condition would be too restrictive even for an account of populations, and is *a fortiori* very unlikely to be met by any real species. Instead, integrative cohesion requires that all organisms of a species are indirectly connected by causal interaction. The chief mechanism thought to be responsible for integrative cohesion in species is gene-flow. The idea here is that organisms share an evolutionary fate because they are genotypically and phenotypically similar to a high degree, and this similarity is created and maintained by gene-flow between populations and, more generally, the transmission of genes from parents to offspring (Barker, 2007; Barker & Wilson, 2010). In other words, because genes are distributed throughout the species, different members of the species tend to be similar in many respects. Because they are similar, they tend to be impacted by selection in the same way, and because members of the same species continue exchanging genes, they tend to be impacted by drift in the same way. Thus, gene-exchange keeps the organisms of a species on the same evolutionary path.

The process-based view on species which explains cohesion by referring to gene-flow has been highly influential in the history of biology, most importantly in the form of Ernst Mayr's Biological Species Concept (BSC) (Coyne & Orr, 2004; Mayr, 1982). Mayr defines species as interbreeding groups of organisms that are reproductively isolated from other such groups. However, over the last decades much criticism has been raised against the BSC, even by proponents of process-based species concepts. One particularly significant problem of the BSC is that in many of the groups that biologists recognize as species, gene-flow seems neither necessary nor sufficient for creating cohesion between organisms. Barker (2007) gives a wide range of examples of sexually reproducing animal and plant species which consist of populations that never exchange genes, but clearly continue to be on a shared evolutionary trajectory. Similarly, he shows that there are many cases where intensive interbreeding between organisms fails to create and maintain both cohesion and shared evolutionary fate. These problems suggest that the second type of cohesion, called 'response cohesion', is needed to account for process-based species. A group of organisms has response cohesion when they are subject to the same or similar causal processes and respond to these processes in a similar way. Response cohesion is weaker than integrative cohesion because it does not require direct or indirect interaction between the organisms. It only requires that different

organisms of a species are affected by common causes in such a way that they have a shared evolutionary trajectory. On this view then, the shared evolutionary trajectory of process-based species is caused by mechanisms such as shared selection pressures, ecological regimes and developmental homeostasis. While gene-flow was traditionally considered the only or most important source of intraspecific cohesion (Barker (2007) calls this ‘the View’), it is now widely accepted that responsive cohesion plays a major role in the cohesion of evolutionary units. One important example of a concept that explicitly refers to such response cohesion is Van Valen’s Ecological Species Concept. This concept defines species as lineages of organisms subject to the same ecological regime (Valen, 1976). Thus, it emphasizes the fact that unity within groups is often the result of common selection pressures and occupation of a common niche, regardless of whether the organisms are interacting with each other.

Barker and Wilson (2010, p. 73) point out that there is no reason to assume that a single mechanism should be held responsible for creating and maintaining cohesion in process-based species. Moreover, there is also no reason to believe that these mechanisms should be precisely the same in all species. Thus, they argue, any realistic process-based view should characterize species both with interactive and response cohesion, depending on the particular mechanisms that form the causal basis of the species under investigation.⁸ Many process-based species do indeed incorporate various cohesion-creating mechanisms in their species definition. The ESC, for example, simply defines species as lineages with a unique shared evolutionary fate, without specifying the mechanisms responsible for this (De Queiroz, 2005a; Wiley, 1978). I will not discuss these species concepts more extensively here, but will come back to some of these issues in the arguments for radical species pluralism in chapter 2.

1.3. The similarity approach

The final major approach to species is the similarity approach, which defines species in terms of similarity between the organisms that form the species. The similarity in question can be of any character, be it morphological, behavioral, karyological, or genetic. While now largely abandoned in the theoretical literature on species (e.g. Ereshefsky, 2001; Ghiselin, 1999, p. 499; Hull, 1997, p. 367; Sober, 2000), similarity-based concepts played an important role throughout the history of systematics. One important example is the Morphological Species

⁸ This fits well with a ‘Homeostatic Property Cluster’ view on species, which I discuss in section 3.2 of this chapter, and which Barker and Wilson also explicitly refer to.

Concept, which Mallet (2001) attributes to Darwin. This species concept defines species on the basis of morphological similarity between organisms of the same species and a morphological gap between distinct species. A second important similarity-based concept is the Phenetic Species Concept, which defines species in terms of overall similarity based on a multivariate statistical analysis of characters (Jardine & Sibson, 1968; Sneath & Sokal, 1973). Finally, a more recent similarity-based species concept (closely related to the Phenetic Species Concept), is Mallet's (1995) Genotypic Cluster Concept, which defines species on the basis of genetic and morphological gaps between populations.

There is broad consensus among philosophers and biologists that these similarity-based approaches are not theoretically suitable for most contemporary evolutionary research. This does not mean, however, that these approaches do not play an important and epistemically fruitful role in practice. Mallet argues that his Genotypic Cluster Concept is often applied in models of sympatric speciation and as an operationalization of the Biological Species Concept. On a similar line, Lewens (2012a, p. 159) points out that despite wide rejection of the Phenetic Species concept, 'phenetic approaches to taxonomy are still alive and well in biological practice'. Rosselló-Mora and Amann (2001), and Ereshefsky (2010b), finally, point out that microbial taxonomy often relies at least partly on a similarity approach, for example in research that adopts the popular Phylo-Phenetic species concept.

2. Natural Kinds and Individuals

The previous section provided an overview of the various kinds of groups that biologists recognise as species. Broadly speaking, the goal of this thesis will be to investigate what factors play a role in the individuation of these groups, and more specifically, whether, and in what role, values are among them. Before we turn to the role of values, it is crucial to take a closer look at how species classification is determined by the world. This is important because it is clear that regardless of the role that values play, species classifications should be dependent on what the world is like. This is suggested by the fact that the approaches discussed above all define species in terms of features of the organic world, and by the fact that careful observation and data-collection form an important part of all taxonomic research. Moreover, as will be discussed below, this connection between species classifications and the world is also needed as an explanation for the epistemic usefulness of species classifications. In short then, any potential role for values in species classification can only be understood in

combination with an understanding of how species classifications are dependent on the structure of the organic world.

The role of the world in species classification has been the subject of much debate in the philosophy of biology. While most authors agree that species taxa are real in some sense, i.e. that species classifications represent real features of the organic world, there is no consensus about what kind of features they represent. More precisely, there is a longstanding debate about whether species are natural kinds or individuals (Ghiselin, 1974; M. H. Haber, 2016; Hull, 1978, 1981; Kitts & Kitts, 1979; Reydon, 2003; Ruse, 1987). This section will argue that for the theoretical purposes of this thesis it does not matter whether the discussion is framed in terms of natural kinds or individuals, and that for practical reasons it is best here to adopt the language of natural kinds.

The view that species are individuals is easiest to explain by referring to a paradigmatic example of an individual like a human organism. A human organism has a particular beginning and end in space and time, consists of distinct parts such as cells or organs, is cohesive and integrated, and can undergo unified change such as growth and aging. Proponents of the view that species are individuals claim that species have precisely these same characteristics: they are spatiotemporally located and restricted, consist of different parts (organisms) that are cohesive and integrated through evolutionary processes, and can change as a unit through evolution because of this (Ghiselin, 1974; Hull, 1978).

Other authors argue that species are natural kinds. Natural kinds are traditionally defined as groups of particulars that share the same essence, which is a set of core properties responsible for the other shared properties of a kind, and jointly necessary and sufficient for membership of the kind. As various authors have pointed out, this essentialist view of kinds does not fit well with current biological knowledge about species (Ghiselin, 1974; Hull, 1965a, 1965b, 1976; Sober, 1980). Most importantly, there is no set of genetic or phenotypic properties that all and only the members of a species share. Evolutionary theory also suggests that even if there were such properties, they could change at any time in any organism in the next generations through mutation. These deviations from the essentialist view are widely observed in the organic world, and are a direct consequence of the theory of evolution. It is no surprise then that after Ghiselin and Hull popularized the species-as-individuals thesis, the consensus quickly grew that species are individuals.

However, several authors have since proposed alternative views on natural kinds that can accommodate the intraspecific variation, changeability and historicity of species. Most importantly, in section three of this chapter I will discuss various recent views on natural kinds which hold that genealogical relations and the evolutionary processes cited by the process-approach can be seen as mechanisms that cause the stable clustering of properties in species. These views (e.g. Boyd, 1999; R. A. Wilson, Barker, & Brigandt, 2007) claim that kinds are characterized by a set of properties and mechanisms, but that none of these is necessary for kind membership. Thus, they are compatible with the intraspecific variation and changeability of species, and with the fact that they are spatiotemporally restricted. Moreover, these views put the causal relations cited by the process-approach and the genealogical approach at the centre of their definition of natural kinds. In that sense, they fit very well with the species recognised by taxonomists in practice.

Various authors have argued that if we adopt these recent views on natural kinds, both the species-as-individuals and the species-as-natural kinds view are plausible (Boyd, 1999; Brigandt, 2009; Dupré, 1993; see also LaPorte, 2004). In line with this, I argue that for the purposes of this thesis, it does not matter whether the relation between species classification and the world is articulated in terms of natural kinds or individuals. As will become clear in chapter 2, my arguments will focus on the two approaches that classify species on the basis of causal relations between organisms, i.e. the process approach and the genealogical approach. These two approaches fit very well both with a view of species as natural kinds and with a view of species as individuals. This is because both views define belonging to a species in terms of having particular causal relations. Thus, successful individuation of species, be it as kinds or as individuals, depends on classifications somehow tracking these chunks of the causal structure of the world. The question about the role of values in species classification then becomes whether these causal structures fully determine classification, or whether (and how) values play a role too. In order to answer this question, it does not matter whether the chunks of this causal structure on which individuation depends are best called 'individuals' or 'kinds'.⁹ Either way, we need to clarify how genealogical relations or causal cohesion individuate groups of organisms as species.

⁹ For a very similar point, see Boyd (1999, p. 163), who points out that 'successful induction and explanation depend just as much on the accommodation of our individuating practices for individuals to relevant causal structures as on the accommodation of those practices for kinds.'

As it does not matter for my purposes whether species are individuals, kinds or both, it is, for reasons of simplicity, better to discuss the arguments either in terms of natural kinds or individuals, rather than making equivalent points for both. Because there is a long philosophical tradition of discussing scientific classification in terms of natural kinds, I will adopt this language in the remainder of this thesis. Not only is the literature on natural kinds more extensive than that on individuals, it also engages more with questions about values in classification, which is the main topic of this thesis.

3. Natural Kinds

The previous section argued that for the purposes of this thesis it is best to articulate the relation between the world and classifications in terms of natural kinds. One might argue here that by framing my discussion in terms of natural kinds, I risk inheriting the wide range of philosophical problems and debates that have accumulated in the debates about natural kinds over the past centuries. Hacking (2007; see also Ludwig, forthcoming) argues that this is the case for any contemporary discussion of natural kinds. He claims that discussions of natural kinds are so heavily 'philosophy-laden', i.e. riddled with metaphysical and semantic assumptions, that they only obscure the philosophical questions and scientific practices we are trying to understand. Hacking speaks of the 'scholastic dawn' of the natural kind concept, where the adjective 'scholastic' refers to an 'inbred set of degenerating problems that have increasingly little to do with the issues that arise in a larger context' (Hacking, 2007, p229). Not surprisingly then, Hacking holds that philosophical and scientific debates about kinds would be better off if all mention of natural kinds was avoided.

It is hard to deny some truth in Hacking's points. The debate about natural kinds is entangled with a range of philosophical, mostly metaphysical, problems that at times are far removed from the concerns of scientists engaged in classification. It is also true that the notion of natural kinds seems to have had little success in resolving these issues. However, this is not necessarily a reason for abandoning the concept of natural kinds entirely. Even if some parts of the debate on natural kinds can be characterised as scholastic, it might well be that the notion can still serve a philosophical purpose (MacLeod & Reydon, 2013). Indeed, I believe that when suitably understood, the notion of natural kind can be useful.

What ‘suitably understood’ and ‘useful’ mean here, is of course dependent on the purposes of this thesis. As this thesis aims to investigate species classifications as they are constructed by taxonomists, a useful view on natural kinds should be able to account for these classifications. The previous sections have briefly discussed various relevant properties of these classifications. More specifically, section 1 described the principles on which these classifications are based in taxonomic practice. Section 2 investigated how we can think about the relation between species classifications and the world. To this we can add, as already briefly noted above, that species classifications are epistemically useful. In the context of natural kinds, philosophers often specify this epistemic usefulness in terms of projectibility (Khalidi, 2015; Slater, 2014; Boyd, 1991). Even though I will articulate my points in terms of natural kinds, I do not mean to imply that the epistemic usefulness of species classification is limited to projectibility, or, conversely, that all species can serve as the basis for generalizations (MacLeod, 2013; Reydon, 2006). As MacLeod (2013), Magnus (2012, pp. 160–165) and Love (2009) point out, scientific classification serve many other varied epistemic aims, including understanding, prediction, simplicity of modelling, explanation, and information storage and retrieval among others.

Assuming that species are natural kinds, it follows that any suitable notion of natural kinds should be able to account for these features. This gives us three criteria to evaluate views on natural kinds as suitable for the philosophical purposes of this thesis. First, any suitable view of natural kinds should recognize as natural kinds most of the species recognised by taxonomists (i.e. the three views described in the first section). In addition, such a view should see genealogical relations, cohesion relations and/or similarity not merely as accidental features of species, but as the features that individuate a group of organisms as a species. Second, any suitable view on natural kinds should incorporate the idea that the world determines at least partially which classifications are legitimate. That is, a good view on natural kinds should explicate the relation between the world and good species classifications. Third, any suitable view of natural kinds should be able to account for the epistemic usefulness of species classifications. This third condition is closely related to the second condition, because it is the connection between natural classification and the world that is usually invoked to explain the epistemic success of these classifications.

In short, then, any acceptable view on natural kinds should be able to account for existing good scientific kinds and the way these are epistemically useful, and it should do this, at least partially, by referring to the relation between natural kinds and the world. The next three subsections will discuss three broad classes of views on natural kinds, and evaluate the extent to which they meet these three criteria.¹⁰ In line with much of the recent literature, I will argue (section 3.1) that the first view, i.e. essentialism, successfully provides an account of the epistemic success of natural kinds and how they are connected to the world, but fails to capture the kinds effectively recognised by taxonomists. I then argue that the other two views, which I will call the causal view (section 3.2) and the bare property cluster view (section 3.3), both provide views that meet all three criteria, albeit to somewhat different extents.

3.1. Essentialism

Essentialists hold that a group of particulars form a natural kind if they share an essence or a set of essential properties. These essential properties are taken to specify the nature of the members of the kind. Even though different authors characterize these essential properties in different ways, there are two criteria that recur in nearly all essentialist accounts. First, essential properties form the necessary and sufficient conditions for a particular to be member of a kind (e.g. Ellis 2001, Soames 2002). That is, a particular cannot be of a certain kind unless it has all essential properties, and any particular with all essential properties of a kind is a member of that kind. Second, the essential properties of a kind should be responsible for the non-essential properties typically associated with the kind. Because of this, the essential properties can be used to explain and predict these non-essential properties. Ereshefsky (2016) gives the example of the natural kind ‘gold’ that has a certain atomic structure as its essence. This atomic structure can be used to explain and predict non-essential properties of the kind, like its melting point.

Besides these two crucial characterizations of essential properties, a number of other criteria are often associated with essentialist natural kinds. First, essential properties are often claimed to be intrinsic properties of the members of a kind. This means that the properties that make a particular member of a kind are independent from that individual’s relations to anything else. For example, the weight of a cube of gold is an extrinsic property, because it

¹⁰ See Reydon (2009a), Magnus (2015) and Ludwig (forthcoming) for similar and more detailed overviews of different views on natural kinds.

depends on the local gravitational acceleration. Its atomic number, on the other hand, is independent from any contextual factors, and therefore is an intrinsic property. Second, essential properties are often assumed to be microstructural properties, i.e. properties underlying the various macrolevel properties of the members of the kind. The malleability of gold, for example, is a macrolevel property that can be explained by its microstructure. Finally, essential properties are often argued to be modally necessary. That is, the essential properties are properties the kind has in any possible world, and any member of the kind possesses these properties in any possible world in which it exists as a member of that kind. Unlike the two criteria discussed in the previous paragraph, not all essentialists accept intrinsicity, modal necessity and microstructure as necessary criteria for essential properties. This is because many well-established scientific kinds meet only some or even none of these criteria. Accepting these requirements for natural kinds then excludes many scientific kinds from natural kind status (see Khalidi (2013, p. 12) and Magnus (2012, pp. 13–50) for further discussion of these criteria).

The appeal of the essentialist account lies in the way it meets criterion 2 (relation to the world) and, particularly, criterion 3 (explaining epistemic success). Essentialists hold that essences are part of the fundamental structure of the world. Classification based on these essences thus reflects this fundamental structure. This connection with the world is then employed to explain the epistemic usefulness of kinds. More specifically, because the essential and non-essential properties of natural kinds are ontologically connected, it is epistemologically legitimate to make inferences based on these essential properties. Moreover, as every member of a kind has all essential properties, these inferences are guaranteed to be valid for all members of the kind.

While essentialism scores well by criteria 2 and 3, it does far worse by criterion 1 (fit with scientific kinds). The main problem with essentialism is that its criteria for natural kindness are very strict. Many of the kinds used by scientists do not have a set of properties the possession of which is necessary and sufficient for kind membership. And even if such kinds exist in some areas of science (which is far from obvious), it is now generally accepted that most if not all biological species fail to meet the strict requirements of essentialism. This is suggested by empirical research, which strongly suggests that both phenotypic and genotypic properties fail as candidates for essences; organisms share many properties across species-boundaries, and

many species are characterized by substantial intraspecific variation. The failure of essentialism is also implied by evolutionary theory, which holds that the properties of species can change over time, and that intraspecific variation plays a crucial role in how species come into existence (Sober, 1980). It is no surprise then that species essentialism is generally taken to be untenable in post-Darwinian systematics. One might argue that at least on similarity-based approaches, species are defined purely on the basis of shared properties, which thus might qualify for essentialism. However, these approaches usually do not consider any particular set of properties necessary and sufficient for membership. Moreover, these approaches do not make claims about causal relations between various properties, and thus fail to meet the requirement that the essential properties explain the non-essential properties of a kind.

These problems are further exacerbated when we take the requirements of intrinsicity, microstructure and modal necessity into account. The first section of this chapter argued that on the main approaches to species classification, species are defined in terms of genealogical relations, causal relations or similarity relations between organisms. This means that the species recognised by taxonomists are not characterized by microstructural or intrinsic properties. While most similarity-based species are usually defined in terms of intrinsic properties, these are not always microstructural properties. And even if present-day similarity-based methods often use genetic properties, this is usually motivated by practical considerations, and not by the belief that these genetic properties are more important than macrolevel properties. The modality-requirement is equally problematic. Given the path-dependent and contingent nature of evolutionary trajectories (Beatty, 1995), it seems unlikely that there are any properties that any particular species would have in any possible world in which it exists.

Because essentialism cannot accommodate the kinds recognized by taxonomists, it fails to meet the requirements I set out above. I will therefore not discuss it further in this thesis. This is in agreement with almost all contemporary literature on species. The authors who do argue for species essentialism adopt a notion of 'essence' that is notably weaker than the traditional notion discussed here. Devitt (2008, p. 371), most importantly, endorses intrinsic essentialism but admits that there may not be a 'tiger gene', i.e. a 'neat and tidy' essence perfectly instantiated in each member of the tiger-kind. His argument for intrinsic essentialism relies on

the fact that robust explanations of species' traits require citing the microstructural properties of those species. However, as Lewens (2012b) and Ereshefsky (2010c, 2016) point out, this renders the notion of essence so weak and permissive that 'it barely merits the name' (Lewens, 2012b, p. 754). Other authors (e.g. Griffiths, 1999; Okasha, 2002; R. A. Wilson, 1999) have defined essences in relational or historical terms and also weakened the requirement of necessary/sufficient properties and explanatory force. While I do not object to these accounts, I believe they are better categorized as causal accounts of natural kinds, which I discuss in the next section.

3.2. Causal kinds

The essentialist account of natural kinds discussed in the previous section starts from the metaphysics of natural kinds. It first specifies what kind of things natural kinds are, and then utilizes this characterization to evaluate scientific kinds and explain their success. Given this starting point, it is not surprising that the essentialist account fails to recognise many scientific kinds as natural, as explained above for the case of species. In response to the failure of the metaphysics-first essentialist views, various authors defend accounts of natural kinds that start from an epistemological perspective instead (Boyd, 1991, 2000; Ereshefsky & Reydon, 2014; Khalidi, 2013, 2015; Magnus, 2012; Slater, 2014). Broadly speaking, these views define kinds as those categories that allow us to make reliable inferences and explanations (and the other epistemic roles mentioned above). Starting from the ways kinds figure in science, these accounts then propose metaphysical views that fit with and explain these epistemic roles.¹¹

The relation between the metaphysics of natural kinds and their epistemic roles is most clearly captured by Boyd's (1999, pp. 146–156, 2000, p. 55) accommodation thesis. Rather than assuming, like essentialist views, that natural kinds are part of the furniture of the world that science sets out to discover, the epistemic views see the construction of scientific classifications as the accommodation of classificatory practices to the structure of the world. Thus, natural kinds are seen as the combination of the 'workmanship of women and men' and the structure of the world (Boyd, 1999, p. 175). The contribution of the world lies in the particular structures of the world that scientific language is accommodated to. The

¹¹ Note that not all 'epistemic' views on natural kinds have a metaphysical component (see e.g. Ereshefsky & Reydon, 2014; Franklin-Hall, 2015; Magnus, 2012). As I am interested in the relation between the world and classification here, I will not discuss these purely epistemic accounts in detail.

contribution of women and men lies in constructing categories that capture the particular structures of the world that are relevant to the explanatory and inferential practices of a particular scientific domain. Reydon (2015a, pp. 60–63) calls this the ‘zooming-in model’ of natural kinds, as the aim of classification is seen as zooming in on the parts of the world that are relevant for explaining the phenomena that we want to explain.

This subsection discusses a particularly prominent subclass of these ‘epistemic accounts’, namely, those that hold that scientific classification ought to be accommodated to the *causal* structure of the world. These ‘causal views’ on the metaphysics of kinds are generally characterized by two features. First, natural kinds are characterized as clusters of properties that repeatedly co-occur. Unlike essential properties, these clusters do not have to be shared perfectly by all members of a natural kind. This means that two members of the same causal kind can have widely different properties, and might even, as Wilson et al. (2007) point out, have no properties in common at all. The second feature that characterizes causal kinds is a causal basis that explains the stable clustering of these properties. Like the property cluster, this causal basis can vary over time and may also differ between different members of the same kind.

It is clear that, despite the different starting points of the causal views and essentialism, they are importantly similar. More particularly, the causal basis plays a role similar to that of essences in the essentialist account. First, just like essences, the causal basis distinguishes arbitrary clusters of properties from natural kinds: property clusters are natural kinds because their clustering is the result of an underlying causal basis. To pick out natural kinds then, we should not just consider shared properties but rather the causes responsible for this clustering. In other words, proponents of causal views claim that natural classifications are determined and constrained by the causal structure of the world. Second, the causal basis is responsible for the surface properties of a kind, and can consequently serve as an explanation for these properties and their clustering. This way, the causal views offer a strong explanation for the epistemic success of natural kinds: natural kinds are epistemically useful because there is a causal basis for the stable clustering of their properties. It is this causally sustained stability that allows us to use natural kinds for inferences and generalizations.

Despite these general similarities between the causal views and essentialism, there are three crucial differences between the two. First, while traditional essentialism often claims that

essential properties are microstructural, intrinsic properties, the causal views explicitly allow for the property cluster and causal basis of a kind to consist of extrinsic and macrolevel properties. This is particularly important for the discussion here, as species are commonly individuated on the basis of relations between organisms rather than their intrinsic, microstructural properties. Second, while traditional essences are generally understood to be timeless and unchanging, the causal views assume that both the causal basis of kinds and their properties can vary over time. Finally, and most importantly, the causal views allow us to recognise particulars as a member of a natural kind even if these particulars do not have all the properties and mechanisms generally associated with the kind. Unlike essentialism then, the causal views deny that there are any properties necessary for kind membership. By loosening the stringent essentialist criteria for natural kind status, the causal views can account for many successful scientific kinds that are not natural according to essentialist views.

I have so far consistently referred to the causal views in plural rather than singular form. The reason for this is that there are multiple distinct views that fit the general description given in the previous paragraphs. These views mostly differ in the way they characterize the causal basis of natural kinds. At least two different views should be distinguished here: Boyd's Homeostatic Property Cluster (HPC) view, and Khalidi's simple causal view.

Boyd's Homeostatic Property-Cluster (HPC) view of natural kinds was the first and still is the most prominent epistemic causal view of natural kinds.¹² According to Boyd, natural kinds are property-clusters stabilized by a set of homeostatic mechanisms, which form the causal basis of the kind. Mechanisms are usually defined as 'entities and activities organized together such that they do something' (Craver, 2009, p. 582). The term 'homeostatic' is borrowed from Bernard, who uses it to refer to the maintenance of an internal state in face of environmental changes. Extending this meaning, homeostatic mechanisms are mechanisms that keep a cluster of properties stable over a range of circumstances. Boyd refers to biological species as paradigmatic examples of such HPC kinds; species share a long range of properties, and both

¹² Of course, other authors have connected natural kinds to epistemological questions (e.g. Goodman, 1975) and causality (e.g. Locke). However, the recent trend of epistemic views and re-popularization of natural kinds in philosophy was without any doubt started by Boyd's HPC view.

this similarity and the stability of similarity are due to mechanisms such as interbreeding, shared selection pressures and common developmental mechanisms.

Various authors (Ereshefsky & Reydon, 2014; Khalidi, 2013, 2015; Slater, 2014) argue that the HPC view is overly restrictive because there are many causally sustained kinds that are not associated with homeostatic mechanisms. Chemical kinds, perhaps the paradigmatic example of natural kinds, can serve as an example. While there are good reasons to see gold as a natural kind, there does not seem to be a way to explain the stability of that property-cluster by reference to a homeostatic mechanism (Khalidi, 2015, p. 8). More importantly here, the causal basis of species does not seem to consist of homeostatic mechanisms. For example, it is not clear how the genealogical relations invoked by the genealogical account can be seen as a mechanism. Common descent results in similarity between members of a species because of what may best be described as phylogenetic inertia or causal isolation, i.e. ‘the absence of potentially perturbing causal pathways from the here and now to the there and then’ (Slater, 2014, p. 392). Such causal isolation does not fit the commonly accepted view of mechanisms. Moreover, even if the causal bases could be described in terms of mechanisms, it is not clear why these should always be described as homeostatic. Selection, for example, can act in a disruptive fashion as well as in a homeostatic way. Yet, disruptive selection is in some cases plausibly invoked as a mechanism for explaining properties and their distribution in species.

Khalidi (2013, 2015) argues that in order to capture scientific kinds not stabilized by homeostatic mechanisms, we should adopt a less restrictive view of the causal basis of natural kinds:

When it comes to natural kinds, causal relations among properties (or more properly, property instances) constitute the ontological ground for the projectibility of the corresponding predicates. But there need be no single causal mechanism that leads these properties to be co-instantiated nor need there be any kind of feedback process that ensures that these properties do not depart from an equilibrium state of co-instantiation. So Boyd’s account has to be loosened in such a way as to retain the emphasis on causality without the mechanism or the homeostasis. (Khalidi, 2015, p. 8)

Briefly summarized, Khalidi defines natural kinds as a cluster of core causal properties that gives rise to a cluster of derivative properties. This causal basis need not solely consist of mechanisms, and need not keep the natural kind in homeostasis. It should be noted here that

it is not always clear to what extent Boyd's and Khalidi's views are actually different. Boyd regularly suggests that his term 'homeostatic mechanism' should not be interpreted literally, but metaphorically. This also seems to be how most authors in the field interpret the term. Ereshefsky and Reydon (2014, p. 971), for example, write that 'one should not read too much into the term "mechanism" in "homeostatic mechanism," however. [...]. A homeostatic mechanism can be anything that causes (in the broadest sense of the term) a repeated clustering of properties'. Khalidi (2015, p. 8) also points out that his simple causal view converges with such a loose interpretation of the HPC view.

The reason why it is hard to distinguish between these views is in part because both Boyd and Khalidi are not very clear in their characterization of the causal basis of natural kinds. Boyd uses the metaphor of homeostasis, but does not provide a clear non-metaphorical description of the causal basis of kinds. Khalidi argues that Boyd's views, when interpreted narrowly, are too restrictive, and claims that kinds are natural as long as they are 'nodes' of the causal nexus. As he does not clarify his notion of causation or what a node of the causal nexus is precisely, this does little to clarify the causal view on natural kinds (see Reydon, 2015b). As my arguments below are not dependent on any particular interpretation of the causal basis of natural kinds, I will not distinguish between Khalidi's and Boyd's views in the remainder of this thesis, and simply refer to them as the 'causal views'. In line with this, I will use the terms 'mechanism' and 'causal basis' in a broad, loose sense to refer to the causal factors that are responsible for the clustering of properties.

It is clear that causal views on natural kinds fit better than essentialism with the kinds recognized in science. This is because the causal views, unlike essentialism, allow members of the same kind to differ to some extent both in properties and causal basis. Causal views on kinds also fit well with the genealogical and process approach to species classification. The process-approach defines species in terms of current or past causal relations between organisms, such as interbreeding or common selection pressures. The genealogical approach defines species in terms of the causal history of the members of the species, and more particularly their genealogical relations.¹³ These causal relations are then invoked to explain the long range of properties that the members of a species have in common.

¹³ One could argue that the genealogical approach does not refer to causal mechanisms or a general causal basis, but to the particular causal history of a group of organisms. However, as Khalidi (2015, p. 17) points out,

At the same time, causal views on natural kinds meet criteria 2 (explicate the relation between world and classification) and 3 (explain the epistemic success of natural kinds). According to proponents of the causal views, the difference between natural and non-natural categories lies in the fact that the former reflects the causal structure of the world, while the latter does not. This relation with the causal structure of the world also provides an explanation for the epistemic success of natural kinds. Natural kinds are useful for investigating the world because they track the causal structure of the world. We can expect the set of properties associated with a kind to cluster in a stable manner because there is a causal basis underlying this clustering.

3.3. Bare property cluster kinds

Various authors (Häggqvist, 2005; Slater, 2014; see also Hacking, 1991, pp. 119–120) have argued that it is unclear why natural kind status should be restricted to property clusters with a causal basis. Assuming the epistemological approach to natural kinds, what matters in the first place is whether property-clusters are sufficiently stable to support our epistemic practices. Whether this stability is the result of a causal mechanism is not of primary importance to this epistemological outlook. Lipton (1996, p. 439) states a similar view when he comments on the HPC-view that ‘essences are supposed to hold together observable properties in stable clusters, but it is not made clear why this should make for a more inductively knowable world than one where that stability is a brute fact’. This objection leads to a Bare Property Cluster (BPC) view, which holds that natural kinds are simply property clusters that are stable across a range of circumstances. According to this view, the world just happens to be such that some properties co-occur with different degrees of stability, which allows us to make reliable inferences and generalizations. Chakravartty (2011) expresses this as the ‘sociability’ of properties. Such sociability is all that is needed for explaining the epistemic success of natural kinds. Even if, in many cases, this sociability has a causal basis, there is no reason to require that this is always the case.

The main strength of the BPC view lies in its inclusiveness. Various authors have argued that there are many scientific kinds that lack a causal basis. While these do not qualify as natural

‘etiological individuation is also a matter of discerning the causal structure of the world’. This is not to say, of course, that there are no interesting differences between etiological kinds and kinds individuated on the basis of a synchronously existing set of causal properties (see Magnus (2012, 2014) for a discussion of the differences between these ‘token-causal kinds’ and ‘type-causal kinds’).

kinds on the causal view, they do on the BPC view. This is significant here because there are some species, more specifically those recognised by similarity-based approaches, that do not straightforwardly meet the criteria of the causal view. The Phylo-Phenetic Species Concept, for example, individuates species on a partially non-causal basis, similar to the phenetic tradition (Ereshefsky, 2010b; Ereshefsky & Reydon, 2014; Rosselló-Mora & Amann, 2001). On a more general level, MacLeod (2013; see also Reydon, 2006) argues that there often is no easy identification between the causal basis of species, and the epistemic uses they are put to. That is, inferences and generalizations about organisms on the basis of species membership often do not straightforwardly follow from the mechanistic basis of the species, but rather require further background knowledge, and often apply to only parts of the species. It might be better, then, to think of species not in terms of their mechanistic basis, but rather simply as units of classification. Species concepts then simply 'classify organisms into groups that together contribute an enormous stock of general information' (MacLeod, 2015, p. 112). Such units of classification do not require a causal basis, but instead could simply be bare property clusters.

At the same time, the BPC view does not rule out causal kinds. Even if natural kind status does not require a causal basis, it might well be that many scientific kinds are characterized by such a basis. In this sense, the bare property cluster view can be seen as a highly abstract view on natural kinds that contains all previous accounts as subclasses: some property clusters might be sociable because they have essences, some might have homeostatic mechanisms, some might have another causal basis, and others, finally, might just be sociable as a brute fact about the world (Slater, 2014, p. 407). An additional advantage of this account is that it need not assume that all species are kinds of the same type. This is important here because the species recognized by biologists are of various types, including groups recognised on the basis of their causal history (genealogical view), current causal relations (process view) and similarity (pheneticism). An account like the bare property cluster view that can accommodate this heterogeneity is more likely to accommodate all successful scientific kinds.

One might think that another advantage of the BPC view is that it requires very few metaphysical assumptions. All that is assumed is the natural sociability of properties in the world. Unlike the other accounts, the bare property cluster view does not have to make further assumptions about the nature of this sociability, but still offers some explanation for

the epistemic success of natural kinds. However, this lack of metaphysics also forms the main weakness of the view. While the sociability between properties is part of the structure of the world, it offers very little explanation for the epistemic success of natural kinds. More precisely, while on this view the world still constrains classification through the sociability of properties, there is no explanation for why these properties are sociable in the first place.¹⁴ Thus, what the BPC gains in scope and inclusiveness, it loses in explanatory power.

Let us summarize this section on natural kinds then. I have argued that, given the goal of understanding scientific species classifications, we need an account of natural kinds that meets three criteria: it should be able to account for the kinds recognised by scientists; it should make sense of the distinction between arbitrary and non-arbitrary classifications by referring to the world; and it should be able to explain the epistemic success of scientific kinds. This section argued that the essentialist account meets the second and third criterion, but scores very badly on the first. The causal view on kinds, on the other hand, meets all three criteria. It provides an explanation for the success of natural kinds by referring to the world, and can account for most of the species recognised by taxonomists. The Bare Property Cluster View, finally, also meets all three criteria. As it can account for all kinds recognized by taxonomists, it scores very well on the first criterion. And even though the causal views clearly do better on the second and third criterion, the BPC still provides a form of explanation for the epistemic success of kinds by referring to the sociability of properties.

I conclude then that both the causal views and the BPC view provide suitable accounts of natural kinds for the purposes of this thesis. Which of these is preferable depends on which of the criteria is considered more important. If we want a view that can account for all scientific kinds, then the BPC is clearly preferable over the causal views. On the other hand, if we want a view that provides a strong explanation for the epistemic success of natural kinds, then the causal views are preferable over the Bare Property Cluster view.

¹⁴ It is worth noting here that on Khalidi's simple causal view, there is no requirement that various core causal properties are also causally connected. This means that this view too lacks an explanation for the sociability of properties, and converges with the BPC view. However, while not taking causal connections between core properties as a requirement for natural kind status, Khalidi (2015) does state that this is the case for most natural kinds.

4. Summary

This chapter has presented three broad approaches to species classification. These approaches differ with respect to the features of the world they rely on for individuating species. The genealogical approach individuates species on the basis of genealogical relations, the process-approach on the basis of cohesion-promoting mechanisms, and the similarity-approach on the basis of bare similarity relations. I then argued that it is most convenient here to think of these features of the world in terms of natural kinds. These natural kinds can either be understood as causally sustained property-clusters, or as bare property clusters.

The next chapter will argue that if we assume these views on species classification and natural kinds, then it follows that a radical form of species pluralism is true. This radical pluralism will then be used in the remaining chapters of the thesis to argue that values play a substantial role in species classification.

Chapter 2: Radical Species Pluralism

This chapter argues for a radical form of species pluralism. If we assume the principles of individuation for species discussed in chapter 1, then there are innumerable good classifications of organisms into species. The argument for this relies on the empirical claim that genealogical relations, cohesion-promoting processes, and similarity-relations are numerous, multi-faceted and complex. I argue that this empirical claim implies that there is no single classification that can represent all relevant genealogical relations, cohesion-promoting processes or similarity relations. Instead, there are very many good classifications that represent different aspects of these features of the organic world.

For some of the views on kinds and species presented in chapter 1, this argument is not controversial, and might even be called trivial. It is widely accepted that both the similarity approach to species classification and the BPC view on natural kinds lead to radical forms of pluralism. This is because these views recognise kinds purely on the basis of similarity relations. Assuming that similarity relations in the organic world are ubiquitous, it follows that there are innumerable equally legitimate classifications based on similarity relations and hence radical pluralism is true. Indeed, this is a criticism that has been raised frequently against these views. Magnus (2014), for example, writes about the BPC view on natural kinds that ‘a problem with this approach is that similarity is ubiquitous. Any arbitrary collection of things is similar in some respects and any bundle of properties is sociable to some degree’. Similarly, David Hull (1997, p. 360) writes about pheneticism that ‘any group of organisms can be classified in indefinitely many ways using various clustering techniques, and no reasons internal to these methods exist for choosing among these classifications’.

The argument for radical pluralism is less trivial if we restrict natural kinds to kinds with a causal basis, and, thus, species classification to the genealogical-view and the process-view. This is because these causal views only consider similarity relevant for classification if it has an appropriate causal basis. Moreover, such causal bases often explain many similarity-relations at the same time. It follows that these causal views only accept a subset of the similarity-based kinds as legitimate species. The idea then is that scientists can rely on the causal structure of

the world to regiment scientific classifications and distinguish important similarity-relations from irrelevant ones; while similarity may be ubiquitous, causally sustained similarity is rarer. Classificatory practices accommodated to this causal structure of the world are thus likely to recognise fewer classifications than those based on bare similarity.

These considerations seem to be reflected in taxonomic practice and the philosophical debate about species. While some approaches to species delimitation arguably rely on a similarity-approach, the lion's share of taxonomic work delimits species on the basis of the genealogical approach and the process approach. At first sight, the reliance on these two causal approaches seems to lead to monism or a moderate form of pluralism: While disagreement about species delimitation of particular taxa is common, these debates never concern the choice between more than a few different hypotheses. In line with this taxonomic practice, many philosophers subscribe to causal views on species classification, be it genealogical or process-based, and adopt a moderate form of species pluralism. One might think, then, that adopting the causal view on classification fends off the radical pluralism that inevitably threatens the similarity-based view.

This chapter argues that even if we restrict species classification to the genealogical and process-based views, and natural kinds to causal views, there are still innumerable ways of classifying organisms into species. In other words, I claim that if species are individuated on the basis of the causal structures referred to by the genealogical approach and process approach, then radical species pluralism is true. If this argument holds, Hull's criticism against the phenetic approach quoted above applies to all approaches to species classification. In that case, 'any group of organisms can be classified in indefinitely many ways' and 'no reasons internal to these methods exist for choosing among these classifications' (Hull, 1997, p. 360). Note that this claim does not imply that literally any group of organisms could be recognised as species by these approaches. Genealogical relations and evolutionary processes still rule out many classification schemes as wrong. However, and that is the point here, the constraints set by these features of the world are still compatible with very many distinct classifications.

The structure of the chapter is as follows. The first section introduces species pluralism, and presents a popular and plausible argument for a moderate form of pluralism. The second and third sections then argue that this argument for moderate pluralism slides off into an argument for radical pluralism. Section 2 sketches how this is the case for causal views of

natural kinds. Section 3 then applies the same argument to the main approaches to species classification. Section 4, finally, discusses other variants of species pluralism in the literature and compares them to the variant defended here. Before I turn to these arguments, it is worth noting that the focus on the causal views on species and natural kinds in this chapter should not be taken to imply that I commit to these causal views rather than to similarity-based views. Instead, I do this for reasons of space, and because radical pluralism is already widely considered plausible for similarity-based views. Moreover, given that kinds and species recognised by the causal views are a subset of those recognised by the similarity-based views, this does not limit the scope of the argument.

1. Moderate Species Pluralism

I will take pluralism about a particular domain to be the view that there are multiple legitimate scientific classifications of that domain. Thus, species pluralism is the view that there are multiple scientifically legitimate ways of classifying organisms into species. This claim can still be interpreted in two ways. One could claim that while each organism is part of only one species, not all these species are classified as such on the same grounds; some groups might be interbreeding species, while others are ecological species or phylogenetic species (see Mishler & Brandon, 1987; Mishler & Donoghue, 1982 for this kind of pluralism). On a stronger reading, species pluralism is the view that there are multiple, overlapping ways of organising the same collection of organisms into species. On such a view, two organisms that are conspecific in one legitimate classification need not be in another. The pluralism defended in this chapter is of the latter kind.¹⁵

This kind of species pluralism has been discussed, defended and criticized by many philosophers and biologists (e.g. De Queiroz, 2005a; Dupré, 1981, 1993, 1999, Ereshefsky, 1992, 1998, 2001, Hull, 1997, 1999; Kitcher, 1984b; Reydon, 2005; Stanford, 1995). Those who adopt it do so for various reasons. Some have epistemic motivations (e.g. A. Rosenberg, 1985,

¹⁵ Such species pluralism should also be distinguished from pluralism about the species category. The latter is the claim that the groups included in the species category are not all of the same kind. Even though pluralism about the species category often goes hand in hand with pluralism about classification, these two positions are logically independent. Ereshefsky (1992, 2001, 2010a), for example, argues for a combination of both types of ontological pluralism. Mishler and Brandon (1987), on the other hand, argue that there is a single best classification of organisms into species, but that these species are not all of the same type.

1994); they claim that because we do not and cannot know precisely what the world is like, all we can do is come up with multiple inaccurate accounts that shed light on parts of this world. Others argue for or against pluralism on the basis of pragmatic considerations. Hull (1987, p. 181) and Dupré (1999, p. 18) for example, defend species monism by pointing out the practical benefits of adopting a single species concept or a single legitimate classification. However, the majority of authors that defend species pluralism do this on ontological grounds (e.g. Dupré, 1981; Ereshefsky, 1998; Kitcher, 1984b). They claim that the world is such that there are multiple legitimate ways of classifying organisms into species.

The discussion of the various views on species classification in chapter 1 provides a straightforward argument for such ontological pluralism. This argument, which is also commonly used in the literature (see e.g. Ereshefsky, 2001, pp. 138–145; Kitcher, 1984b), consists of two premises. The first premise holds that there are various approaches to species classification that pick out species relevant to the explanatory and inferential demands of biology. The second premise of the argument holds that these different approaches do not always lead to the same classifications. Thus, two organisms might be conspecific on one approach, but not on another. It follows from these two premises that there are multiple legitimate ways of classifying organisms into species. Note that there is no agreement about what approaches are acceptable. Kitcher recognizes nine different approaches. Ereshefsky argues that an ecological approach, an interbreeding approach and a genealogical approach all pick out real species. My discussion in chapter 1 collapsed the ecological approach and interbreeding approach into a single process-approach. However, this disagreement does not affect the structure of the argument.

Part of this argument is empirical, as it relies on the claim that these approaches in fact pick out different groups. However, this empirical claim is not controversial. Examples confirming this empirical claim can be found in abundance in discussions and defenses of species pluralism such as Ereshefsky (2001), Reydon (2005), and Kitcher (1984a). Section 3 of this chapter will discuss more cases that illustrate this. In addition, these authors (and section 3) also provide theoretical reasons for why divergence between these approaches should be expected. Thus, the empirical claim on which the argument is based seems plausible. The argument then is that if the organic world is anything like what our current biological

knowledge tells us it is, then the various approaches discussed in chapter 1 (and further subdivisions of them) actually lead to different classifications that crosscut the organic world.

This argument for species pluralism fits well with the causal views on natural kinds. Pluralism about species as natural kinds simply implies that different approaches latch onto distinct parts of the causal structure of the organic world that are useful for different aspects of the explanatory and inferential demands of biology. For example, the genealogical approach tracks the causal history of a group of organisms, while the process-approach tracks causal factors such as interbreeding relations and selection pressures. In Reydon's (2015a) words, different approaches 'zoom in' on different parts of the causal structures relevant to the phenomena that scientists want to explain using the species category. This is also what Boyd (1999, pp. 170–171) expresses when, discussing the relation between HPC kinds and species pluralism, he writes that 'a plurality of species level classificatory schemes contribute significantly to achieving (different aspects of) the accommodation of inferential practices in biology to relevant causal structures'.

This 'ontological argument' (Ereshefsky, 2001, p. 138) shows that there are as many legitimate classifications as there are distinct legitimate scientific approaches to species classification. In that sense, it supports a moderate form of species pluralism. Even though there is no general consensus about how many scientific approaches to species classification there are, it is agreed that the number is relatively low. Ereshefsky (2011) recognizes three valid approaches, Kitcher (1984b, 1984a, 1987) recognizes nine, and one could argue that all modern species concepts, about thirty in total, lead to distinct natural classifications. Regardless of which of these positions is adopted, the resulting pluralism is of a relatively moderate kind.

Many species pluralists, including Ereshefsky and Kitcher, explicitly distinguish their position from more radical forms of species pluralism which hold that there are innumerable natural ways of classifying organisms into species. They do this because radical forms of species pluralism seem particularly vulnerable to the objection that they lead to 'anything goes' relativism (see e.g. Ereshefsky, 2001, pp. 158–162; Hull, 1987, p. 179; Sober, 1984, p. 334). The idea is that because the radical species pluralist accepts that there are innumerable many legitimate classifications, she does not have the tools to discern scientific classifications, such as the ones based on the genealogical history of organisms, from non-scientific classifications, like the ones based on creationist thinking.

My concern here is not whether the anything goes objection against radical pluralism is valid.¹⁶ Rather, I want to argue that the ontological argument for species pluralism presented above fails as an argument for moderate species pluralism, and instead leads to radical pluralism.¹⁷ In order to see this, consider the argument for moderate species pluralism again. There are two ways to understand this argument. The first, weak interpretation holds that there are *at least* three, nine, or thirty legitimate ways of classifying organisms into species. On this interpretation, the argument is undeniably valid but not strong enough to rule out radical pluralism; the claim that there are at least three natural classifications is perfectly compatible with the claim that there are innumerably many. The second, stronger interpretation holds that there are *only* three, nine, or thirty natural ways of classifying organisms into species. As the number of natural classifications is explicitly fixed, this interpretation does rule out radical species pluralism. However, the problem with this strong interpretation is that it simply does not follow from the premises. These premises state that there are multiple scientific approaches to species classification that pick out overlapping, legitimate species. Thus, it only follows that there are at least as many classifications as there are scientific approaches to species classification. To defend moderate species pluralism then, the ontological argument requires an additional premise.

One might think that moderate species pluralists need to assume that no new legitimate approaches to species classification will be discovered in the future. The idea is that this would increase the number of legitimate classifications, and thus, the extent of the pluralism supported by the argument. However, this assumption is not needed. Even if such radical new approaches to classification were to be found, they would only increase pluralism by one natural classification at a time, and so would hardly push the position towards a radical form of species pluralism. It is also worth noting that this thesis investigates the way taxonomists currently classify organisms into species. If biological knowledge changes drastically, and classificatory practices change accordingly, then the conclusions of this thesis become less relevant. However, such unexpected and drastic changes fall outside the scope of this investigation.

¹⁶ I come back to the anything goes objection in chapter 5, section 1.

¹⁷ I discuss the importance of the difference between radical and moderate species pluralism in section 4 of this chapter, and in section 4.2 of chapter 4.

However, there is another assumption that the moderate pluralist needs to make to avoid radical species pluralism. This assumption is that each of the three, nine or thirty approaches to species picks out a single classification of organisms into species. If this assumption is false, and these approaches each comprise many distinct but equally legitimate classifications, then the argument presented above does not lead to moderate species pluralism, but to radical species pluralism. The remainder of this chapter will argue against this assumption, and, hence, show that the main approaches to species classification lead to radical pluralism. To do this, section 2 returns to the causal account of natural kinds and discuss problems with individuating kinds on this account. It is argued that in causally complex parts of the world (like the organic world), the causal view of natural kinds leads to a radical form of natural kind pluralism. Section 3 then applies this argument to the causal approaches to species classification, and argues that the principles of kind individuation assumed by these approaches are compatible with innumerable different classifications.

2. Radical Pluralism and Causal Kinds

Given the dual structure of causal kinds, there are two possible grounds for kind individuation: the cluster of properties shared by the members of a kind, and the causal basis that grounds the stable clustering of these properties. It is clear that the property-cluster cannot serve as the only ground for individuation, as this would reduce the causal view of natural kinds to a Bare Property Cluster view. Moreover, the causal view of natural kinds explicitly holds that there is no set of necessary and sufficient properties that can serve to pick out all members of a kind. Instead, to be a member of a kind, a particular only needs to have some subset of all properties associated with the kind. Thus, the individuation of causal kinds must primarily rely on the causal basis that explains the stable property-clustering. As Carl Craver (2009, p. 575) puts it in his discussion of HPC kinds, 'kind concepts cut nature at its joints, and [...] nature's joints are located at the boundaries of mechanisms'.

This section will argue that if we take natural kinds to be individuated by the causal structures relevant to any particular scientific domain, then radical pluralism about natural kinds is likely to follow. I provide two reasons for this, based on Craver's (2009) discussion of the individuation of HPC kinds. Subsection 2.1 argues that radical pluralism follows from the fact that different combinations of causal mechanisms lead to different natural kinds. Thus, by splitting and combining mechanisms moderate pluralism quickly turns into radical pluralism.

Subsection 2.2 argues that using mechanisms to individuate kinds requires an account of the individuation of mechanisms. It is argued that such an account is likely to lead to radical pluralism about mechanisms, and hence, radical pluralism about natural kinds.

2.1. Selection of mechanisms

Accommodation of inductive and explanatory practices to the causal structure of the world requires that relevant aspects of this structure are reflected in scientific classifications. Craver phrases this accommodation-demand in terms of lumping and splitting of kinds. He writes that ‘if you find that a single cluster of properties is explained by more than one mechanism, split the cluster into subset clusters, each of which is explained by a single mechanism’ and ‘if you find that two or more putatively distinct kinds are explained by the same mechanism, lump the putative kinds into one’ (Craver, 2009, p. 581). For example, if it turns out that similar symptoms are caused by two widely different mechanisms, then we should recognise two diseases rather than one. If, on the other hand, the same mechanism leads to different symptoms, then we should recognise one disease rather than two.

It is uncontroversial to say that scientific kinds are rarely individuated by a single causal mechanism. This is perhaps most clearly illustrated by the case of species. Boyd (1999, p. 165) explicitly refers to the role of multiple mechanisms, such as gene exchange, reproductive isolation, and developmental constraints, in stabilizing similarity within species. Other proponents of the HPC view of kinds and species likewise emphasize that most species are characterized by a combination of mechanisms (Barker & Wilson, 2010; R. A. Wilson et al., 2007, p. 204). Indeed, it is taken to be one of the strengths of the epistemic views on natural kinds that the causal basis of natural kinds consists of multiple mechanisms that need not be instantiated in each member of the kind. This way, the causal view on natural kinds can account for the many scientific kinds that fail to meet the strict essentialist requirement of necessary conditions for kind membership.

While it is thus plausible that the causal basis of natural kinds sometimes consists of multiple distinct but interacting mechanisms, it also presents moderately pluralist proponents of the causal view of kinds with a puzzle. Given that various parts of the causal basis of a kind are considered distinct mechanisms, it must be that they make a relevant causal difference. Hence, it must be that the property clusters maintained by these mechanisms play different roles in meeting the explanatory and inferential demands of biology. It follows from this that

there are good reasons to recognise at least a distinct kind for each distinct mechanism, and perhaps additional kinds on the basis of combinations of these mechanisms.

Take, for example, gene kinds (Reydon, 2009b, 2015a; R. A. Wilson et al., 2007). Stretches of DNA of different organisms are individuated as genes of the same kind on the basis of their function, the particular stretch of bases they consist of, their genome locus, and their lineage. When various instances of the same gene-kind are considered, for example in different but related taxa, there often is divergence between these criteria for individuation. Two particular stretches of bases of the same lineage may have acquired changes since their lineage split. Some of these changes may have no functional effect while others may change the amino acid that the gene codes for. Similarly, two genes of the same lineage may be found on different loci in the genome or may occur in more or fewer copies throughout the genome. The point here is that including, combining or disregarding different mechanisms leads to different stretches of DNA being included or excluded in the gene-kind. Given that these kinds track different parts of the causal structure that is considered relevant for individuation, they all seem to be legitimate kinds. Assuming that the causal basis of many scientific kinds consists of a high number of mechanisms, it follows that radical pluralism of natural kinds is likely to obtain.

Craver thinks that this problem affects all causal kinds. He writes that ‘most putative kinds are entangled with myriad mechanisms. By attending to some of these, one is led to lump. By attending to others, one is led to split’ (Craver, 2009, p. 584). The worry seems particularly potent for parts of the world that are causally complex, like the organic world. If similarity-relations are maintained by countless mechanisms that are entangled in various ways and to various degrees, it is likely that even within one domain of research there are also countless ways of carving up these mechanisms and property clusters into natural kinds.

2.2. Individuating mechanisms

One straightforward way to avoid the problems discussed in the previous subsection is to deny that natural kinds can be affected by multiple mechanisms. Instead, one could argue that each causal mechanism simply corresponds to a single natural kind. For example, in the case of species one could see selection pressures, interbreeding, and developmental constraints as part of one integrated causal basis that affects different organisms in a similar way. A particular selection pressure, say, the presence of a particular predator or the aridity of the

environment, does not count as a distinct mechanism then, but rather as one aspect of an integrated causal basis.

In order to maintain the flexibility characteristic of the causal view on kinds, this 'one mechanism, one kind' view must allow the causal basis of different members of a kind to be different, as long as it is similar in the right aspects. This raises the question of how precisely such causal bases are to be individuated. This is not only important for the view that each mechanism corresponds to a natural kind. Even if natural kinds are characterized by several mechanisms, individuation of kinds still requires that these mechanisms can be individuated too. Craver (2009) argues that closer consideration of the individuation of mechanisms leads us to a radically pluralist view on natural kinds.

This is easiest to see for causal mechanisms that have a distinct instantiation in each member of the kind they are associated with. While some natural kinds are associated with a single mechanism exerting influence over different particulars (e.g. a particular selection pressure), most scientific kinds have a causal basis consisting of numerically distinct mechanisms associated with each member of the kind. For example, when we say that developmental constraints cause similarity within species, it is meant that each organism of the species has a numerically distinct mechanism that causes it to have the relevant properties. This counts as a common causal basis because these numerically distinct mechanisms are all of the same kind. It would be too strong to require that these numerically distinct mechanisms are identical. Even if they share many features, they are also undoubtedly different in many respects. This raises the question of how similar two developmental mechanisms must be in order to count as mechanisms of the same kind and consequently as causal bases of organisms of the same species. If the differences between these mechanisms are observable, they must be causally potent in some way. Grouping different developmental mechanisms together despite their differences thus picks up on some aspects of the causal structure of the world that are considered relevant, and ignores other aspects.

On a more abstract level, what Craver points out is that individuation of kinds cannot be solved by referring to the individuation of mechanisms. This is because the latter requires a view of kinds of mechanisms, and thus leads to an infinite regress of kinds. This could only potentially be stopped by kinds that have clearer conditions of individuation than is the case for causal kinds. For example, if kinds of mechanisms were defined in terms of essences with clear

conditions of individuation, then they could serve as a ground for individuating natural kinds like species. However, the problems for essentialism discussed in chapter 1 apply to kinds of mechanisms as much as they apply to species (for more on this regress problem, see Craver (2009, p. 586) and Slater (2014, pp. 386–387)).

It seems then that there is no easy way of individuating kinds of mechanisms. Depending on which similarities are considered relevant, and which dissimilarities are considered safe to ignore, we end up with different kinds of mechanism, and hence different natural kinds. Craver (2009, p. 584) writes that ‘pluralism about natural kinds would seem the most prudent option: there are as many kinds as there are distinct and dissociable mechanistic entanglements’. Given the complexity of the causal structure of the world, this pluralism is likely to be of a radical form, in which ‘there are many more natural kinds than science will ever find it useful or interesting to recognize, let alone study’ (Craver, 2009, p. 579).

The problems of selecting mechanisms and individuating mechanisms suggest that the causal view of natural kinds leads to radical pluralism in parts of the world that are causally complex. In response to these arguments, moderately pluralist or monist proponents of the causal view of natural kinds could point out that the individuation of natural kinds is not only determined by the causal structure of the world. Indeed, Boyd (2000, p. 57; see also Ereshefsky, 2001; Slater, 2017) explicitly claims that the individuation of kinds is fixed both by the causal structure of the world and by the ‘inductive and explanatory aims’ of the relevant disciplinary matrix. The idea then is that these aims reduce radical pluralism to a more moderate form, or even to monism. This response, which I will call the ‘Classificatory Norms View’, will be discussed extensively in chapter 5. I will argue that radical pluralism applies even if we understand natural kinds as relative to the explanatory interest of a reasonably broad discipline. To illustrate this argument briefly, consider again gene-kinds; because of the particular interests of evolutionary and functional biology, these kinds are not determined by just any causal factor, but only by their genealogical history, function and genome locus. However, even if we restrict individuation to only these mechanisms, there are still many different good classifications since there are many plausible ways of individuating and combining these mechanisms.

Given that the similarity between different members of a species is affected by a wide range of different causes, it seems plausible that Craver’s general argument applies to species

classification too. The next section considers the two causal approaches to species classification in more detail, and shows how the abstract arguments discussed in this section apply to species and their particular causal bases. That is, I argue that even if we restrict species individuation to the particular causal structures that the genealogical and process approach deem relevant, there are still innumerable ways of classifying organisms into species.

3. Radical Pluralism and Species

This section argues that species individuation on the basis of genealogical history or evolutionary processes leads to radical species pluralism. Broadly put, the argument is that the criteria of kind individuation used by the two approaches can be specified in very many different ways. Depending on which fine-grained aspects of the genealogical history and causal relations of organism are considered relevant, different groups of organisms are individuated as species and different species classifications result. Subsections 3.1 and 3.2 explicate this argument for the genealogical approach and process approach respectively. Subsection 3.3 applies the argument to the Evolutionary Species Concept, which combines the other two approaches.

3.1. The genealogical approach and pluralism

The genealogical approach individuates species on the basis of genealogical relations and common ancestry. As the tree of life is a representation of genealogical relations, one can also say that genealogical species are the result of cutting up this tree at its natural joints. These joints are reflected in the branching structure of the tree of life, which allows us to divide organisms into a nested series of exclusive or monophyletic groups. The genealogical approach considers all these taxa to be natural groupings, and recognises a subset of these as species on the basis of an additional ranking criterion.

If we assume that the tree of life forms a perfectly hierarchically ordered branching structure, then it is plausible that the genealogical approach leads to species monism, or, at most, a moderate form of pluralism. This moderate pluralism would be due to three reasons. First, different taxonomists could adopt different ranking criteria (Velasco, 2009, 2010). While some authors might argue that the least inclusive taxa are species, others might recognise more inclusive groups as species. Second, as discussed in chapter 1, different taxonomists and

philosophers interpret exclusivity or monophyly in slightly different ways (Rieppel, 2010; Velasco, 2009). This might lead them to different species classifications. Finally, different taxonomists could adopt a different understanding of species as kinds (or entities) extended over time. More specifically, even if they agree on how to individuate species in organisms living at any point in time, they might disagree on how to individuate the time-extended segments of the tree of life that these organisms are members of (Reydon, 2006). These sources of pluralism are reflected in current genealogical approaches to species delimitation, which vary on these three points. However, while they make moderate pluralism plausible, they do not seem sufficient to establish a radical form of species pluralism that cross-classifies the same organisms in many different ways.

This subsection argues that even if the three sources of moderate pluralism are not taken into account, the genealogical approach to species delimitation leads to radical pluralism. To show this, I evaluate a crucial assumption of the genealogical approach, namely, that the genealogical history of organic nature can be represented as a perfectly branching structure. This assumption is of course trivially false for sexually reproducing organisms, as their genealogical relations form a reticulated pattern. However, such reticulation is often claimed to only occur within species, and not between groups. Thus, proponents of the genealogical approach could argue, the branching structure of the genealogical history of life is retained on the level of species in spite of reticulation within species. This then allows them to recognize reticulating groups of organisms as genealogical species.

To re-emphasize the importance of the treelike structure of evolution for the genealogical approach, consider figure 1, which shows the genealogical relations between five groups. The pattern is not strictly treelike, because two of its branches merge again in B.

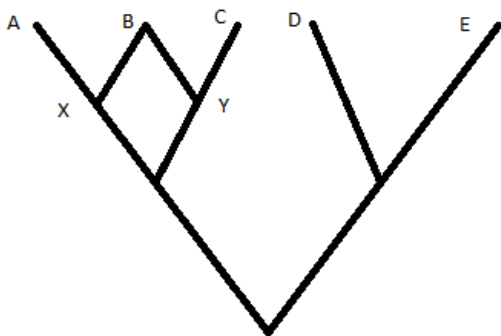


Figure 1. Non-Treelike Evolution. This figure illustrates non-treelike evolution and the problems it causes for the genealogical approach.

Groups A, C, D and E can be defined by their history and position in the tree of life. However, this is not the case for group B. Because this group is the result of two merging branches, not all organisms or genes in the group have the same genealogical history. While some are more closely related to X, others are more closely related to Y. This leads to pluralism about genealogical classification, as there is no single classification that represents all genealogical relations at once. While one classification may represent the genealogical history shared by all organisms or genes in B, that classification would have to omit the shared history of organisms in A and B, and in B and C. The step from here to radical pluralism is small. If such non-treelike genealogical relations occur particularly frequently and in many different and complex ways, it follows that there are very many distinct, cross-cutting classifications that represent various aspects of this complex genealogical history.

The remainder of this section will argue that this is indeed the case. That is, while some aspects of the genealogical history of the organic world can be represented by a branching diagram, there is no single best one that captures all genealogical relations. Moreover, many other parts of the genealogical history of life do not fit the branching structure assumed by the genealogical approach. It follows from this, I argue, that there is no single correct species classification based on genealogical history. Depending on which of the many particular genealogical relations we 'zoom in' on, different genealogical classifications result. To make these arguments, the following paragraphs briefly discuss various phenomena that cause genealogical history to be non-treelike.¹⁸

Lateral gene transfer

A much-discussed and major source of non-treelike evolution is lateral gene transfer (LGT) (Baptiste & Boucher, 2009; Doolittle, 2010; Ereshefsky, 2010b; Soucy, Huang, & Gogarten, 2015). This is the transfer of bits of genetic material not vertically from parent to offspring via reproduction, but horizontally through other mechanisms. Such transfer occurs most frequently between relatively closely related prokaryote taxa, but can occur between organisms from different orders and kingdoms too. It is now widely accepted that LGT has an enormous impact on prokaryote evolution, and it is not implausible that LGT has affected each

¹⁸ I am not the first to rely on non-treelike genealogical relations to argue against monism about genealogical classification (e.g. Baptiste et al., 2009; Hausdorf, 2011; Mallet, Besansky, & Hahn, 2015; O'Malley, 2014; Soucy, Huang, & Gogarten, 2015). Given that these claims are already widely accepted for prokaryotes, my discussion will mostly focus on eukaryotes.

single prokaryote gene over the course of evolution (Baptiste et al., 2009). This suggests a representation of the history of prokaryotes vastly different from a treelike pattern as bits of genetic material are frequently horizontally transferred between unrelated lineages, making some parts of the genome of a group or organism closely related to one group, and other parts to other groups.

Given that prokaryote life outnumbers eukaryote life on practically all ways of counting (O'Malley & Dupré, 2007), the abundance of LGT in prokaryotes poses severe problems for proponents of a genealogical approach to classification.¹⁹ There is also increasing evidence that the transfer of functionally relevant DNA also plays a significant role in eukaryote evolution (Soucy et al., 2015; see also references in O'Malley, 2010). Most importantly, it has been suggested that evolutionary innovation in protists shows important similarities to the processes that cause prokaryote history to be non-treelike, including LGT (Andersson, 2005, 2009; Boenigk, Ereshefsky, Hoef-Emden, Mallet, & Bass, 2012; Ricard et al., 2006). In addition, various authors argue that LGT also affects fungi (e.g. Richards, Dacks, Jenkinson, Thornton, & Talbot, 2006) and even metazoans (Drezen et al., 2017; Soucy et al., 2015). A recent paper by Crisp et al. (2015) reports on hundreds of foreign genes in the genomes of higher animals, including humans. They write that 'far from being a rare occurrence, LGT has contributed to the evolution of many, perhaps all, animals' and that 'the process is ongoing in most lineages' (Crisp et al., 2015, p. 12).

Introgression and hybrid speciation

One might say that at this point we still know very little about prokaryote groups, protists, and LGT in larger eukaryotes, so the conclusions based on these phenomena cannot be very strong and might still change (see also O'Malley, 2010, p. 544). However, even without taking these into consideration, there are several reasons why the genealogical history of these groups cannot be represented in a perfectly treelike diagram. Most important among these are, arguably, introgression and hybrid speciation.

Introgression occurs when genes of one species invade another species through backcrossing of hybrids with one of the parental species. Introgression is generally accepted to occur

¹⁹ There are more reasons why a treelike diagram is unlikely to capture all relevant aspects of lineage divergence in prokaryote groups (see Lawrence & Retchless, 2010). However, given that my claims are widely accepted with respect to prokaryotes, I will not discuss these here.

frequently, definitely between closely related and recently diverged species (Mallet, 2005, 2008, 2010; Mallet, Besansky, & Hahn, 2015). It is widely accepted to play an important role in plant evolution, and recently it has been argued that it plays a larger role than previously thought in animals too (Dowling & Secor, 1997; Hedrick, 2013; Mallet et al., 2015). One reason for the importance of introgression is that it can be a source of adaptive genetic variation. Through hybridization, adaptive alleles that have been under selection for many generations can be exchanged between species.

One notable example of this comes from North American Canids. A study by Monzon, Kays and Dykhuizen (2014) discusses hybridization between dogs, coyotes and different wolf-subspecies. The sampled coyotes in eastern United States had approximately 66% nuclear coyote DNA, 24% wolf DNA and a small proportion of dog DNA. The authors argue that the introgression of wolf and dog DNA into coyotes was particularly beneficial for these coyote populations. Due to their wolf ancestry, the sampled coyotes were significantly larger than regular coyotes, making them more efficient at hunting deer. Indeed, the authors also found that in areas with high deer density, the coyotes tended to have a higher proportion of wolf ancestry, which suggests that selection favored the introgressed genes. The authors also hypothesize that hybridization with domesticated dogs may have helped coyotes to adapt to environments dominated by humans (to the extent that even New York now has its coyote population). A case like this shows that even in groups like mammals, the genealogical approach cannot always be applied straightforwardly. A combination of introgression and adaptation to local selection pressures causes different populations of a species to have different proportions of genes with ancestry from different species.

Even more damaging for the treelike pattern of genealogical relations is hybrid speciation, which occurs when hybridization between two species leads to a new species. This is a relatively common phenomenon in plants, more particularly through allopolyploid speciation. This refers to the process in which hybridization leads to the birth of a new species with chromosomes from both parent species (see e.g. Ma, Li, Vogl, Ehrendorfer, & Guo, 2010). Mallet (2007, p. 281) claims that 2%-7% of speciation events in ferns and flowering plants are due to allopolyploidy, and that this most likely underestimates the real number. There are examples from the animal world too (Dowling & Secor, 1997; Gross & Rieseberg, 2005). Mallet and co-workers (Mallet, 2007, p. 282; Mallet, Beltrán, Neukirchen, & Linares, 2007) point to

examples of sculpid fish, fruitflies, and heliconiine butterflies. Another particularly interesting example consists of so-called klepton species like *Pelophylax esculentus*, which require continued input from a different species for successful reproduction (see Dubois, 1998).

More generally, Mallet claims that hybridization between species, which forms the basis of hybrid speciation and introgression, is far more common in eukaryotes than traditionally assumed. He (Mallet, 2005, 2008) claims that it affects at least 10% of all animals and 25% of all flowering plants. Moreover, like many of the references cited in this subsection, he also points to the evolutionary, adaptive importance of hybrid speciation and traits acquired through introgression. This suggests that it is problematic to assume, like the genealogical approach, that there is a more or less perfectly branching diagram that represents all relevant genealogical history.

Endosymbiotic gene transfer


Another phenomenon that disturbs the treelike pattern of genealogical history is endosymbiotic gene transfer (Timmis, Ayliffe, Huang, & Martin, 2004). It is widely assumed that mitochondria and chloroplasts, both organelles in eukaryote cells, were originally independently living prokaryotes. Thus, this is a case of prokaryotes that came to live in eukaryote cells, and eventually ceased being separate entities. In addition, DNA from those organelles has over the course of time gradually been transferred to the nucleus of the cell. Thus, DNA is transferred between very distantly related lineages. The consequence of endosymbiotic gene transfer is that different genes within one organism (and species) have radically different histories. Again, these genealogical relations are important for the explanatory and inductive demands of biology. For example, recent work on conflictual speciation suggests that transfer of genes within organisms can play an important role in speciation (Crespi & Nosil, 2013).

Incomplete lineage sorting and gene-tree/species-tree discordance

A final reason for the unfeasibility of the single, treelike diagram assumed by the genealogical approach is the ubiquity of incomplete lineage sorting (ILS) and subsequent discordance between species-trees and gene-trees.²⁰ Briefly summarized, ILS occurs whenever there is less

²⁰ For a more extensive discussion of the difficulties that the different levels of the hierarchy of lineages (e.g. genes, cells, populations) cause for constructing a single tree to represent genealogical relations, see Haber (2012; 2016), Degnan and Rosenberg (2006, 2009), and Maddison (1997).

than perfect segregation of different alleles into the various lineages that result when an ancestral lineage splits. This is easiest to explain using an image (see figure 2a). The thick lines represent the species tree of humans, chimpanzees and gorillas. The thinner lines within the species tree represent the history of particular genes. The ancestral population of the three species (H,C,G) was polymorphic for a gene, with a green, yellow and purple variant arising through mutation from the same ancestral allele. When a species splits into several new species, it is not necessarily the case that all different alleles of a gene are represented in all different branches. It is possible that one of the alleles is not present in the founding population of one of the lineages. This is the case for the orange allele in figure 2a. When the ancestral species (H,C,G) split into the gorilla lineage (G) and the Human + Chimpanzee lineage (H,C), the orange allele was not represented in (H,C). Over time, through drift and natural selection, variation in the Gorilla lineage was reduced, such that today the orange allele is the only variant left. Similarly, while the Human and the Chimpanzee lineages originally had both the green and the purple variant, the Human lineage today only has the green variant, and the Chimpanzee lineage only retained the purple variant.



Figures 2a and 2b. Incomplete Lineage Sorting. These figures illustrate the problems that ILS causes for the genealogical approach. The figure is copied and adapted from a blogpost by Venema (2013).

However, before these polymorphisms are sorted, it is not uncommon that there are multiple variants of the same gene in one population. If one of these variants also occurs in other species, then, at least for that gene, some organisms in the species are more closely related to organisms in the other species than to organisms in the same species with a different

variant. For example, at the start of the chimpanzee lineage on figure 2a, chimpanzees with the green variant are more closely related to humans with the green variant than to chimpanzees with the purple variant. Given that incomplete lineage sorting occurs in most if not all species, and often exists throughout the existence of the species, this phenomenon disturbs the treelike structure of genealogical relations.

Incomplete lineage sorting causes a further problem for the genealogical approach. To see this, consider again figures 2a and 2b. On figure 2a the history of the alleles also matches the history of the species. Just like humans and chimpanzees are more closely related to each other than to gorillas, the green and purple variants are more closely related to each other than to the orange variant. The process through which ancestral polymorphisms disappear and alleles become fixed in a species is called lineage sorting. In figure 2b, however, the history of the alleles does not match the species tree. While chimpanzees are more closely related to humans than to gorillas, the blue allele (retained in the chimpanzee lineage) is more closely related to the red allele (retained in gorillas) than to the yellow allele (retained in humans). In fact, about 30% of all human genes have a gene tree that does not match the species tree. Thus, the commonly accepted species tree as depicted in figure 2 represents the history of only 70% of the genes correctly. Generally speaking, incomplete lineage sorting probably occurs to some extent in all sexual species, and is particularly common when different speciation events occur in rapid succession (Hudson, 1992; Maddison, 1997; Pamilo & Nei, 1988; N. A. Rosenberg, 2002; Degnan & Rosenberg, 2006, 2009). This means that any species tree only represents part of the genealogical history of the genes embedded in the tree.

The problem that ILS causes for the genealogical approach is manifested particularly prominently in the discordance between the genealogical history of mitochondrial DNA (mtDNA) and nuclear DNA (nDNA). MtDNA is inherited maternally, while nDNA is inherited from both parental organisms. As a consequence, species trees constructed on the basis of mtDNA can diverge from nuclear gene-trees and vice versa. Take for example *Discoglossus jeanneae*, the Spanish painted frog. *Jeanneae* is the sister taxon of *D. galganoi*, and there is a longstanding debate in the literature about the relation and rank of these two taxa. About half of the specialists argue that the two taxa are subspecies, while the other half argue that they are both species (Busack, 1986; Fromhage, Vences, & Veith, 2004; García-París & Jockusch, 1999; Martínez-Solano, Gonçalves, Arntzen, & García-París, 2004; Velo-Antón, Martínez-

Solano, & García-París, 2008; Zangari, Cimmaruta, & Nascetti, 2006). The cause of this disagreement is conflicting mtDNA and nDNA evidence. The two taxa are reciprocally monophyletic for the former, but not for the latter. One reason for this could be that the effective population size of mtDNA is four times smaller than that of nDNA. Because coalescence times are inversely correlated to the effective population size, reciprocal monophyly arises much faster for mtDNA than for nDNA. The result is that for one set of (mitochondrial) genes, *jeanneae* forms a genealogical group with a unique evolutionary origin, but for another set of (nuclear) genes this is not the case.

In addition to ILS, there are other processes that cause divergence between mtDNA and nDNA. Consider again the case of *jeanneae* and *galganoi*. Some taxonomists argue that the divergence between nDNA and mtDNA is not caused by ILS, but by ongoing gene-flow between the two taxa (García-París & Jockusch, 1999; Velo-Antón, Martínez-Solano, et al., 2008). More specifically, they claim that *jeanneae* and *galganoi* regularly exchange nDNA genes, but not mtDNA genes. This is possible because mtDNA is only maternally transmitted. If dispersal and migration are male-biased, or if female hybrids are less viable, it is possible that two populations are isolated for mtDNA, but freely exchange genes for nDNA.²¹ This would result in reciprocal monophyly for mtDNA, but not for nDNA. If this is the case, then historical approaches have the problem that different parts of the genome lead to different species trees. This divergence between different types of DNA and consequently of the different possible trees can be stable and last as long as the two taxa are not extinct. Again, this means that there is no single branching tree that represents all genealogical relations equally well. Instead, there are various competing trees that all represent different aspects of this genealogical history.

I conclude then that both in prokaryotes and in eukaryotes there are various regularly occurring processes that disturb the treelike pattern of organic genealogical history. Thus, while some parts of the tree of life may be treelike for part of their genealogical history, other parts of the genealogical history of the same groups or organisms severely disturb this treelike pattern. In addition to this, a branching pattern seems entirely implausible for prokaryotes, which make up a large part of all life on earth. The picture of genealogical history that emerges

²¹ Decreased female hybrid viability is expected in *Discoglossus*, because in this taxon females are heterogametic and Haldane's rule states that in hybrids the heterogametic sex is usually less viable (García-París & Jockusch, 1999; Zangari, Cimmaruta, & Nascetti, 2006).

from this is one that is very complex, perhaps more like a web than a tree, with genealogical relations frequently crossing the borders of what is commonly called species (Soucy et al., 2015). The point here is that there is no single classification that can capture all fine-grained genealogical relations at the same time. Depending on which of these particular genealogical relations we focus on, different classifications that all represent some part of genealogical history result. As the web of genealogical relations is very extensive, and non-treelike processes are prevalent, it follows that there are very many such classifications. This means that individuating species on the basis of genealogical relations leads to radical pluralism.

3.2. The process approach and pluralism

The process approach defines species as groups of organisms cohesive in an evolutionary sense, i.e. unified by causal processes such as interbreeding and common selection pressures. Individuation of process species thus happens on the basis of these cohesion-promoting causal processes or mechanisms. In chapter 1, I pointed out that we can generally distinguish between two types of cohesion. The first, integrative cohesion, is the result of causal interaction between organisms. It obtains in a species whenever ‘there are causal interactions between many or all of its components that facilitate their causal unification into a whole’ (Barker & Wilson, 2010, p. 65). The second, response cohesion, is the result of organisms being subject to common causes. It obtains in a species whenever ‘its components respond (or are disposed to respond) as a unit to some kind of intervention’ (Barker & Wilson, 2010, p. 64). This section argues that the process approach leads to radical pluralism for both types of cohesive groups. The argument for this is similar to that in the previous section. Organisms interact in many varied ways, and they are generally subject to even more common causes. Consequently, both response cohesion and integrative cohesion come in many forms and degrees. Depending on which particular, fine-grained causal relations we select, different classifications result. Because there are innumerable many causal relations at play in the organic world, this pluralism is of the radical kind.

This argument is easiest to make for species that are response-cohesive, such as ecological species. Ecological species are cohesive because they are subject to a shared set of selection pressures. These common selection pressures keep the organisms of the species on a shared evolutionary trajectory even if they do not exchange genes. There are two reasons to think that individuating species on the basis of such common selection pressures leads to radical

pluralism. First, organisms are subject to a very high number of selection pressures. While there is undoubtedly considerable overlap between the groups of organisms affected by each of these selection pressures, it is also beyond doubt that they rarely pick out precisely the same groups, and also sometimes pick out radically different groups. One organism might be part of an ecological group that responds similarly to increased aridity of a region, while it is part of another group that responds similarly to predation in that region (for examples, see Barker & Velasco, 2013; Rieseberg & Burke, 2001). Second, it is reasonable to assume that selection pressures affect different organisms to a different degree. The question then is how similar the impact of a selection pressure on two organisms has to be in order to count them as members of the same group. Different answers to this question result in different species classifications.

Barker and Velasco (2013, p. 976) provide a clear example of how individuation of groups on the basis of response cohesion can lead to pluralism. They discuss a study by Barbará et al. (2007) on *Alcantarea imperialis*, a perennial plant species with highly varied traits and with populations located in significantly different environments. They point out that in addition to this, genetic variance and genetic distance between populations of the species is so high that it is near the level of that between various *Alcantarea* species. Together, this suggests that these populations are subject to various sets of selection pressures with different genetic resources responding to these selection pressures. Depending on which selection pressures and which responses are selected and combined, different groups are individuated. Given the high genetic distances between populations of the species, genetic responses to selection pressures are likely to be more similar between some populations of different species than between populations of *Alcantarea imperialis*.

Similar considerations apply for the individuation of groups of organisms that are cohesive in the integrative sense. The main evolutionary process responsible for integrative cohesion within groups of organisms is interbreeding. However, as already argued extensively in the previous subsection, hybridization is common throughout the tree of life. This means that instead of a clear boundary between interbreeding-relations on the species-level and phylogenetic relations above that level, it is better to speak of a continuum of interbreeding (Mallet et al., 2007; Mishler, 1999). Frequency and success of interbreeding relations are a matter of degree; they occur very regularly and are often successful above any supposed

species level (Mallet et al., 2015). While this might be rather limited for species closely related to humans, it occurs very frequently in plants. Similarly, levels of interbreeding are higher on the population-level than on the species-level. Indeed, depending on the time scale considered, it is often the case that populations are reproductively isolated from other populations of the same species. If a population is isolated for a few generations, there are good reasons to say that it forms an interbreeding group when just considering that period of time. However, this can happen frequently, and such brief periods of isolation can still be canceled out when gene-flow between the groups is restored.

In short then, there are innumerable many causal processes involved in the response cohesion and integrative cohesion of groups around the species level. Because these causal processes affect different groups of organisms in different ways, no single classification of organisms into species can capture all of them at once. Depending on the particular combination of fine-grained causal processes and responses, different classifications result. Given the complexity of the organic world and the sheer number of causal processes, I claim that this pluralism is of the radical kind.

3.3. The Evolutionary Species Concept and pluralism

The previous two sections have discussed in detail the causal structures that typically serve as the basis for species individuation. I have argued that these causal structures are so complex and multi-faceted that no classification can capture all of them. Instead, they can be carved up into species in many different ways depending on the particular fine-grained causal relations we focus on. The moderate pluralist could object here that this argument for radical species pluralism neglects an important feature of species classification. Species are usually characterized by a large set of shared properties, which in turn are explained by a set of interconnected causal mechanisms. One could argue then that species individuation should focus on such interconnected complexes of mechanisms, rather than the narrow causal bases considered in the previous subsections (see e.g. R. A. Wilson et al., 2007, p. 204; Brigandt, 2003). And while the arguments in the previous sections show that radical pluralism is true with respect to causal kinds with a narrow basis, it is not clear whether this is also the case if the genealogical and various versions of the process approach combine forces.

In section 2, I have already argued that one such ‘mixed’ approach, namely the HPC view on species (Boyd, 1999; R. A. Wilson et al., 2007), leads to radical pluralism. Another prominent

mixed approach is the 'Evolutionary Species Concept' (ESC) and its methodological implementation called 'integrative taxonomy'.²² This mixed approach defines species as independently evolving lineages (De Queiroz, 2005a; Mayden, 1999; Padial et al., 2009; Schlick-Steiner et al., 2010; Wiley, 1978). The ESC is particularly popular among taxonomists involved in species delimitation. This is because it seems to capture a common ground shared by many other approaches to species delimitation. Some taxonomists and philosophers have even referred to the ESC as the long-sought solution to the species problem (e.g. Padial, Miralles, De la Riva, & Vences, 2010, p. 2; De Queiroz, 2005b; Mayden, 1999). The remainder of this subsection will argue that, like the genealogical approach and the process approach it consists of, this ESC approach to species classification leads to radical species pluralism.

The easiest way to see why the ESC leads to radical pluralism is by looking at its central notion, namely, 'evolutionary independence'. A lineage is taken to evolve independently if the organisms that make up its populations generally face evolutionary pressures as a unit, distinct from other such units (Wiley 1978). Thus, independently evolving lineages are defined as groups of organisms that are cohesive in an evolutionary sense. The result of this evolutionary cohesion is that the organisms that make up these lineages have a shared unique path through evolutionary space.

Proponents of the ESC claim that evolutionary independence is typically instantiated differently in different species. Some species are reproductively isolated, others form an exclusive group, others are subject to a similar ecological regime, and yet others are characterized by various combinations of these mechanisms. Despite the differences between the underlying mechanisms, these groups are similar in that they consist of organisms that are on a shared trajectory through evolutionary space. Thus, this notion of evolutionary independence is intentionally vague in order to capture the wide variety of lineages in the organic world. It is through this deliberately vague notion of evolutionary independence that the ESC integrates different approaches to species classification.

The inclusive character of the ESC is often cited as one of its main strengths. However, this strength comes with an important worry on the other side of the coin. The discussion of the

²² In chapter 1, I classified the ESC as a special case of the process-approach, because it classifies organisms into species on the basis of a shared evolutionary trajectory. I discuss the ESC separately here because of its popularity in taxonomic practice.

genealogical and process approach in the previous sections suggests that these various operationalizations associated with evolutionary independence, such as various fine-grained genealogical relations and shared selection pressures, rarely pick out the same groups. For example, many groups that are morphologically diagnosable are not reproductively isolated, while there are also many morphologically cryptic groups that are reproductively isolated. Indeed, various biologists emphasize that disagreement between these operationalizations is more common than agreement (Padial et al. 2009; Schlick-Steiner et al. 2010; Willis 2017). In practice, this means that different operationalizations of evolutionary independence tend to pick out different groups as species. It follows then that depending on how taxonomists decide to operationalize the notion of evolutionary independence, different species delimitations result.

This suggests that the ESC's integration of various approaches to species pluralism fails to fend off species pluralism. As Ereshefsky (2011, p. 15) writes:

The point is that De Queiroz attempts to unify the species category by asserting that all and only lineages are species. But that just masks the heterogeneity of the species category because what constitutes a lineage has multiple answers, and those answers vary according to which species concept one adopts.

This problem is even more pressing given the arguments made in the previous subsections. They show that diverging classifications can not only come from different species concepts, but also from countless more fine-grained specifications of these concepts.

4. Different Kinds of Pluralism

The previous section argued that radical pluralism obtains if we assume that species are individuated on the basis of genealogical relations, cohesion-promoting mechanisms, or a combination of these. The argument for this radical pluralism mirrored the arguments for pluralism about natural kinds from section 2. I argued that the features of the world that the genealogical approach and process-approach consider relevant for species individuation comprise many distinct fine-grained causal relations, processes and mechanisms. As all these fine-grained causes and mechanisms are part of the relevant causal structure of the world, they all pick out equally real genealogical groups or process-based groups. Moreover, there is no single classification that can represent all these relations at once. This means that radical

pluralism obtains within any genealogical or process-based approach to species classification that is not specified to an absurd degree (e.g. groups cohesive due to selection pressures X at time T with response R etc.). In other words, for all these approaches, it is true that ‘any group of organisms can be classified in indefinitely many ways’ and ‘no reasons internal to these methods exist for choosing among these classifications’ (Hull, 1997, p. 360).

I am, of course, not the first philosopher or biologist to defend species pluralism. As Hull (1999, p. 24) writes, ‘numerous senses of *monism* blend imperceptibly into just as many senses of *pluralism*’ and ‘everyone seems to feel obligated to espouse the position held by all thoughtful scholars – a nuanced pluralism.’ Indeed, over the past four decades, many authors have proposed different varieties of species pluralism and monism that differ in various subtle ways (Dupré, 1993; Ereshefsky, 1992; Kitcher, 1984b; Mishler & Donoghue, 1982; Reydon, 2005; Stanford, 1995). This section will briefly discuss in what respects my radical pluralism differs from other variants of pluralism. In addition to positioning my arguments in the literature, this will also further clarify their scope and import.

Most famous among different variants of species pluralism is John Dupré’s promiscuous realism (Dupré, 1981, 1993, 1996, 1999). According to Dupré, there are very many good ways of cutting up the world into natural kinds. I will not go into Dupré’s views in detail. Here it suffices to say that his position is at the same time more radical and less radical than the pluralism defended here. Dupré’s promiscuous realism is more radical because it grants an equal status to good scientific kinds and good non-scientific kinds (Dupré, 1981, 1996). According to Dupré, both can capture relevant aspects of the world, and there is no reason to think that one is more real than the other. While I do not necessarily object to this position, it falls outside the scope of my arguments. I am only interested in the kinds of groups that scientists classify as species. I argue that there are very many ways of classifying the organic world into those kinds of groups. This position is independent from views on the status of non-scientific kinds.

With respect to scientific kinds, the pluralism defended here is more radical than Dupré’s promiscuous realism. This is because I argue that radical pluralism obtains even if we take into account that species classifications should be accommodated to the explanatory and inferential demands of biology. That is, I argue that even for a relatively narrow domain of research, there are innumerable equally good ways of classifying organisms into species.

Dupré, on the other hand, seems to imply that radical pluralism only obtains if the legitimacy of natural kinds is considered independently from the purposes of classification. He writes that ‘relative to a sufficiently well-articulated set of aims of enquiry there may very well be, and often is, a best way of classifying the phenomena within a domain’ (Dupré, 2002, p. 31).²³ The arguments discussed in this chapter, relying on the complexity of fine-grained causal relations in the organic world, suggest that such aims would have to be specified to such a high degree that they could no longer plausibly be attributed to a domain of scientific enquiry. A second position that requires consideration here is Boyd’s. While he does not discuss pluralism extensively, he points out that his HPC view of natural kinds is compatible with species pluralism if it turns out that accommodation to different parts of the causal structure of the organic world is required for the explanatory demands of different research domains in biology (Boyd, 1999, p. 169). Like Dupré, however, Boyd assumes that classificatory monism or moderate pluralism applies within any one domain of research or ‘disciplinary matrix’. In that respect, Boyd’s position differs from the radical pluralism defended here, which holds that even within one research domain there are still many different equally good ways of classifying organisms into species. For example, even if it were clear that the explanatory and inductive demands of the historical domains of biology require us to accommodate species classification to genealogical history, there are still countless different good ways of zooming in on these structures.²⁴

It is also worth distinguishing my radical pluralism from Boyd’s claim that HPC kinds often have fuzzy boundaries. Boyd claims that ‘HPC kinds sometimes ‘[lack] precisely defined membership conditions’, which leads to the ‘vagueness in the extension’ of kinds (Boyd, 2010, p. 216, see also 1991, pp. 141–142). According to Boyd, this vagueness is not problematic for the HPC view, and even constitutes one of its strengths. He argues that the world just happens to be such that the boundaries of some kinds are vague, so any good view on natural kinds must incorporate this vagueness rather than provide artificial criteria for kind individuation that represent the world as if there are clear boundaries between all natural kinds. This is

²³ It should be noted that it is not entirely clear how radical Dupré’s pluralism is precisely. Elsewhere he writes that biologists might need ‘specialised classificatory schemes for their particular theoretical projects’ (Dupré, 1999, p. 4). This can be taken to support a more radical pluralism than what is suggested by the citation above.

²⁴ I come back extensively to the differences between Boyd’s and my views concerning pluralism in chapter 5.

particularly relevant for species, because the evolutionary processes that cause and maintain species are expected to lead to vague boundaries between groups.

The radical pluralism defended here goes beyond the claim that kinds have fuzzy boundaries. While some of the processes of non-treelike evolution merely lead to fuzziness in the boundaries of kinds, and thus only to problems of line-drawing, this is definitely not the case for all these processes. Various selection pressures can affect widely different groups of organisms, LGT and endosymbiotic gene-transfer connect groups that are far apart on the phylogenetic tree, and hybrid speciation occurs regularly both in plants and animals. Thus, radical pluralism does not only concern classifications that are minimally different, but also classifications that are radically different from other legitimate classifications. In addition, it is worth mentioning that for some explanatory or inferential demands of biology, the difference between different individuations of fuzzy kinds may be important. The mere fact that the difference between some kinds is minimal does not imply that they cannot be different kinds.

Kitcher and Ereshefsky defend a moderate form of species pluralism that fits well with Boyd's view. Kitcher argues that there are about nine legitimate approaches to species classification. According to Kitcher, these different approaches are biologically interesting in different ways because they track different features of the organic world. Kitcher (1984b, p. 325) recognizes an ecological approach, approaches based on reproductive isolation, and even structural approaches among others.²⁵ Ereshefsky (1992, 1998, 2001), on the other hand, only recognizes interbreeding lineages, ecological lineages, and phylogenetic lineages. As both restrict species classification to these approaches, their pluralism is moderate. The pluralism in this chapter differs from Kitcher's and Ereshefsky's in that it holds that each of the approaches they recognize still comprises many distinct classifications. For example, depending on which actual interbreeding events are considered relevant, different classifications based on reproductive isolation will result.

One might worry here that there is no philosophically interesting difference between my radical species pluralism and these more moderate variants. This worry derives from the fact that the difference between radical and moderate pluralism is merely quantitative. That is, the difference simply lies in the absolute number of potentially legitimate classifications: while

²⁵ It should be noted that in his later work, Kitcher briefly refers to a pragmatist view on species and species pluralism that is similar to the view defended here.

the moderate pluralist may accept three, nine or thirty classifications, the radical pluralist accepts, say, thousands of classifications. Any dividing line between these two forms of pluralism then seems arbitrary. However, while it may be that the difference between these two forms of pluralism is merely quantitative and any dividing line is arbitrary, it does not follow from this that the difference is not philosophically interesting. I will argue in chapter 4 (section 4.2) that the difference between radical and moderate pluralism is of crucial importance to the way we should regulate and evaluate values in taxonomy.

Finally, it is worth comparing my radical species pluralism with Kyle Stanford's anti-realism about species. Stanford (1995) argues that different classifications can track different but equally real features of the organic world. Kitcher claims that only nine of these ontologically good classifications are also biologically interesting, and thus defends a moderate pluralism. Stanford argues that this criterion for the legitimacy of classifications implies that there are very many potentially legitimate classifications, depending on the many ways in which our biological interests may change in future. At first sight, this position lies close to the radical pluralism defended here. Stanford argues that the structure of the organic world can be carved in many ways that all represent some part of this structure. And while he calls this position 'anti-realism', his assumption that these classifications track real features of the organic world suggests that it might equally well be called promiscuous realism.

However, Stanford's position is importantly different from the radical pluralism defended here. This is because Stanford relies on the claim that biological interests may change drastically, and thus any classification that is relevant to any conceivable interest is equally legitimate. The argument in this chapter is that even if we assume the explanatory and inferential demands of contemporary biology, there are innumerable different legitimate ways of carving organisms up into species. This radical pluralism does not rely on farfetched potential biological interests, but instead on the complexity of the innumerable fine-grained causal relations that make up the causal structures that are deemed relevant to the explanatory demands of biology.

While none of the major proponents of species pluralism have defended my variant of radical pluralism, there are various authors who defend a similar position in closely related domains. As already pointed out in section 2, Craver argues that the HPC view on natural kinds leads to radical pluralism due to the innumerable ways in which mechanisms can be individuated and

combined. Jacob Stegenga (2016) makes a similar argument about populations. He argues that if we understand populations as groups of organisms unified by fine-grained causal relations between the organisms, innumerable different classifications can be considered legitimate depending on which of these causal relations are considered relevant. Finally, Barker and Velasco (2013) make a similar argument about evolutionary groups in general. They argue that regardless of the processes or patterns that one takes as the causal basis of such groups, these can always be specified in countless different fine-grained ways, leading to countless equally legitimate classifications.

5. Summary

This chapter has argued that if we assume the principles of species individuation adopted by the genealogical and process-based approaches, then radical species pluralism follows. Given that radical species pluralism is trivially true on similarity-based approaches, it follows that all three dominant approaches to species individuation lead to radical species pluralism. Very broadly, the argument for this claim is that the features of the world that these classifications are supposed to latch onto are very multi-faceted and complex. Because of this complexity and diversity, they cannot be represented in a single classification. Instead, different classifications reflect different parts of this complex structure. It follows that radical species pluralism obtains if we assume the genealogical approach, the process approach, and the similarity-based approach as the principles for species individuation.

One might argue that this argument based on the causal complexity of the organic world is trivial. Mishler and Brandon (1987, p. 402), for example, write that:

controversies in evolutionary biology over causal agents generally do not involve claims that all but one favoured agent are impossible. Rather, a number of causal agents are acknowledged to be possible and controversy centers around which agent is the "most important".

What Mishler and Brandon seem to point out is that species individuation should not solely rely on the causal basis of kinds, but also on value-judgments. This brings us back to the guiding interest of this thesis, namely, the role of values in species classification. In the next chapter, I turn to the relation between radical species pluralism and the role of values in species classification.

Chapter 3: The Argument from Ontological Underdetermination

Having made the argument for radical pluralism in the previous chapters, I return to the main concern of this thesis, namely, the role of value-judgments in species classification. Over the next three chapters, I investigate the implications of radical species pluralism for the role that values should play in taxonomy. I will argue that radical pluralism causes values to play a profound role in the construction and evaluation of species classification. This chapter takes the first step towards defending that claim; it offers an argument for the value-ladenness of taxonomy, and sketches the theoretical framework that allows us to gauge how profound this value-ladenness is.

The structure of the chapter is as follows. The first section introduces and clarifies the notion of 'values'. The second section proposes an argument for the value-ladenness of taxonomy on the basis of radical species pluralism. This argument, which I call the argument from ontological underdetermination (AOU), will be the main focus of the remainder of this thesis. The third section considers and refutes three potential objections to the AOU. Having shown that taxonomy should be value-laden, I turn to the role of these values in section 4, which introduces three competing frameworks that one can adopt to think about this value-ladenness of taxonomy: the 'Value-Free Ideal', the 'Aims View' and the 'Epistemic Priority View'. These views all accept that values should influence scientific decisions, but disagree about how profound these value-influences should be. The Value-Free Ideal argues for a superficial role, the Epistemic Priority View for a moderate role, and the Aims View, finally, for a profound role for values in science. These considerations then form the starting point for chapter 4, which investigates which of these three frameworks best captures and accounts for the value-ladenness established by the AOU.

1. Values and Epistemic Standards

To make the importance of questions about values in taxonomy more salient, it will be helpful to start with an example. Consider the hypothetical case of a taxonomist who is greatly

concerned about biodiversity conservation, and more particularly about a morphologically distinct group of birds that she discovered to be living in a small forest that will be cut down for a large real estate project in the near future. Let us assume that the real estate project can be put on hold if there are groups with special conservation status endemic to that forest, and that the local conservation legislation only grants special conservation status to groups formally recognised as species. With her conservation interests in mind, the taxonomist is then likely to set out to collect data and investigate whether the group should be recognised as a species. This case illustrates the entanglement of values and taxonomy, and raises a long range of questions about the appropriate role for values in species delimitation: Is it appropriate that the taxonomist decides to investigate the organisms in the hope of recognising them as a species? If the evidence is undecided about the status of the group, should she recognise them as a species anyway, knowing that not doing so will lead to their extinction? If the evidence points towards not recognising the group as a species, do her conservation values provide legitimate reasons to postpone judgment? What if instead of environmental values, the taxonomist is particularly concerned about the financial interests of the real estate developer?

Questions such as these have received increasing attention over the past few decades, and the debate on values in science now constitutes an important and vibrant subfield in the philosophy of science (e.g. Anderson, 2004; M. J. Brown, 2013; Douglas, 2009; Intemann, 2001; Kitcher, 2001; Kourany, 2003; Longino, 1990; Steel, 2010). Because of the enormous potential and actual impact of science on all areas of life and society, it is clear that values and science should be and inevitably are closely connected. At the same time, innumerable examples of values corrupting science, such as the Lysenko case (Graham, 1987, pp. 102–150) and the Vioxx case (Biddle, 2007), suggest that there are dangers connected to values influencing scientific decisions. The authors engaged in these debates about values in science attempt to clarify what role values can and should play in science, and when they form a threat to the proper functioning of science. One important aim of these debates, then, is to develop a framework for evaluating the role of values, and to act as a guide for scientists in future scientific work.

One of my aims in this thesis is to contribute to developing such a framework, and to shed light on the kinds of questions raised by the example above. One might think that in order to

do this, it is necessary to provide a precise definition of what values are. Given the diversity of ways in which the term 'values' is used both within and outside philosophy of science, this would not be an easy task. A long tradition of work in the philosophy of science has shown the difficulty of providing a precise criterion that demarcates values from other factors that can influence scientific decisions, such as empirical evidence and rules of inference (e.g. Dewey, 1929; Dupré, 2007a; Putnam, 2002). Fortunately, I do not need such a general definition to make my arguments. This is because despite the lack of agreement about the definition of values, there are more than enough examples of considerations that are generally accepted to be values, such as fairness, equality, or, more relevantly in the context of taxonomy, the importance of biodiversity conservation and improving people's health. Much of the debate about the legitimacy of values in science relies on using such examples and proceeds without defining values more precisely. In line with this, I will use the term 'values' to refer to considerations like these paradigmatic examples. The question then is whether considerations about conservation, health, equality, fairness, etc. should influence species classification, and how precisely they should do this.

It is important to note here that the term 'values' is also commonly used to refer to a set of normative considerations that are thought to single out scientific outcomes as epistemically good, such as empirical adequacy, internal consistency, and explanatory power (Douglas, 2013; Laudan, 2004; McMullin, 1982; Steel, 2010). Because these so-called epistemic values are often considered to be connected to the attainment of truth, their role in science is relatively uncontroversial and they are generally accepted to influence scientific decisions throughout the research process. This stands in stark contrast to the social, political and moral values referred to above, whose role in science is far more controversial. Indeed, most of the controversy concerning values in science focuses on these social, moral, and political values rather than on the obviously epistemic ones. In the remainder of this thesis I will use the term 'values' to refer to the social, moral, political and environmental normative considerations that may influence scientific decisions. I will use the term 'epistemic standards' to refer to the normative considerations for which there is a prima facie good case linking them to the attainment of truth, and which are accepted to be truth-attaining by many, such as empirical adequacy, consistency, inter-theoretic coherence and explanatory strength.

By adopting this vocabulary, I do not mean to commit to an absolute distinction between these two kinds of normative considerations. Indeed, it is far from clear whether the distinction between epistemic standards and values is tenable. While there are relatively uncontroversial cases for each of these – for instance, it is commonly agreed that empirical adequacy and internal consistency are epistemic standards, and political and moral concerns are values – it is unclear where precisely to draw the line between the two groups. The problem is that attempts at specifying the distinction either include too many values in the category of epistemic standards or vice versa. Furthermore, there are many normative considerations that are neither clearly epistemic standards nor values, such as simplicity or scope, and even considerations typically considered to be non-epistemic values are sometimes truth-attaining. It is not surprising then that many authors have argued that the distinction between epistemic standards and values is problematic and ultimately untenable (e.g. Longino, 1995, 1996; Rooney, 1992).

Daniel Steel (2010; see also McMullin, 1982) has recently attempted to revitalize the distinction between epistemic standards and values by defining the former as any value that is truth-attaining. He thus rejects the opposition between social values and epistemic standards, and claims that any social value that is truth-attaining in a particular research project should be considered an epistemic standard. This understanding of epistemic standards solves some of the problems raised above, but comes with the important downside that epistemic standards can no longer straightforwardly be used to evaluate truth. Rather, it seems that we need to know the truth before we can decide whether a particular value was an epistemic standard. This makes the distinction between epistemic standards and values unhelpful for evaluating the role of values in any ongoing research project in which it is still unclear where the truth lies (Douglas, 2016, p. 12). I will not go deeper into these issues here and instead assume, for the sake of the argument, that the distinction holds. I do this because it is assumed by the positions I argue against. I want to show that even if we assume this distinction to hold, it is still the case that values play a profound role in taxonomy. If it turns out that the standards/values distinction is indeed untenable, this does not invalidate my points, but rather reinforces my claims about the role of values in taxonomy.

In short then, the claim that taxonomy is profoundly value-laden should be understood as the claim that moral, political, environmental, social, etc. considerations should play a substantial

role in taxonomic decisions. This claim is intended to capture the role of these values independently from the truth-attaining effect they may have. While it may be that the influence of social values on scientific decisions sometimes makes a true outcome more likely, I argue that this is not the only reason why these values should influence taxonomy. In other words, I argue that values should play a profound role in taxonomy even if they do not contribute to attaining the truth.

2. The Argument from Ontological Underdetermination

The previous section clarified the notion of values, and restricted it to the moral, social and political considerations that are typically the subject of debates about values in science. The aim of this section is to show that these considerations should and unavoidably do play a profound role in making scientific decisions in taxonomy. To do this, I rely on what I will call the argument from ontological underdetermination (AOU). This argument holds that values influence taxonomy by determining the choice between classifications that all pick out real but different groups. This argument follows directly from radical species pluralism. Values are needed in taxonomy to choose between innumerable genealogical and process-based species. If the world had been such that there was only one good way of classifying organisms into species (e.g. the creationist view), or if the species concept had been reserved for one particular kind of group (e.g. groups of organisms that are reciprocally monophyletic for 50% of their genes) then species classification would not require such value-laden ontological choices. As it is, however, the organic world is causally extremely complex, and the species concept is used to refer to groups individuated on the basis of various combinations and selections of these causes. Given that the world does not determine the choice between these groups, value-judgments are required to settle on one or a few classifications that are considered legitimate.

The AOU has a normative and descriptive side. The descriptive argument shows that values inevitably play a role because species classification is ontologically underdetermined. This descriptive argument implies that any species classification contains value-judgments that, together with the world, shape this classification. The normative argument holds that because the ontological choices required for species classification have non-epistemic consequences, values should be taken into account when making these choices. That is to say, ontological choices determine the result of research, and therefore influence the ways in which the results

of research are used in further research and applications outside science. Assuming that, like any other person, taxonomists are responsible for the consequences of their actions, they should take the foreseeable consequences of their ontological choices into account when making these choices. Imagine, for example, that a scientist has to choose between recognising a phylogenetic and a biological species, and that choosing the phylogenetic species would increase the survival chances of the organisms more than choosing the biological species would. The normative argument from ontological underdetermination holds that these conservation consequences should be taken into account when deciding which of these two groups to recognise.

To illustrate the argument, consider again the hypothetical case of the taxonomist and the real estate project. Spurred by her conservation interests, the taxonomist sets out to investigate the group of organisms using a diagnosability-based phylogenetic species concept. She conducts her research using commonly used methods and sampling commonly used genes, and concludes that the group is a species. At the same time, those managing the real estate project hire another taxonomist to investigate the same group of organisms. The second taxonomist adopts the BSC, uses other well accepted methods and samples other genes, and concludes that the group is not a species. This illustrates, in line with the AOU, that the decision of whether the group is a species is underdetermined by the world, and hence requires value-judgments to be settled. The first taxonomist might rely on her conservation values and choose to accept the group as a species. The second taxonomist might rely on the financial interests of the real estate project and reject the species. Either way, values enter taxonomy through the choice between two ontologically equally legitimate options. Note that both normative and descriptive arguments are at work here. According to the descriptive argument, any classification we end up with is value-laden because classificatory decisions are not fully determined by the organic world. According to the normative argument, taxonomists should take the consequences for conservation biology or the real estate project into account when choosing between the PSC and BSC.

The argument from ontological underdetermination has received significantly less attention than other arguments for values in science. David Ludwig (2015) discusses the relevance of value-considerations for ontological choices in species classification, psychiatry, intelligence testing and disease classification. Others have discussed ontological choices implicitly, or in

close combination with other aspects of the role of values in science. Anderson (1995b), for example, argues that the choice of evidence and the representation of this evidence plays a crucial role in theory choice. As there is no such thing as complete evidence or the complete truth concerning a phenomenon, collecting evidence, representing the evidence and assessing the evidence require value-laden judgments concerning the significance of potential evidence. While Anderson does not explicitly refer to ontological underdetermination, the following quote illustrates how it plays a crucial role in her argument.

[T]he world is too complex and messy to be organized into a few layers of all-inclusive and mutually exclusive classifications that account for all causal regularities. For each classification that supports some causal regularity, there are likely to be some other crosscutting ones in the neighbourhood that bear a causal relation to some other phenomenon. So criteria of epistemic significance alone do not tell us which classifications to base our theory on. (Anderson, 1995b, p. 45)

A version of AOU is also often made in discussions of species pluralism, and, more generally, ontological pluralism. Kitcher (2001), for example, argues that values play a role in determining which of several species concepts to use. Similarly, Boyd points out that the explanatory and inductive goals of a disciplinary matrix co-determine the kinds recognised in that discipline. Dupré (1993, 2001), perhaps most famously, also uses his promiscuous realism about natural kinds to argue that classifications are always tied to their particular purpose; he writes that

There is no God-given, unique way to classify the innumerable and diverse products of the evolutionary process. There are many plausible and defensible ways of doing so, and the best way of doing so will depend on both the purposes of the classification and the peculiarities of the organisms in question [...]. (Dupré, 1993, p. 95)

These views, which I will call the 'Classificatory Norms View', are similar to the views presented here. I will come back to the differences between the classificatory norms view and my view extensively in chapter 5. Here it suffices to say that my view differs from this view in two significant ways. First, while most of these authors defend a moderate form of pluralism, my argument is based on radical pluralism and hence leads to a radical form of underdetermination: values are needed to choose between innumerable legitimate classifications. Chapter 4 will argue that this has direct implications for the strength of the

AOU. Briefly put, I will argue that on a moderately pluralist view of species, it is easy to see the value-laden choices between multiple ontologically valid classifications as a part of the external rather than internal stages of science. The external stages of science concern the choice of research topics and the application of scientific outcomes, while the internal stages concern the characterization of the evidence and the evaluation of the hypothesis on the basis of the evidence.²⁶ I will argue that because of the radical nature of species pluralism, it is impossible to see the value-laden choices between multiple ontologically legitimate options as part of the external stages only. This is important, as the legitimacy of values in the internal stages of science is far more controversial than the legitimacy of values in the external stages. The second difference between the AOU and the Classificatory Norms View is that the latter is typically connected to the claim that value-judgments can be reduced to the generally shared goals of the discipline. While these goals are value-judgments, this interpretation of the argument mitigates the impact of values on species classification. On this view, values only play a role in selecting one classification from a set of ontologically good classifications that can be found in the world. In chapter 5, I will argue that this view does not fully capture the role of values in taxonomy. I will argue that in addition to these generally accepted goals of taxonomy, other value-judgments should play a role too. These additional value-judgments are not only involved in selecting the classification that is appropriate for our purposes, but also in shaping the species recognised by these classifications.

3. Three Objections to Ontological Underdetermination

In order to develop the AOU in more detail, it is worth considering three objections that one might raise against the argument. The first objection is that values may be needed for choosing between multiple options, but that this influence of values can easily be avoided simply by accepting all ontologically legitimate options. For example, if we accept both a classification of PSC species and a classification of BSC species, taxonomists need not rely on values to choose between them. This would mean the AOU fails to establish a role for values in taxonomy.

²⁶ I discuss the distinction between the internal and external stages of science extensively in section 4.1 of this chapter.

This objection faces several problems. First, given that the AOU is based on a radical form of pluralism, the strategy of accepting all ontologically valid options would lead to severe practical problems. Taxonomists would not only have to recognise a PSC and a BSC classification, but also one for many other species concepts, and for the innumerable operationalizations of each of these species concepts. Not only would this multiply the already immense task of discovering and recognising all species, it would also make it significantly more difficult for conservation biology, and health and trade legislation to use the results of taxonomic research. A second problem for this objection is that taxonomists do not face ontologically underdetermined decisions only in the final stage of research, where they accept or reject species-hypotheses. As I will argue in chapter 4, the taxonomic research process consists of a long range of ontologically underdetermined decisions, each of which influences later decisions. For example, the choice between two different methods of analysis, sampling strategies or data-types may be partially ontologically underdetermined. To accept all options in each of these cases would make finishing any research project nearly impossible. In short, there are very strong practical reasons for scientists to choose between multiple ontologically good options rather than to accept all. As this means that they have to rely on value-judgments, it follows that the first objection fails.

The second objection to the AOU holds that scientists need not necessarily rely on value-judgments to choose between multiple ontologically legitimate options. If there are other factors available that can determine this decision, taxonomy might remain value-free despite ontological underdetermination. Most importantly, one could argue that taxonomists can rely on epistemic standards for these decisions. I argue that this objection fails because epistemic standards are not relevant for ontologically underdetermined choices. The AOU holds that values enter taxonomy through the choice between multiple ontologically good options. Assuming that all these options are already true descriptions of the world, it is not clear how truth-attaining standards could help in choosing between them. The point here is that the aim of truth underdetermines the choice of classification. Assuming this, it is clear that truth-attaining standards cannot fully determine this choice either. In other words, epistemic standards are in principle unable to determine ontologically underdetermined decisions. Thus, even if we accept that epistemic standards guide scientific decisions, taxonomists must also rely on value-judgments for those decisions that are not fully determined by the aim of truth.

The final objection to the AOU holds that even if we accept that classification is underdetermined by both the world and epistemic standards, it does not follow that values have to fill the remaining gap. In response to the similar argument from epistemological underdetermination in the context of climate models (see Winsberg, 2012), Wendy Parker (2014) argues that the gap between theory choice and evidence is usually closed by pragmatic considerations instead of values.²⁷ Similarly, Morisson (2014) points out that subjective decisions due to epistemological underdetermination may be embedded in the ‘nooks and crannies’ of models (Winsberg, 2012, p. 130), but that there is no reason why these subjective decisions should be connected to value-judgments. Parallel to this, one could argue that scientists usually make ontologically underdetermined choices on the basis of pragmatic considerations. For example, familiarity with a particular method, the availability of particular software, or the ease with which particular data can be collected are all pragmatic factors that plausibly influence ontologically underdetermined decisions without directly relying on social or moral values. Similarly, it may be that these decisions are influenced by the taxonomists’ mood that day, the weather outside, or even a hangover (Biddle, 2013, p. 131). The point of the objection is that while the AOU shows that it is undeniable that some subjective factors influence species classification, there is no reason to assume that these subjective factors are necessarily social (political, moral, environmental, etc.) value-judgments. Indeed, it seems plausible that many ontologically underdetermined decisions are in practice made on the basis of practical factors without consideration of values.

I hold that this last objection fails for two reasons. First, the distinction between pragmatic factors and social (moral, political, environmental, etc.) value-judgments is hard to hold. One major reason for this is that these pragmatic factors are inevitably tied to economic considerations, which in turn entail social values (Wilholt, 2017). Thus, a decision on the basis of the availability of a particular computer program, the accessibility of data or even a hangover may indirectly be connected to value-judgments. The second problem for the objection is that it only rejects the descriptive form of the AOU. Even if, as a matter of fact, these ontologically underdetermined choices are often determined by pragmatic

²⁷ The arguments from epistemological underdetermination for the value-ladenness of science will be discussed in section 4. Very briefly, they hold that at least in ampliative reasoning, scientific decisions are never fully determined by the evidence, so scientists need to rely on value-judgments to make these decisions under uncertainty.

considerations, the argument that taxonomists *should* take value-judgments into account still stands. In other words, this objection fails to argue against the ideal that values should influence taxonomy, and only shows that perhaps in practice they do not (Steel, 2016).

Assuming that these three objections fail, I hold that taxonomy should be value-laden due to the role of values in choosing between multiple ontologically legitimate options. This conclusion fits well with the increasing consensus in the philosophy of science about the importance of values in science. However, this conclusion is typically drawn from different arguments, most prominently the argument from inductive risk (Douglas, 2000; Rudner, 1953; Steel, 2010), the gap-argument (Howard, 2006; Longino, 1990; Neurath, 1983), and the argument from value-laden concepts (Dupré, 2007a; Putnam, 2002).²⁸ While I believe that these arguments are relevant for taxonomy too, this thesis will primarily focus on the value-ladenness due to ontological underdetermination. Apart from pragmatic concerns of space, I do this for three reasons. First, while there is widespread interest in other arguments for the value-ladenness of science, little attention has been paid to values that enter science through ontological underdetermination. One aim of this thesis is to make a first step towards filling this relative gap in scholarship. Second, and more fundamentally, the value-ladenness due to ontological underdetermination is particularly significant in species classification. This is due to the radical character of ontological species pluralism, and, in relation to this, due to the immense variety of methods of species delimitation that taxonomists rely on. Understanding and evaluating the impact of values is thus particularly important for understanding the practice and results of species classification. Finally, as will become clear in the remainder of this chapter and the next, I will argue that ontological underdetermination causes taxonomy to be value-laden in a particularly profound way. Indeed, one of the main arguments will be that ontological underdetermination causes species classification to be more profoundly value-laden than is established by the commonly used argument from inductive risk and the gap argument. As this profound character of the value-influences has important implications for the way we should think about values in science, it is worth paying extra philosophical attention to the AOU.

²⁸ I discuss the gap argument and the argument from inductive risk in section 4.2. The argument from value-laden concepts holds that values enter science through concepts that are inherently value-laden, such as rape, well-being, or even inflation (Alexandrova, 2016; Dupré, 2007b).

4. Three Views on Values in Science

The previous sections have argued that value-judgments should play a role in constructing and evaluating species classifications. If we think of the taxonomic research process as a series of decisions (choice of research subject, methodology, data collection, analysis, interpretation of the evidence, etc.), this can be understood as the claim that value-judgments should weigh in on at least some of these decisions. Thus, the AOU places values next to the other group of normative considerations, namely, epistemic standards, which are already widely accepted to influence scientific decisions. Of course, the mere claim that values have some role to play in some decisions does not imply that values have the same role as these standards, or that values are equally important. More precisely, epistemic standards are generally accepted to play a particularly profound role in science; it is far from clear whether values should play an equally profound role. Thus, the claim that values have some role to play naturally raises the question of how profound this role is, and how it relates to that of epistemic standards. The aim of the remainder of this thesis is to answer this question.

The question about the appropriate role for values in science can easily be overlooked in what Matt Brown (2017, pp. 5–6) calls ‘ideal cases’, when value-judgments are in line with epistemic standards. Take again the hypothetical example of the taxonomist and the endemic bird species. If the claim that the group is a species turns out to be more empirically adequate than the claim that the group is not a species, then the question about the appropriate role for values does not come up. Given that accepting the group as a species is also in line with the value-judgments of the taxonomist, all factors are pushing the decision in the same direction. According to Brown, the question only comes up in non-ideal situations, when values push in a different direction from epistemic standards. If it turns out that, given the evidence we have, rejecting the bird species is more empirically adequate than accepting it, then there is a clear clash between epistemic standards and value-judgments. In this situation, the taxonomist needs to decide whether epistemic standards or values should be prioritized. Similar clashes between epistemic standards and value-judgments can come up regularly throughout the research process, ranging from ethical concerns about experimental methodologies to scientific outcomes not in line with what scientists had hoped to find. Thus, once it is accepted that values should play some role in scientific decisions, a strategy or framework that tells us

how to deal with these non-ideal cases is needed. Such a framework should clarify what kind of role values should play and what kinds of decisions they can and should influence.

Over the past decades, philosophers of science have proposed various such frameworks for regulating and evaluating the role of value-judgments in scientific decisions. The remainder of this thesis will investigate the implications of the AOU for the choice of such a framework. This section sets a first step in this by taking a closer look at three prominent philosophical frameworks for thinking about values in science. The first view, which will be called the ‘Value-Free Ideal’, holds that values can legitimately weigh in on scientific decisions provided that their influence is limited to the external, non-epistemic stages of science (Lacey, 2005; McMullin, 1982). The second view, which will be called the ‘Epistemic Priority View’ holds that values play a legitimate role in science as long as they do not trump epistemic standards in the internal parts of science (Douglas, 2009; Steel, 2010; Steel & Whyte, 2012). The final view, which will be called the ‘Aims View’, holds that values influence science legitimately as long as they are in line with the aims of science (M. J. Brown, 2013; de Melo-Martín & Intemann, 2016; Elliott & McKaughan, 2014; Kourany, 2010). After presenting each of these views in a separate subsection (3.1, 3.2, 3.3), the section ends (3.4) by taking a closer look at the main points of agreement and disagreement between these views. This will then be the starting point of chapter 4, which considers the implications of the AOU for the choice between these frameworks.

4.1. The Value-Free Ideal

The first and arguably oldest view on values in science is what is commonly called the ‘Value-Free Ideal’. Simply put, this view holds that values should not play a role in the ‘internal stages of science’, i.e. the parts of science that are concerned with making inferences, characterizing the data, and interpreting the evidence. To understand the Value-Free Ideal and, more particularly, its appeal, it will be helpful to return to the example of Wellington and Hoser, the two Australian taxonomists discussed in the introductory chapter. Among other things, these taxonomists meddled with the data and even made up data in order to recognise a high number of ill-supported species. They did this because of the prestige connected to recognising species and to express their discontent with the taxonomic practice in Australia. Unsurprisingly, the work of these two taxonomists is widely rejected by other taxonomists, and their practices are considered a corruption of taxonomy.

This example is relevant here because of the prominent and particular role that values played in the work of these taxonomists. These taxonomists disregarded empirical evidence, and recognised species merely because this was the outcome that they wanted to have. Thus, value-judgments played the role that is usually played by empirical evidence. This practice is usually referred to as ‘wishful thinking’, which occurs whenever scientists accept a claim because they want it to be true, and not because the evidence suggests so (M. J. Brown, 2013, 2017; de Melo-Martín & Intemann, 2016; Douglas, 2009). It is widely accepted among scientists and philosophers that wishful thinking constitutes an illegitimate influence of values in science. Indeed, many prominent examples of values playing an illegitimate role, such as the Lysenko case, are cases of wishful thinking. Unsurprisingly, then, nearly all participants in the values in science debate mention wishful thinking as an important threat of allowing values to play a role in science, and claim that any acceptable view on values in science should rule out the possibility of wishful thinking.

The threat of wishful thinking forms the direct motivation for the Value-Free Ideal as a strategy for regulating values in science. Proponents of the Value-Free Ideal point out that the problem of wishful thinking lies in the fact that scientists rely on value-judgments instead of empirical evidence or epistemic standards to make decisions. This opens the door to accepting hypotheses on the basis of what they want the world to be like, rather than on what it actually is like. To avoid such corruption of science, proponents of the Value-Free Ideal claim, values should be banned from these inferential decisions. This way, empirical evidence, epistemic standards and logical rules of inference can guide decisions without the danger of wishful thinking.

It is important to note that this value-free ideal is not claimed to apply to all scientific decisions, but only to those where value-influences could lead to wishful thinking. For example, taxonomists typically rely on value-judgments when choosing a research subject. Groups that are suspected to be endangered and require formal recognition for increased protection are typically prioritized over other research-subjects. Similarly, value-judgments often come into play when scientists use the results of their research in practical applications. For example, the regulation of trade in organic materials extensively relies on taxonomic research. In such cases, economic, social and political considerations are inevitably entangled

with the classifications they rely on. However, it seems that in these examples values can influence scientific decisions without leading to wishful thinking.

In order to avoid ruling out such harmless and even important value-influences in science, proponents of the Value-Free Ideal crucially rely on a distinction between two parts of the typical research process (e.g. Anderson, 2004; Douglas, 2009, 2016; Kitcher, 2011; Lacey, 2005; Longino, 1990). Some authors speak of 'steps', 'phases' or 'contexts' of research. I will use the terms *research stages* or *stages of science*. There are variations in the number and description of stages that different authors adopt, but the different taxonomies of research stages are all compatible with a distinction between what I will call the *internal* stages and the *external* stages of science. The internal stages of science are concerned with the question of what counts as evidence, the characterization of the evidence, whether and to what degree the available evidence supports the hypothesis, and the acceptance or rejection of hypotheses as scientific knowledge (Douglas, 2009). These aspects of science are often seen as the work of the scientist *qua* scientist, and they are also sometimes referred to as the 'inferential' and 'epistemic' stages, or even as the 'core' or 'heart' of science. The external stages (also 'peripheral' or 'pre- and post-epistemic' stages), on the other hand, simply consist of all aspects of science that are not part of the internal stages. Most importantly, the external stages include the choice of research topics, initial formulation of the hypothesis, choice of concepts, design of the research methods and the potential use of scientific results in applications or further research.

The precise extension of the internal and external stages is the subject of some debate. Various authors have argued that the acceptance and rejection of hypothesis is not part of the internal stages (Elliott & Willmes, 2013; Giere, 2003; Lacey, 2005; Mitchell, 2004). To make this point, they introduce a distinction between two cognitive attitudes that scientists may take towards a scientific claim, namely, belief and acceptance. Belief is then considered to be the purely cognitive attitude of taking a claim to be true. Acceptance, on the other hand, consists in relying on a particular claim both in practical reasoning and further theoretical reasoning. With this distinction in hand, these authors argue that the internal stages are only concerned with determining the truth of scientific claims, and hence, with belief. Acceptance, on the other hand, is closely connected to practical reasoning, and thus may involve considerations about how sure one has to be about a claim before it is suitable as a basis for

action, how much evidence we should collect before acceptance given limited time and funding, and when a claim is suitable as a basis for further research. The point then is that these practical decisions are not part of the work of the scientist qua scientists, and instead could be made by policy-makers or scientists playing that role. For example, scientists could report probabilities (belief), and policy-makers could then decide whether that particular degree of belief is sufficient as a basis for action (acceptance).

In response to these arguments, Daniel Steel (2013) points out that on any reasonable interpretation of the cognitive attitude of belief, it still contains practical components.²⁹ While a purely cognitive notion of belief may be theoretically conceivable, he claims that it does not play any significant role in scientific practice. It follows from this that if we want an interpretation of the internal/external distinction that is helpful for understanding actual scientific practice, we have to adopt an interpretation of the internal stages that involves cognitive attitudes that have practical components. In line with this argument, I will continue to take the acceptance and rejection of hypotheses as part of the internal stages of science. This is also the position taken by the majority of authors in the literature on values in science.

Even if we disregard the issue of the acceptance of hypotheses, there is no consensus on where precisely the border between the internal and external stages lies. Douglas (2009, 2016) defends a narrow interpretation of the internal stages, which limits these to the characterization of the evidence, the interpretation of the evidence, and the acceptance of hypotheses. Steel (2017; see also Steel & Whyte, 2012), on the other hand, adopts a particularly broad interpretation of the internal stages, which adds methodological decisions and the dissemination of scientific results to Douglas' interpretation. These differences are important for the Value-Free Ideal because the profoundness of the role of values in science it allows depends on the extension of the external stages. If we adopt a broader characterization of the external stages, as Douglas does, then the Value-Free Ideal allows values to play a profound role in a larger part of science. If we adopt a narrow characterization, as Steel does, then this role is more limited. As my arguments are not dependent on a particular interpretation of the internal/external distinction, I will not go deeper into the debate about the scope of the internal stages here. As there is relatively wide agreement that

²⁹ Steel makes this argument in different terms. He uses the opposition between a cognitive notion of acceptance ('belief' here) and a behavioural notion of acceptance ('acceptance' here).

at least the characterization of data, the interpretation of evidence and the acceptance of hypotheses are part of the internal stages, I will use this characterization in the remainder of this thesis.

In short then, proponents of the Value-Free Ideal adopt the distinction between the internal and external stages of science, and argue that values should only play a role in the latter. The idea is that values only risk distorting scientific outcomes if they play a role in the part of science where beliefs about the world are inferred from the evidence. There is no reason then, according to proponents of the Value-Free Ideal, why values should not be allowed in the external stages. This conclusion is in line with scientific practice, where value-judgments already play a substantial role in the external stages. The distribution of funding over various branches of science and research subjects, with very little money invested in taxonomic research compared to cancer research or military research, provides one example here. Other examples include the ethical requirements for research involving living organisms, and arguments for censorship of work that can potentially be used for terrorism.

Despite its intuitive appeal and wide usage, the internal/external distinction has not gone uncontested. Various authors (Brigandt, 2015; Elliott, 2012; Hicks, 2014) argue that the distinction is not helpful for thinking about values in science because values that influence decisions in the external stages, such as the choice of research topics and ethical restrictions on methodology, continue to influence the internal stages. I will come back and add to these arguments in chapter 4. However, as the internal/external distinction is presupposed by the Value-Free Ideal and the Epistemic Priority View (which I discuss in the next subsection), I will accept it for now. Assuming then that the internal/external distinction is tenable, the Value-Free Ideal can be characterized as the view that the internal stages of science should be free of values like concerns about conservation and health.

To illustrate the Value-Free Ideal, it is instructive to consider how it would deal with the example of the bird species and the real estate project from the introduction of this chapter. As proponents of the Value-Free Ideal do not deny that values can influence the external stages of taxonomy, the conservation-interests of the taxonomists form a legitimate reason for her to study the group. Similarly, ethical principles of the taxonomist could be a legitimate reason to only rely on tissue that can be collected without harming the organisms. However, given that the Value-Free Ideal does not allow values in the internal stages of science, values

would not be allowed to influence the research process after these decisions. Assume that the taxonomist collected nDNA data, and found weak evidence for gene-flow between the group and a sister lineage living somewhere else. The Value-Free Ideal holds that conservation values are not relevant when deciding whether weak evidence for or against species status is enough for the taxonomists to make a decision. Similarly, the Value-Free Ideal does not consider conservation values a legitimate reason for the taxonomist to recognize the species despite the nDNA evidence, and neither can they form a reason to doubt the validity of the evidence.

4.2. The Epistemic Priority View

Despite its intuitive plausibility, the Value-Free Ideal has been challenged from various angles over the past decades. These challenges all point out that values have some legitimate role to play in the internal stages of science too. Most prominent among these challenges are two arguments based on the fact that some, most, or even all scientific decisions are taken under uncertainty (Biddle, 2013; Steel, 2010). Broadly put, these arguments hold that inductive and other forms of ampliative reasoning leave a gap between the evidence and the scientific decisions taken on the basis of this evidence. Values are then needed to bridge this gap. As this uncertainty is as much part of the internal stages as it is of the external stages, these arguments are claimed to show that values play a role in the internal stages of science. Note that these arguments have the same structure as the AOU, but rely on epistemological rather than ontological underdetermination.

There are two variants of the argument from epistemological underdetermination. The first of these is commonly called the ‘gap-argument’ (Biddle, 2013; M. J. Brown, 2015; Howard, 2006, 2009; Kourany, 2003; Longino, 1990; Neurath, 1983). This argument holds that some scientific decisions are not fully determined by evidence, epistemic standards and rules of inference. Instead, there are often multiple options that seem epistemically equally good. In such cases, the argument goes, values are needed in addition to these epistemic considerations to determine scientific decisions. Its name, ‘the gap-argument’, derives from the fact that values fill the gap between epistemic considerations and theory choice. Like the AOU, the argument has a normative and a descriptive version. In the descriptive sense, it shows that values are needed to bridge the gap between evidence and theory in case of underdetermination. In the normative sense, it states that if theory choice is not determined

by epistemic considerations, we should take the non-epistemic consequences of scientific decisions into account when making these decisions.

The second variant of the argument from epistemological underdetermination is the argument from inductive risk. This argument is often considered the strongest support for the view that values should play a role in the internal stages of science, and has received by far the most attention in the literature (Betz, 2013; Biddle, 2013; de Melo-Martín & Intemann, 2016; Douglas, 2000, 2009; Jeffrey, 1956; Steele, 2012). Although the name of the argument stems from Hempel (1965), it was most famously raised by Rudner (1953). The argument focuses on the fact that because scientific decisions are often underdetermined by the evidence, scientists need to make a judgment about the required degree of certainty. The point then is that this required degree should depend at least partially on the possible consequences of making the epistemically wrong decision. Thus, values should play a role in scientific decisions whenever these consequences are non-epistemic. Like the gap-argument, the AIR has both a descriptive and a normative side. In the descriptive sense, it shows that values are inevitably needed to determine the required degree of certainty. In the normative sense, it holds that whenever there are foreseeable consequences of error, these should be taken into account when choosing the degree of certainty required for scientific decisions.

While these two arguments put pressure on the Value-Free Ideal, they do not by themselves imply an alternative view on values in science. Indeed, taken by themselves these arguments are compatible with a range of different views that accept a role for values in the internal stages. Despite this uncommitting character, these arguments are in fact often used to articulate and support one particular view on values in science. On this view, the epistemological arguments express the *only* role that values can legitimately play in the internal stages. Those who defend this interpretation then commonly couple the arguments to the cautionary claim that wishful thinking threatens whenever we allow values in the internal stages to play any role other than filling the epistemologically underdetermined gap. Thus, like the Value-Free Ideal, these views are motivated by allowing a role for values in science while ensuring that they cannot lead to wishful thinking. The difference from the Value-Free Ideal is that these views hold that there is one way in which values can influence the internal stages without leading to wishful thinking.

To understand these views, it is necessary to clarify what precisely it means to claim that values should *only* play the role supported by these arguments. As discussed above, the positive part of this claim is that values should bridge the gap of epistemological underdetermination. The negative, cautionary part typically comes as the claim that values should never infringe on a particular set of epistemic standards. Whenever the influence of value-judgments leads to breaching these standards, it is claimed, values have played an unacceptable role. Douglas (2013) proposes empirical adequacy and internal consistency as such minimal epistemic standards. Steel (2017; see also Steel & Whyte, 2012) adds to these the severity principle, which holds that ‘a test counts as evidence for a hypothesis only if that test puts the hypothesis at risk of a contrary result’ (Steel & Whyte, 2012, p. 168; see also Douglas, 2009, p. 100; Mayo, 1996). For example, if the toxicity of Bisphenol A, which is associated with the substance’s similarity to human oestrogen, is tested using a strain of rat that is known to be particularly insensitive to oestrogen, then the hypothesis that Bisphenol A is toxic has not been tested severely (vom Saal & Hughes, 2005; Wilholt, 2009). The point here is that Douglas and Steel hold that values can play a legitimate role in the internal stages of science as long as this role is limited to bridging the underdetermination-gap, and as long as values do not interfere with decisions already determined by epistemic standards such as severity of testing, empirical adequacy or internal consistency. Because this implies that epistemic standards take priority over value-judgments whenever they clash, I will refer to these views as ‘Epistemic Priority Views’.

At this point, one might wonder how the positive part (bridging the epistemological gap) and negative part (priority of epistemic standards over values) of the Epistemic Priority View connect to each other. This is easiest to see for the gap-argument, which holds that values enter science in choosing between multiple epistemically good options. If values can only influence decisions that are not determined by epistemic considerations, then it is in principle impossible for them to clash with epistemic standards. If it were the case that one option is epistemically preferable, then the gap-argument would no longer apply. In line with this, Daniel Steel (2010, 2017) relies on the gap-argument to propose a view on the legitimacy of values in science. He argues that values can only legitimately influence scientific decisions in the internal stages when these decisions are not fully determined by epistemic standards. In such cases, he argues, scientists should take the non-epistemic consequences of their

decisions into account and choose the option that is most desirable. Steel (2010, p. 27) points to the trade-off between the speed and reliability of scientific theories as an example. If of two competing theories one produces more reliable results and the other produces results at a much faster rate (thus allowing further research to proceed), it is not directly clear which of these is preferable on epistemic grounds. In that case, scientists must rely on values to choose either; in cases where policy-action is urgent, the speedy theory may be preferable. In cases where erroneous results are very costly, the reliable theory might be preferable. However, if one of the two theories is clearly more reliable and faster, these epistemic considerations always take priority over value-judgments in theory choice. In that sense, the Epistemic Priority View holds that epistemic considerations trump value-judgments.

To explain the connection between the argument from inductive risk and epistemic priority, it is necessary to introduce Douglas' distinction between the direct and indirect roles of values in science. Values play a direct role when they 'serve as reasons in themselves for the choices being made and directly assess those choices' (Douglas, 2016, p. 10). For example, when a scientist decides to study a disease hoping to find a cure, her values have played a direct role in the decision. Values play an indirect role when they are not reasons in and of themselves, but only influence how strong the reasons or evidence for a decision have to be. This distinction fits well with the argument from inductive risk, which shows that values are needed to determine the degree of certainty required for a particular scientific decision.

Douglas takes the direct/indirect distinction as the guiding principle of her view on the legitimacy of values in science. She holds that values cannot influence decisions about how certain we are about a hypothesis, but should determine how certain we should be about something before accepting it. Thus, values can only legitimately influence decisions in the internal stages if they play an indirect role. In this sense, she writes (Douglas, 2016, p. 11),

They [epistemic standards] have what Matthew Brown has called "lexical priority," along with the evidence (Brown 2013a). Once these values have been utilized to assess how much uncertainty we think there is, the other values (social, ethical, *and cognitive*) must compete to help weigh whether the evidence (and its relationship to the theory) is enough.

In other words, it is only after epistemic considerations have done their work that values come in to decide the required degree of certainty. Thus, epistemic considerations are logically

(though not necessarily temporally) prior to value-judgments. Because values only determine the required degree of certainty and cannot affect the epistemic evaluation, it is guaranteed that values do not trump the epistemic and lead to wishful thinking. Indeed, Douglas points out that it is impossible for value-judgments to trump epistemic standards without playing a direct, and hence forbidden, role in the internal stages. This is because these epistemic standards play a crucial role in assessing how strong the evidence for a claim is. Given that the indirect role of values is concerned with deciding how strong the evidence *should* be rather than how strong it actually is, any values trumping these epistemic standards would be overstepping their legitimate domain.

Steel and Whyte (2012) and Elliott (2011a, 2013) argue that the distinction between direct and indirect roles for values is not a helpful basis for a view on values in science. There are at least two related reasons for this. First, allowing free rein to an indirect role for values seems incompatible with Douglas' claim that certain epistemic principles always take priority over value-judgments. Going back to the example from above, imagine that the taxonomist considers the costs of falsely rejecting the bird species extremely high, and hence requires a degree of certainty for rejecting it that is so high that it is nearly impossible to meet. The consequence is that the group will be regarded as a species almost regardless of the test we conduct. Hence, the test would not meet the severity principle, which Douglas also accepts as a basic epistemic standard. Second, this example also shows that in extreme cases the distinction between direct and indirect roles seems to disappear or become irrelevant for evaluating the role of values in science. While at first sight values play an indirect role here (setting the required degree of evidence), it is easy to misuse this as a way of directly influencing the research outcome. In that sense, values would be playing a direct role, as the species would be accepted simply because we want it to be accepted. In other words, giving free rein to an indirect role for values in the internal stages seems to leave the door open to wishful thinking.

While it is thus unclear whether Douglas' indirect/direct distinction is a helpful starting point for an Epistemic Priority View, I will not go deeper into these issues here. As Douglas explicitly writes that epistemic standards have lexical priority over value-judgments (see the quote above), I will continue to take her view to be broadly in line with Steel's. If committing to the direct/indirect distinction means that value-judgments might sometimes trump epistemic

standards in the internal stages of science, then her view should be considered an example of an Aims View, which I discuss in the next section. While this would mean that my current taxonomy of views on values in science is incorrect, it would not harm my arguments.

Let us now take stock of the subsection so far. I have argued that the two epistemological arguments against the Value-Free Ideal lead to an Epistemic Priority View on values in science. This view holds that values play a legitimate role in the internal stages only if they do not interfere with epistemic standards. This means that they can legitimately play the roles exposed by the gap-argument and argument from inductive risk. At this point, there is a potential objection that needs addressing. One might remark that interpreting the epistemological arguments in terms of priority of epistemic standards over values relies on an incorrect interpretation of these arguments. These arguments show that all scientific decisions faced with uncertainty require both values and epistemic standards to be made. Given that both are necessary, it does not really matter which one comes first. Moreover, given that both are necessary for different aspects of scientific reasoning (standards to determine the truth, and values to fill the remaining gap), one could claim that it does not make sense to think of values and epistemic standards as taking priority over each other; rather, they do not seem to compete at all.³⁰

In response to this objection, it is worth emphasising that epistemic standards and values are both normative considerations. Hence, they are capable of functioning in the same kinds of reasoning and push against each other or in the same direction. This response echoes considerations by Steel (2017), who explicitly states that values and epistemic standards both act as reasons weighing in on the same decisions. Indeed, he even holds that sometimes epistemic standards come in the shape of social and political values. Douglas, on the other hand, may be taken to adopt a sharp distinction between the roles of values and epistemic considerations in the shape of the direct/indirect distinction. However, as argued above, while this distinction may seem feasible in theory, it is hard to maintain in scientific practice where the two roles often seem to fade into each other (Elliott, 2011a; Steel & Whyte, 2012).

Second, and regardless of whether we accept Douglas' distinction between direct and indirect roles, I claim it is wrong to think that simply because both factors are taken to play different

³⁰ See Brown (2013, p. 835), who refers to Magnus making a similar remark in personal communication.

roles (assessing the truth – filling the gap), it follows that they cannot clash. When the taxonomists' environmental value-judgments push her to accept the species, and considerations of empirical adequacy push her to reject the species, these factors push in different directions. Even if we concede that epistemic standards are only relevant to deciding whether the group is a species, and value-judgments are only relevant to deciding how strong the evidence should be, it remains the case that empirical adequacy might favour an outcome that goes against our value-judgments. The point is that when values and standards clash, be it in the same or different roles, Douglas and Steel hold that the latter should trump the former. This effectively is an Epistemic Priority View.

Finally, this objection seems to assume a temporal understanding of 'priority'. On this temporal understanding, epistemic priority implies that scientists should first consider epistemic standards and only after that turn to value-judgments. However, proponents of the Epistemic Priority View do not deny that value-judgments concerning inductive risk often precede the collection of evidence and the use of epistemic standards. Rather, priority is used here to refer to the claim that standards should trump values whenever these two push in different directions for a particular decision. For example, if one hypothesis is more empirically adequate while a competing hypothesis is more in line with relevant value-judgments, the Epistemic Priority View holds that scientists should always choose the empirically adequate one.

Having responded to this potential objection, let us briefly consider how the Epistemic Priority View proposes to regulate values in the case of the real estate project and the bird species. Like the Value-Free Ideal, the Epistemic Priority View accepts that the conservation values of the taxonomist legitimately influence her decision to study that particular group. Unlike the Value-Free Ideal, the Epistemic Priority View holds that values have a further role to play. They hold that conservation values can legitimately motivate the taxonomist to minimize the chance of falsely rejecting the group as a species, even if by doing that she increases the chance of falsely accepting the group as a species. The reason for this is that given her set of conservation values, falsely rejecting the species would have particularly grave consequences. Similarly, the Epistemic Priority View also accepts that if the evidence is inconclusive and there is no time or funding to conduct further research, the taxonomists' conservation values form a legitimate reason to accept the group as a species. These roles comprise all legitimate ways

in which, according to the Epistemic Priority View, values can influence the decisions of the taxonomists in the internal stages. Most importantly, if the evidence and epistemic standards suggest that the group is not a species, the Epistemic Priority View does not consider her conservation values a valid reason to doubt this evidence and reject this conclusion.

4.3. The Aims View

While the Epistemic Priority View recognises that values play a crucial role in the internal stages of science, a number of authors have argued that it does not go far enough. Like proponents of the Value-Free Ideal and the Epistemic Priority View, these authors recognise the danger of wishful thinking connected to allowing values in the internal stages of science. Unlike these two other views, they hold that values can play a profound role in the internal stages without necessarily leading to wishful thinking. Indeed, they hold that because of the enormous impact and importance of science in society, it is crucial that values guide science in the right way. They express this as the requirement of a 'joint necessity of evidence and values' (M. J. Brown, 2013, p. 837; Kourany, 2010). They hold that to be fully acceptable, scientific decisions should meet epistemic criteria *and* value standards. When this is not possible, scientists need to compromise either on the values-side or the evidence-side, but neither of these takes unconditional priority over the other. Sometimes scientists might have to settle for an outcome that goes against the accepted values, but in other cases scientists might have to reconsider their methodology, experimental set-up, evidential standards, or their decision to publish the results, because the outcomes do not meet the desired value-standards.

To clarify this view, it is helpful to return to the example of the taxonomist investigating the status of the bird species. Assuming that she collects nDNA-evidence and finds that there is no substantial support for the species-status of the group, she has two choices. One option is to go against her value-judgments, and recognise that the group is not a species. Alternatively, she could at least temporarily reject the results because they are not in line with her value-judgments, and reconsider the type of evidence she collected, the chosen DNA-locus, her method of analysis, or even the chosen species concept. According to views that defend such a joint necessity-thesis, there is no single priority-rule that can correctly guide the use of value-judgments in all scientific decisions. Hence, they argue, the Epistemic Priority View should be

rejected, and replaced by an ideal that allows values to trump epistemic considerations in certain cases.

Various views adopting such a joint-necessity position have been proposed over the past decade. Janet Kourany (2010) argues for an ideal of 'socially responsible science', which holds that any scientific decision should meet high value-standards as well as the appropriate epistemic standards. Matthew Brown (2012, 2013, 2015, 2017) adopts a Deweyan pragmatic position which views science as resolving particular problems, and which holds that scientific success relies on meeting both the value-standards and epistemic standards necessary for a good solution to these problems. Brigandt (2015) argues that scientific decisions should meet broad conditions of adequacy, which comprise both epistemic standards and value-judgments. Similarly, Intemann and De Melo-Martín (de Melo-Martín & Intemann, 2012, 2016, Intemann, 2001, 2005; Intemann & Melo-Martín, 2010) argue in various papers that both values and epistemic standards should weigh in on epistemic justification in science. Finally, Elliott (2013) and Elliot and McKaughan (2014) argue that non-epistemic values play a deep role in science through the many non-epistemic aims that motivate scientific research. Assuming that these non-epistemic aims are no less important than science's epistemic aims, they argue that the normative considerations connected to the epistemic aims should not unconditionally trump the values connected to the non-epistemic aims.

I will not go deeper into the precise arguments for these positions, or the differences between them. Rather, I want to point out that these views are significantly similar. Unlike the Epistemic Priority View, these views do not evaluate values in science on the basis of the kinds of decisions they should influence. Like epistemic considerations, value-judgments are accepted to weigh in on any decision, and what matters instead is whether the right values influence scientific decisions. What counts as 'the right values', proponents of these views hold, depends on the epistemic and non-epistemic aims of science. They argue that the social, moral, political and environmental aims that motivate scientific research should all constrain and determine scientific decisions in a way that ensures that science contributes to the projects, outcomes and aims that society values. In line with the connection between these views and the aims of science, I will refer to them as the 'Aims Views' on values in science.

4.4. A continuum of views

The previous subsections have presented three views on the role of values in science. Before we return to the AOU and consider which of the three views is supported by that argument, it will be helpful to take a closer look at the main points of agreement and disagreement between these views. This will help in clarifying what kind of considerations could be used to argue for or against these views. To start off, it is worth noting that despite their differences, the three views also share two important assumptions about the legitimacy of values in science. First, as pointed out above, all three views accept that values can legitimately influence any decision in the external stages of science, be it the choice of research subjects or decisions concerning the application of scientific outcomes. The Value-Free Ideal even takes this assumption to cover all legitimate ways in which values can influence science; the Epistemic Priority View and Aims View, on the other hand, only take this assumption to cover a part of the legitimate value-influences in science.

Second, the three views also agree that there is a particular kind of role that values should not play in science. That is, they agree that values play an illegitimate role and corrupt science whenever they lead to wishful thinking. Such corruption occurs whenever scientists accept a claim because they want it to be true, and not because the evidence suggests so. Wishful thinking is often connected to the closely related problem of dogmatism, which occurs whenever scientific research is structured so as to lead to a predetermined result. Thus, it constitutes a threat to the ability of science to generate surprising results. The problems of wishful thinking and dogmatism derive from the same deeper worry. In both cases, values are claimed to play a problematic role because scientific claims are not appropriately dependent on what the world is like. Instead, value-judgments are used in the role that empirical evidence ought to play. Proponents of all three views assume that values should not play such a role. I will refer to this assumption as the ‘world-dependency-requirement’. This requirement holds that any scientific claim about what the world is like (or any decision involved in the process of making such claims) should at least partially be dependent on what the world is like, and not merely on value-judgments.³¹

³¹ Note that I write ‘partially’ here to emphasize that the world-dependency criterion does not require that the world fully determines these claims.

Beyond these two common assumptions, the commitments of the three views diverge. First, there is a divide between the Epistemic Priority View and Aims View on one side, and the Value-Free Ideal on the other. While both the Epistemic Priority View and the Aims View hold that values have a legitimate role to play in the internal stages, the Value-Free Ideal holds that these internal stages should remain value-free. More precisely, the Epistemic Priority View and Aims View argue that values can at least play the tie-breaker role. Thus, despite further differences between the Epistemic Priority View and Aims View, the two views also share a common basis rejected by Value-Free Ideal.




Figure 3. A Continuum of Views. This figure represents the various views on values in science discussed in this chapter.

On the basis of these points of agreement and disagreement, we can represent the three views on a continuum that shows the profoundness of the role for values allowed for by each view. The red arrow close to the right end of the continuum indicates the world-dependency criterion as an assumption that indicates the point beyond which views allow values to play the same role as empirical evidence. Given that all three views accept the world-dependency criterion, they are all placed to the left of this. One possible view further to the right is a hypothetical ‘anything goes’ view, which holds that values can play any role whatsoever, including wishful thinking. Similarly, constructivists who believe that the world depends significantly and directly on what we think about the world, may be placed to the right of this. On the other side of the continuum we find the Value-Free Ideal, which holds that values should only play a role in the external stages. The only view further to the left is a hypothetical naïve Value-Free Ideal, which holds that values have no role at all to play in science. To the right of the Value-Free Ideal is the ‘tie-breaker role’, which is assumed by both the Epistemic Priority View and Aims View, but rejected by the Value-Free Ideal. It indicates the point on the

continuum beyond which all views at least allow values to act as a tie-breaker in the internal stages.

Together with the world-dependency criterion, the tie-breaker role can be understood to delimit the space within which the debate between the Epistemic Priority View and Aims View takes place. Both views agree that any acceptable view on values in science should be positioned somewhere between these two points (indicated in red), but disagree on where precisely this should be. The Epistemic Priority View is placed very closely to the tie-breaker point, as it claims that this is the only role that values should play in the internal stages of science. More precisely, it claims that values should only influence decisions in the internal stages if these are not determined by basic epistemic principles. The Aims View, on the other hand, is placed closer to the world-dependency criterion, as it holds that values can play any role that is in line with the aims of science and that does not hinder the world-dependency of the decisions.

This continuum provides a clear representation of the debate between these three views and at the same time reveals what is needed for the AOU to support one of these views or argue against them. If the AOU only shows that values should play a role in the external stages of science, then it provides direct support for the Value-Free Ideal. If the AOU shows that values can influence the internal stages of taxonomy by breaking ties between epistemically equally legitimate options, then it provides direct support for the Epistemic Priority View and forms an objection against the Value-Free Ideal. Finally, if the AOU shows that values should stand on equal footing with epistemic standards in the internal stages of science, then it provides direct support for the Aims View and forms an objection against the Value-Free Ideal and the Epistemic Priority View. In short then, the debate between the Value-Free Ideal and the other two views hinges on the question of whether the internal stages can be value-free. The debate between the Aims View and Epistemic Priority View hinges on the question of whether epistemic principles should take priority over value-judgments in the internal stages.

The next chapter will argue that the AOU has important implications for both questions. I argue that it answers both questions negatively: the internal stages cannot be value-free, and epistemic standards should not take unconditional priority over value-judgments. Before we turn to these arguments, it is worth pointing out here that the debate between the Epistemic Priority View and the Aims View is sometimes framed differently. While the emphasis in this

thesis is put on whether epistemic standards should trump value-judgments in the internal stages, one could also focus on the question of whether value-judgments can provide *epistemic* reasons to weigh in on scientific decisions (de Melo-Martín & Intemann, 2012, 2012, 2016, Intemann, 2001, 2005, 2015). Proponents of the Epistemic Priority View then claim that they cannot, while proponents of the Aims View claim that they can.

I believe there is no profound difference between these two ways of framing the debate. Both the Aims View and the Epistemic Priority View seem to agree that the aim of the internal stages of science is to further certain non-epistemic and epistemic goals by advancing knowledge. Proponents of the Epistemic Priority View hold that advancing knowledge only requires satisfying certain epistemic standards. Hence, they hold that the only necessary criteria for adequate science are these epistemic standards. It follows from this that avoiding the corruption of science requires that we prioritize these standards over value-judgments. Proponents of the Aims View, on the other hand, believe that advancing knowledge requires satisfying both these epistemic standards and certain value-judgments. Hence, they hold that the necessary criteria for good science are broader than just epistemic standards. As this places values and standards on equal footing, it follows that there cannot be unconditional priority of one over the other. This means that sometimes value-judgments can trump epistemic standards.³² In order to avoid discussing the meaning of the term ‘knowledge’, the remainder of this thesis will frame the debate only in terms of values trumping epistemic standards.

5. Summary

This chapter has considered the implications of radical pluralism for the role of values in species classification. The first three sections considered the question of whether values should play a role in species classification. I argued that values should and inevitably do play a role in the choice between many ontologically legitimate options. The fourth section then turned to the question of what kind of role for values the argument from ontological underdetermination establishes. More precisely, the question is how profound this role

³² The use of ‘epistemic standards’ is confusing here, because what is at stake is precisely the meaning of the term ‘knowledge’. Here, I use the term (in line with my usage of the term in the rest of the thesis) to refer to those principles that the Epistemic Priority View considers as the minimal requirements for adequate science (empirical adequacy, severity principle, consistency). It should not be understood as ‘those standards that indicate knowledge’, given that the two sides of the debate understand ‘knowledge’ in different ways.

should be. In order to answer that question, the fourth section presented three prominent views on values in science. The Value-Free Ideal holds that values should only influence decisions in the external stages of science. Arguing against this, both the Epistemic Priority View and the Aims View hold that values have a legitimate role to play in the internal stages too. The former holds that they can do this as long as they do not trump epistemic considerations. The latter argues for a stronger role and holds that values and epistemic considerations stand on equal footing in guiding scientific decisions. The next chapter will investigate the implication of the AOU for the choice between these frameworks.

Chapter 4: Against the Wedge Strategy

The previous chapter introduced the argument from ontological underdetermination (AOU), and presented three frameworks within which to think about the legitimacy of these value-influences in science. These three frameworks differ in how profound a role for values they allow. This chapter considers the implications of the AOU for the choice between these three frameworks. More precisely, I investigate which of these three frameworks is most suitable for thinking about values that enter science through ontological underdetermination. This, in turn, will directly reveal how profound a role these values play in taxonomy.

At first sight, one might think that the AOU supports the Value-Free Ideal. Proponents of the Value-Free Ideal could point out that the choice between multiple ontologically legitimate classifications is made when taxonomists choose their topic of research. Some taxonomists choose to study interbreeding groups, while others choose to investigate ecological groups, units of evolution or genealogical groups. Alternatively, they could also say that this choice is made when the results of taxonomy are applied, for instance, when conservation biologists, ecologists and policy-makers decide whether they should rely on phylogenetic species, units of evolution, or other kinds of groups delimited by taxonomists. Either way, they could argue, the ontologically underdetermined decision is not made by the taxonomists *qua* scientists as part of the internal stages of taxonomic research. Instead, it is a non-epistemic decision made in the external stages, where values can and should legitimately do their part.

Alternatively, one might think that the AOU supports the Epistemic Priority View. This is because the role for values established by the AOU shows significant similarities with that established by the epistemological arguments. In both cases, values act as tie-breakers between multiple options that are epistemically equally good. In the case of the epistemological arguments, the decision is between multiple equally supported options (gap argument), or between multiple trade-offs between false positives and false negatives (inductive risk). In the case of the AOU, the decision is between multiple arguably true options. In all these cases, it seems that values only come in when the epistemic does not fully determine the decision. As the Epistemic Priority View holds that the only role for values in

the internal stages lies in tie-breaking epistemically equally good options, the AOU seems to fit well with this view.

This chapter will argue that, despite these *prima facie* plausible ways in which the AOU seems to fit nicely with the Epistemic Priority View and Value-Free Ideal, it actually supports the Aims View. In other words, I argue that the AOU supports a role for values in science that is more profound than what is considered legitimate by the Value-Free Ideal and Epistemic Priority View. I will do this by bringing to light and criticizing two crucial assumptions of the Epistemic Priority View and Value-Free Ideal, which I will refer to as ‘the first wedge’ and ‘the second wedge’. This terminology is in reference to Matt Brown (2015, p. 62), who points out that these views adopt a ‘wedge-strategy’ for thinking about and regulating values in science. Simply put, both frameworks propose to think about values in science by driving a wedge between the part of science that is directly concerned with finding truth, and all the rest of science. This wedge then acts as a limit for value-influences: values can legitimately play a role in all parts of science except in the epistemic core.

Section one presents the first wedge, namely, between the internal and external stages of science, and shows how it is assumed by the Value-Free Ideal and the Epistemic Priority View. The second section argues that this wedge is not tenable due to implications of the AOU. The third section then presents a second wedge, between epistemic and non-epistemic reasons, and the fourth section argues that this wedge too is untenable due to implications of the AOU. Section 5 considers a potential response from the Epistemic Priority View, and argues that it reduces the view to a trivial position. Section 6, finally, concludes the chapter and suggests that the AOU also forms a positive argument for the Aims View.

1. The First Wedge: Internal and External Stages

The first assumption of the Epistemic Priority View and the Value-Free Ideal that I want to bring to light is the distinction between the internal and external stages of science.³³ This distinction is reminiscent of the classic distinction between the context of justification and the context of discovery, but should not be taken to imply a commitment to the internal stages being the only proper subject of the philosophy of science. Typically, the internal stages are taken to include the characterization and interpretation of evidence, the evaluation of

³³ I have already discussed this distinction in chapter 3, section 4.1.

theories on the basis of this evidence, and the acceptance and rejection of hypotheses as scientific knowledge (M. J. Brown, 2013; Douglas, 2009 chapter 5; Hicks, 2014). The external stages, on the other hand, simply cover all the rest of science. Traditional examples of scientific decisions in the external stages include the choice of research topics, the ethical permissibility of experiments, and the way scientific results are applied. Douglas (2009) explicitly states that this distinction reflects the temporal order of the typical research process, with the external stages preceding and following the internal stages. However, this temporal interpretation should not be taken too strictly, as in scientific practice the two are often mixed up (Elliott, 2012; Hicks, 2014).

The internal/external distinction plays a crucial role in how the Value-Free Ideal proposes to regulate values in science. It drives a wedge between the external and internal stages, and claims that values should be restricted to the former. This results in a superficial role for values, because they are banned from all scientific decisions concerning the interpretation of evidence and the evaluation of hypotheses. Indeed, one could argue that the choice of research topics, decisions concerning ethical restrictions on experiments and decisions concerning the application of science in society are the business of funding bodies, ethicists and policy-makers, and consequently need not enter the scientists' minds unless they take up these roles. On the Value-Free Ideal value-judgments would then be banned from the decisions of the scientist *qua* scientist.

The two epistemological arguments, which form the basis of the Epistemic Priority View, suggest that the Value-Free Ideal's simple application of the wedge-strategy fails to appropriately regulate values in science. This is because scientists sometimes have to make decisions in the internal stages that are not fully determined by purely epistemic considerations. Proponents of the gap-argument point out that values can act as an arbiter in the internal stages when there are various options that are equally supported by the evidence and epistemic standards. Proponents of the argument from inductive risk argue that values have to set the degree of certainty required for scientific decisions. Because these arguments show that values should play a role in the internal stages of science, they seem to go against the first wedge. However, closer consideration of these arguments and the Epistemic Priority View shows that rather than refuting the first wedge, they incorporate it. While these arguments reject the internal/external distinction as the border between the value-laden and

the value-free parts of science, they do not reject the distinction altogether. Instead, the internal/external distinction is adopted as a distinction between the parts of science in which values can play any role whatsoever (the external stages), and the parts of science where the role of values should be constrained (the internal stages).

This wedge is clearly part of Steel's and Douglas' views on values in science. Steel and Whyte (2012, p. 164) explicitly state that epistemic priority should only apply within the range of what is 'ethically permissible and practically feasible'. Moreover, they do not include the application of scientific results and the choice of research subjects in their overview of the scientific decisions in which the epistemic should take priority. Similarly, Douglas (2016, p. 10) explicitly states that values should play a direct role in the choice of research subjects, the application of scientific results, and even methodological choices. She holds that the role of values should be restricted to an indirect role only in the internal stages of science. Thus, like the Value-Free Ideal, the Epistemic Priority View holds that values should have free rein in the external stages, while further restrictions apply in the internal stages.

2. Problems for the Internal/External Distinction

This section argues that the distinction between the internal and external stages is not helpful for thinking about the values that enter taxonomy through ontological underdetermination. It does this by discussing three examples of value-judgments in the external stages. The point of these examples is to show that values that enter in the external stages continue to have profound influence on scientific decisions in the internal stages. Thus, unless we want to ban values from the external stages, it is hard to deny that values can legitimately influence decision in the internal stages too. This poses a problem for any view, like the Value-Free Ideal and the Epistemic Priority View, that relies on the internal/external distinction for regulating values in science.³⁴

As pointed out in chapter 3, there is some debate about the precise extent of the internal and external stages (Douglas, 2016; Hicks, 2014; Steel & Whyte, 2012). More precisely, while some authors take methodological choices and the dissemination of results to be part of the internal stages, others believe these are part of the external stages. I will not take a position in this

³⁴ Note that I do not aim to argue that the internal/external distinction is unviable *tout court*. I only aim to show that it is not helpful for thinking about values in science. For an argument similar to the one here, see Brigandt (2015).

debate. Instead, I start with examples of value-influences that are uncontroversially part of the external stages, namely the choice of research topics (2.1) and the ethical constraints on scientific research (2.2). Section 2.3 concerns a more controversial example, namely, values that influence the dissemination of the results of scientific research. However, this lack of consensus does not affect the argument. If, contrary to what I assume here, decisions about the dissemination of results are part of the internal stages, the arguments in 2.3 still form a problem for the Value-Free Ideal and the Epistemic Priority View as they illustrate a direct role for values in the internal stages.

2.1. Selection of research topics

The selection of research subjects is a prime example of a scientific decision in the external stages in which values play a prominent role (Kitcher, 2001). This is acknowledged by proponents of all three views on values in science. Due to the unfortunate combination of a high number of undiscovered species and a pressing biodiversity crisis, taxonomy is an excellent example of this role for values in science. Estimates of the proportion of undiscovered species range between 90% and 50% even if prokaryote species are not taken into account (Mora, Tittensor, Adl, Simpson, & Worm, 2011). In addition to this, there are longstanding debates about the classification of many of the species that have already been described. Given that taxonomists have to select their research topic from a practically endless range of topics, it is clear that this decision is not fully determined by the world and instead requires value judgments. For example, legal protection of a group of organisms depends on that group of organisms being recognized as a taxon. Hence, conservation values motivate us to study those groups that require urgent protection. Other plausible value-influences on the choice of research topics include the charisma of a species, familiarity with the group through previous work or education, the usefulness of the species for commercial purposes and agriculture, and practical factors such as the ease of data collection (see e.g. Giangrande, 2003; Isaac et al., 2004; Linnean Society, 2014; J. R. Wilson et al., 2007).

I argue that the value-judgments that influence the choice of research topics sometimes continue to influence decisions in the internal stages of taxonomy. An interesting example of this comes from the debates about taxonomic inflation, which I discussed in the introductory chapter. Harris and Froufe (2005) argue that a historical geopolitical bias in the choice of research topics is one of the causes of the recent strong increases in species numbers. They

point out that species are often delimited on the basis of the genetic distance between groups of organisms. Very simply put, two groups of organisms are considered distinct species if the genetic distance between them reaches a particular threshold. Harris and Froufe show that this threshold is set on the basis of taxonomic work on groups in North-America and Western Europe. The reasons for this are simple: when DNA-based methods first came to be used for species delimitation and were still very expensive, there were simply more taxonomists and funding opportunities in these wealthier parts of the world, and taxonomists mostly focused on organisms that were readily available nearby. While these value-based influences are by no means illegitimate, they have strongly influenced decisions in the internal stages of science. Both North-America and western Europe are genetically depauperate because of relatively recent recolonization after the most recent glacial period. Because genetic diversity is lower in these northern regions, the genetic distance threshold for species status was set rather low too. When applied to diversity-rich tropical areas, these criteria result in high numbers of new species.

The point here is that because taxonomists focused on groups living in diversity-poor areas, they adopted criteria for species status that pick out small, recently diverged groups as species. Because value-judgments directly influenced the choice of research subject, they are directly linked to setting the thresholds too. That is, if taxonomists had preferred to study diversity-rich areas, then their criteria might have picked out a different kind of group, namely, larger, more distinct groups instead of small, recently diverged groups. These criteria and their application are clearly part of the internal stages of taxonomy, as they determine how the DNA-evidence is interpreted and which classification is accepted as scientific knowledge. Thus, legitimate value-influences due to ontological underdetermination in the external stages influence decisions in the internal stages too.

2.2. Ethical constraints

A second way in which values legitimately affect the external stages is through ethical constraints on the design of research methodologies. For example, even though we could benefit epistemically from testing potentially carcinogenic substances on humans instead of animals before commercialization, it is clear that this is not permissible for ethical reasons. Similar considerations play a role in taxonomic research. One clear example comes from the taxonomic debate on *Discoglussus jeanneae* and *D. galganoi*, which I have referred to in

previous chapters and which I will continue to use for most of the examples in this chapter. In order to investigate the relation between *D. galganoi* and *D. jeanneae*, Fritz et al. (1994) compare the DNA content of the two taxa. They motivate their choice for this research method by saying that '[s]pecies determination by means of DNA content can be carried out without killing specimens' (Fritz et al., 1994, p. 143). This distinguishes their research from earlier research carried out by Busack (1986), which required killing specimens for allozyme data.

The value-laden decision to rely on research methods that do not require killing organisms has significant influence on the outcome of the research. Fritz et al. (1994) find no significant difference between the DNA content of *D. galganoi* and *D. jeanneae*. Hence, they conclude, these two groups should not be considered as distinct species. Their conclusion goes against later research on the basis of mtDNA-sequences (García-París & Jockusch, 1999), and earlier research on the basis of allozymes and morphometrics (Busack, 1986; Capula & Carti, 1993). This suggests that the ethical restrictions on the methodology impact whether *jeanneae* and *galganoi* are recognized as distinct species or lumped together. As this concerns the acceptance or rejection of the species-hypothesis, this is arguably part of the internal stages of taxonomy. Again, this shows that legitimate value-influences in the external stages have continued impact on decisions in the internal stages too.

2.3. Dissemination of results

The two previous examples both concerned decisions in the external stages that generally take place before scientists turn to decisions that are part of the internal stages. The example in this subsection concerns decisions that take place after the internal stages, namely, the publication and dissemination of scientific results. It is clear that values play a role in these decisions. For example, one could argue that for moral or social reasons, a scientist should refrain from publishing research that could be interpreted to justify or encourage racist attitudes. Similarly, if the scientist were to publish this research, there would be good reasons to phrase and explain the research in such a way that minimizes the chance that racist conclusions will be drawn from it. Values play a similar role in taxonomic research. For example, Veló-Anton et al. (2008) presented a study on the basis of mtDNA and nDNA in *D. galganoi* that uncovers two deeply divergent mtDNA lineages and a single nDNA lineage. Thus, they face the ontologically underdetermined choice between recognising two mtDNA lineages

as species or recognising a single nDNA lineage as a species. They point out that while the decision is epistemically arbitrary, it is best to maintain the status of the two taxa as species because that is the way they currently function in conservation legislation in Spain (Velo-Antón, Garci-Paris, et al., 2008, p. 532). Changing their status might alter their level of protection, or might require legal changes that cause a delay in the protection of these groups. These value-judgments about the dissemination of the results did not impact the internal stages of Velo-Antón et al.'s research, as they were made at the very end of the research process. However, these value-judgments did influence scientific decisions in the internal stages of subsequent research relying on Velo-Antón et al.'s results. Various later studies used the species-status of *jeanneae* and *galganoi* as data through measures such as species richness and species density (S. B. Carvalho, Brito, Pressey, Crespo, & Possingham, 2010; Sillero, Brito, Skidmore, & Toxopeus, 2009; Trochet et al., 2014). Similarly, a study by Romero et al. (2014) on a related taxon investigates the impact of taxonomic uncertainty on the results of distribution forecast models. They used two different classification-hypotheses of the heavily debated taxa *Triturus pygmaeus* and *T. marmoratus*, and obtained significantly different results. This shows that publishing the results of a scientific study in a particular way can end up influencing the decisions in the internal stages of subsequent research. Thus, the legitimate influence of values on the dissemination of results directly impacts scientific decisions in the internal stages of taxonomy.

2.4. Failure of the internal/external distinction

The examples in the previous subsections suggest that values that influence the external stages of taxonomy due to ontological underdetermination also continue to influence decisions in the internal stages of the same or later research. In other words, value-influences simply are not restricted to the stage of research in which they enter, but influence decisions across the various stages of research. This poses a problem for views like the Epistemic Priority View and the Value-Free Ideal, which rely on this wedge for the regulation of values in science. If values can play any role whatsoever in the external stages and then influence the internal stages through the decisions made in the external stages, it does not make sense to put restrictions on values in the internal stages that do not apply to values in the external stages. This would just result in arbitrarily regulating values on the basis of when they first entered the research project rather than on the extent and nature of the influence they have.

One could object here that this point is trivial. For example, the decision to study frogs rather than birds, which arguably was determined by value-judgments, also played a role in rejecting or accepting *galganoi* as a species. But clearly, the thought goes, this influence is not of the profound kind that should worry proponents of the Value-Free Ideal or Epistemic Priority View. The point here is that any scientific decision upstream, be it value-free or value-laden, can be said to play some role in all scientific decisions downstream, and consequently also in the results.

However, there is an important difference between the examples discussed here, such as the choice for DNA content on the basis of ethical considerations, and an example such as the choice of frogs over birds. To explain this difference, it is helpful to consider the following three decisions and the various outcomes of the decisions:

Decision **Ex1**: Use mtDNA (**Ex1A**) or DNA content (**Ex1B**) for species delimitation?

Decision **Ex2**: Study birds (**Ex2A**) or study frogs (**Ex2B**)?

Decision **Int**: Recognise *galganoi* as a species (**IntA**) or as a subspecies (**IntB**)?

The Proponents of the Value-Free Ideal and Epistemic priority view hold that decisions **Ex1** and **Ex2** should be determined by value-judgments, while decision **Int** should be determined by epistemic considerations only. The point I want to make here is that decisions **Ex1** and **Ex2** are interestingly different with respect to the way they (and the value-judgments involved in making the decision) influence decision **Int**. The reason for this lies in the relation between these decisions. While the choice between **Ex2A** and **Ex2B** can be said to have contributed to choosing **IntA** or **IntB**, it is clear that it does not make either of these two latter choices more likely. Whether we recognise *galganoi* as a species or subspecies is entirely independent from whether we decide to study the group in the first place. The relation between **Ex1** and **Int** is different. While **Ex2** is impartial towards **IntA** and **IntB**, the outcome of **Ex1** pushes the result to either of these. Thus, the value-judgments involved in **Ex1** can be said to play a more profound role in the internal stages than the value-judgments involved in **Ex2**.

Assuming then that the objection fails, I hold that the arguments in this section show that the value-judgments in the external stages sometimes continue to influence scientific decisions in the internal stages. This poses a problem for any view that allows values to have a profound influence in the external stages, but not in the internal stages. In other words, these

arguments show that the Value-Free Ideal, which is built around this distinction, is not helpful for regulating values in science. Moreover, they also put pressure on the Epistemic Priority View, which is also dependent on this distinction. Indeed, the failure of this distinction presents proponents of the Epistemic Priority View with a choice between Scylla and Charybdis; either they should constrain the role of values in both the internal and external stages, or they should abandon the restrictions in the internal stages. The former is undesirable because it would go against philosophical consensus and widespread successful scientific practice, and the latter is undesirable for the Epistemic Priority View because it would effectively mean adopting the Aims View.

3. The Second Wedge: Epistemic and Non-Epistemic Reasons

The previous section argued that the distinction between the internal and external stages is not helpful for evaluating the role of values in science. Given that the Value-Free Ideal depends on the possibility of the internal stages to be free of value-judgments, this argument implies that this ideal is not helpful for regulating values in science. While the failure of the internal/external distinction also puts pressure on the Epistemic Priority view, there is a seemingly plausible response available for proponents of this view. They might argue that even though the examples in the previous section show that values from the external stages infiltrate the internal stages too, it is still feasible to recognise and identify their influence. This in turn would allow scientists to consider epistemic considerations before these value-judgments influence their decisions. For example, they could point out that we can reconsider the empirical adequacy of the threshold of genetic distance independently from the values that influenced it. Similarly, they might consider whether ethical restrictions on methodology made us rely on a data-type that is not informative. If it turns out that these value-judgments do not clash with epistemic standards, they form no threat to the Epistemic Priority View. If they do clash, proponents of the Epistemic Priority View can simply claim that the epistemic considerations should take priority.

This section presents and illustrates this strategy for dealing with values that enter the internal stages of taxonomy. I do this using table 1 (see p. 125), which lists four important decisions that are part of the internal stages of the taxonomic research process: the choice of data-type, operationalization of the data-type, method of analysis, and interpretation of the results. I argue that these decisions require value-judgments because they are ontologically

underdetermined, and elaborate briefly on how the Epistemic Priority View proposes to deal with these value-influences in the internal stages.

The influence of value-judgments due to ontological underdetermination is perhaps the clearest in the choice of data-type (D_1); a quick look at the scientific debate on the classification of *D. galganoi* and *D. jeanneae* suggests that there is a wide range of data-types that taxonomists can choose from. Over the course of about thirty years, taxonomists have used morphometrics (Busack, 1986; Capula & Carti, 1993), allozyme data (Busack, 1986; Zangari et al., 2006), nDNA (Gutiérrez-Rodríguez, Salvi, Geffen, Gafny, & Martínez-Solano, 2014; Pabijan et al., 2012; Velo-Antón, Garci-Paris, et al., 2008), mtDNA (Fromhage et al., 2004; García-París & Jockusch, 1999; Martínez-Solano et al., 2004), karyological data (Pabijan et al., 2012), DNA content (Fritz et al., 1994), acoustic data (Glaw & Vences, 1991; Vences & Glaw, 1996), ecological data (Capula & Carti, 1993; Real, Barbosa, & Garci-Paris, 2005), and biogeographic data (Busack, 1986; García-París & Jockusch, 1999; Vences et al., 2014). Given that these data-types track different parts of the relevant causal structure of the world, it follows that the choice between these is not fully determined by what the world is like.

Like D_1 , the choice of operationalization of the selected data-type (D_2) is ontologically underdetermined. By operationalization of the data-type I mean the particular parts or aspects of the chosen data-type that taxonomists decide to focus on. For example, if mtDNA is the data-type of choice, taxonomists have to decide which genes to study. Similarly, if acoustic data is chosen, taxonomists have to decide which aspects of mating calls to focus on, such as intensity, duration, intervals, and variation due to the temperature in the environment. Finally, the environment of organisms consists of innumerable ecological variables, and taxonomists studying ecological species have to decide which of these to include in their research. The point is that even within one data-type, different operationalizations of that data-type often track different causal features of the organic world. Assuming that these are all equally real, the choice between them is ontologically underdetermined. Different genes reflect different parts of the genealogical history of the investigated organisms, different variables from the environment are connected to different selective forces, and different morphological characters might be adaptations to different factors in the environment.

Third, the choice of a method of analysis (D_3) is also ontologically underdetermined. Most importantly, there are countless methods for analysing DNA-data that often lead to different results for the same data (Satler, Carstens, & Hedin, 2013). Part of these differences can be explained by the fact that different methods tend to pick up on different causal processes or kinds of groups.³⁵ This is most particularly so for model-based approaches to species delimitation, because these rely on the assumptions of the model of evolution they adopt. Take the increasingly popular methods using the multi-species coalescent model (Camargo & Sites, 2013; Fujita, Leaché, Burbrink, McGuire, & Moritz, 2012; Leavitt, Moreau, & Lumbsch, 2015). These are methods that, on the basis of a model of evolution, describe the likelihood of various possible histories of a set of alleles under investigation. Simply put, these methods track independently evolving lineages by using population genetic models to trace the genealogy of alleles backward in time. Most of these coalescent-based methods assume neutrality of the loci they study, and assume the absence of gene-flow between lineages. The result is that they tend to trace allopatric speciation through drift rather than sympatric speciation through strong selective pressures. Similarly, the commonly used GMYC-model tends to recognise well-supported haplotype-clades as independently evolving, while many other methods would interpret these clades as population structure and lump them into larger groups (Leavitt et al., 2015; Satler et al., 2013). Choosing between various methods thus implies tracking different parts of the relevant causal structure of the world.

Finally, taxonomists are faced with ontological underdetermination in the interpretation of the results (D_4). Because many of the methods of analysis referred to in the previous paragraph come with built-in criteria for species status, D_4 and D_3 are in practice often collapsed into one complex decision. I distinguish between these two decisions here because D_4 constitutes a source of values in taxonomy that is significantly different from D_3 . Simply put, while in D_3 taxonomists decide which causal structures to focus on, D_4 concerns the question of what the implications of these causal structures are for the taxonomic status of the group under investigation.

³⁵ The choice of a method of analysis also seems heavily underdetermined in an epistemological sense. That is, there are countless slightly different methods in use, and it is often unclear which of these is best in any particular research project. This epistemological underdetermination is an additional source of values in this step of research.

This ontologically underdetermined decision is clearly illustrated by the case of *D. jeanneae* and *D. galganoi*. On the basis of the numerous studies and data-types employed over the past decades, a complex image of the relations between the two groups emerges. Participants in the debate generally assume that the two groups formed a single species around 5 million years ago, after which a number of processes of divergence and cohesion produced the complex pattern that we find today. Let us start with the causes of divergence. Morphological differences and differences between the typical environment in which both groups are found suggest that divergence through selection-driven, ecological speciation may be ongoing (Capula & Carti, 1993; Real et al., 2005). It should be noted, however, that, independently from the differences between the two groups, there might also be a continuous cline that explains part of these morphological differences (Capula & Carti, 1993; Vences & Glaw, 1996). The largely allopatric or parapatric distribution and divergence in neutral mtDNA markers suggest additional divergence through genetic drift (Velo-Antón, Garci-Paris, et al., 2008). Finally, Vences et al. (2014) report numerous findings of mitochondrial genes of both groups within a single specimen in the contact zone of the two groups. This suggests either mitochondrial genes that moved to the nuclear genome, or the occurrence of more than one set of organellar genome within a specimen. While this could point towards the selective advantage of these mitochondrial genes, it could also be a sign of intragenomic conflict, and hence the role of conflictual speciation in the divergence between the two groups (Crespi & Nosil, 2013). It is also worth noting that these processes of divergence do not only affect the mitochondrial lineages typically identified as *jeanneae* and *galganoi*. For example, Martínez-Solano et al. (2004) show that there are also at least three mitochondrial lineages within *galganoi* that show divergence through drift.

On the other hand, there are various processes that cause cohesion between *galganoi* and *jeanneae*. Most importantly, there is evidence for current gene-flow over the contact zone, and specimens with the mtDNA of either of the groups have been found in the same population (Gutiérrez-Rodríguez et al., 2014; Velo-Antón, Garci-Paris, et al., 2008; Vences et al., 2014; Zangari et al., 2006). Such gene-flow is likely to have occurred repeatedly over the history of the two groups, probably intermitted by repeated isolation in glacial refuges (Fromhage et al., 2004; Velo-Antón, Garci-Paris, et al., 2008). This is likely to have produced a complex pattern of shared nuclear haplotypes with a weak geographical pattern (Velo-Antón,

Garci-Paris, et al., 2008). The gene-exchange generally does not extend to mitochondrial genes, which show two deeply divergent lineages and a genetic distance that, even though low compared to other *Discoglossus* species, is sufficient for species status. The discordance between mtDNA and nDNA lineages is likely to be a combination of male-biased dispersal, reduced viability of female hybrids (due to Haldane's rule), and slower sorting of ancestral polymorphisms in the nDNA due to the effective population size of nDNA being four times larger than that of mtDNA (García-París & Jockusch, 1999).

It is clear, then, that the relations between *D. galganoi* and *D. jeanneae* are highly complex. Assuming that the above picture correctly reflects the various processes at work in the relation between these groups, it follows that the world does not fully determine their classification. Depending on which of these processes taxonomists focus on, different classifications result. Accordingly, about half of the taxonomists engaged in this debate have recognised the two groups as species, while the other half concludes that subspecific status is more appropriate.³⁶

Given that the four decisions described in this subsection are all ontologically underdetermined, the AOU implies that value-judgments will be required for making them. The Epistemic Priority View proposes to evaluate and regulate the role of these values by distinguishing between epistemic and non-epistemic reasons weighing in on scientific decisions. Values then play a legitimate role if they only influence decisions that are not fully determined by epistemic reasons. To illustrate this strategy, the second and third columns of table 1 (see p. 125) present a selection of respectively epistemic and non-epistemic reasons directly relevant to making the four ontologically underdetermined decisions. The epistemic reasons mostly refer to background knowledge and explanations for observations. Thus, they are based on the epistemic standards such as inter-theoretic coherence and empirical adequacy. The non-epistemic reasons rely on value-judgments such as conservation concerns, ethical constraints, and practical considerations. According to the Epistemic Priority View, values play a legitimate role as long as the reasons in column 3 do not override the epistemic reasons in column 2.

³⁶ For another excellent illustration of the impact of ontological underdetermination on the interpretation of results, see Willis (2017).

Decision	Epistemic	Non-Epistemic
D₁ : ‘What data-type should we collect?’	<ul style="list-style-type: none"> a. Data-type X reliably tracks species in the same genus. b. Background knowledge about the physiology of this group suggests that data-type X is reliable. 	<ul style="list-style-type: none"> c. We are interested in recently diverged lineages, and data-type X tends to track these. d. Official species recognition requires a description in terms of data-type X. e. Data-type X has been collected before, and is accessible. f. Data-type X can be collected without killing specimens.
D₂ : ‘How should we operationalize the chosen data-type?’	<ul style="list-style-type: none"> a. Gene X has a suitable rate of change for tracking genealogical history in recently diverged groups. b. DNA-fragment X is neutral/under selection. c. Environmental variable X significantly affects survival-chances of the organisms. 	<ul style="list-style-type: none"> d. Gene X can be typed from highly degraded and processed tissue-samples. e. Environmental variable X has been recorded previously and is available. f. For conservation reasons, we are interested in adaptive change and ecological speciation; gene X is under heavy selection pressure.
D₃ : ‘What method should we use to analyse the samples?’	<ul style="list-style-type: none"> a. Guide tree X has better evidential support than guide tree Y. b. There is gene-flow between groups A and B; Method X can incorporate gene-flow, while method Y cannot. c. Species delimitation method X is unreliable when data is strongly geographically structured. 	<ul style="list-style-type: none"> d. Model X assumes neutrality and is unlikely to track recent ecological speciation; given our interest in adaptive evolutionary change, this model should not be used here. e. Method X tends to split, and method Y tends to lump; given our interest in recently diverged lineages, method X is preferable. f. We are familiar with software X, and have the computational power to use software X with this data-set.
D₄ : What classification should we infer from the results of analysis?	<ul style="list-style-type: none"> a. MtDNA – nDNA divergence can be explained by ILS/introgression/sex-biased dispersal/sex-biased hybrid viability. b. Morphological stasis is common in related taxa, and can explain the lack of diagnosability. c. Group X is reciprocally monophyletic for three of the sampled genes, but not for the two other genes. 	<ul style="list-style-type: none"> d. Current conservation legislation and field atlases accept species status for group X. It would be practically convenient to maintain this status. e. Given our interest in adaptive change, proof of the latter should trump lack of divergence through drift and evidence of gene-flow. f. Given that species-status is required for special conservation status, weak signs of divergence should be sufficient to recognise group X as a species. g. To avoid the practical problems of inflation, we should higher the species-level threshold of genetic divergence.

Table 1. Epistemic and Non-Epistemic Reasons. This table distinguishes between hypothetical epistemic and non-epistemic reasons that influence decisions in the internal stages of a taxonomic research process.

For example, take D_{1a} and D_{1b} ; the research on *Discoglossus* suggests that there are many data-types that reliably pick out species in that genus (e.g. nDNA and mtDNA), and other data-types that are not reliable (e.g. DNA content). Hence, epistemic considerations limit the set of legitimate options but do not determine the choice completely. According to the Epistemic Priority View, taxonomists can then rely on value-judgments about the type of group we would like to delimit (D_{1c}), the availability of data (D_{1e}) and whether to choose a method that involves harming specimens (D_{1f}) to choose between the epistemically acceptable data-types. Similarly, the choice between tracking allopatric speciation through drift or sympatric ecological speciation can be determined by values (D_{2f}), as long as we have good reasons to believe that a particular set of environmental factors affects a group of organisms (D_{2c}) or that a particular gene is neutral or under selection (D_{2b}).

4. The Entanglement of Epistemic and Non-Epistemic Reasons

The previous section illustrated how the Epistemic Priority View proposes to regulate and evaluate the role of values in science. Very simply put, the strategy is to distinguish epistemic reasons from value-judgments and to prioritize the former over the latter. This section argues against this strategy by showing that epistemic reasons and non-epistemic reasons are so tightly entangled that it is impossible in practice to separate them. The reason for this is that epistemic reasons take value-laden outcomes of earlier decisions as their starting point. Thus, the epistemic parts of these reasons are entangled with value-judgments that influenced these earlier decisions. In this way, value-judgments influence the outcome of scientific decisions through the epistemic reasons. While this might not show that the distinction between epistemic and non-epistemic reasons is theoretically unsound, I believe it is enough to show that the Epistemic Priority View, which relies on scientists and philosophers being able to make this distinction, is not helpful in thinking about values in science.

To make these arguments, I once again rely on the examples from table 1. Subsection 4.1 shows that the epistemic reasons cited in column 2 are in various ways entangled with value-judgments from column 3. Subsection 4.2 argues that due to radical species pluralism, this

entanglement is likely to be pervasive and very difficult to undo in practice. It follows that any view that relies on distinguishing epistemic and non-epistemic reasons is not viable.³⁷

4.1. The non-epistemic in the epistemic

There are at least two ways in which epistemic reasons are entangled with value-judgments from decisions earlier in the internal stages.³⁸ The first is that value-judgments determine what evidence and epistemic reasons are relevant to the various decisions that scientists make throughout the research process. Take epistemic reason (D_{4c}), which holds that the taxon is reciprocally monophyletic for three of the sampled genes. D_{4c} is only relevant given the choice for DNA as data-type in D_1 , the choice to sample those particular genes in D_2 , and the choice to analyse the data with a particular method in D_3 . All these decisions (D_{1-3}) are the result of both value-judgments and epistemic reasons. This way, these value-influences constrain what epistemic questions are relevant, and hence in what direction the epistemic questions push the decision. This is important because it could be that this direction, which could have been different given different value-judgments in earlier decisions, clashes with the value-judgments relevant to that decision. According to the Epistemic Priority View, epistemic reasons should trump the non-epistemic reasons in such a conflict. However, given that earlier value-judgments co-determine the direction of the epistemic reason, portraying this decision as a conflict between purely epistemic and non-epistemic reasons oversimplifies the clash of reasons at hand.

The second way in which value-laden decisions in earlier stages influence the epistemic decisions in later stages concerns cases where value-judgments in earlier decisions directly determine the meaning of terms figuring in the epistemic reasons for later decisions. Consider D_{4a} , which sums up various reasons for discordance between mtDNA and nDNA evidence. García-París and Jokusch (1999) point out that there are likely to be several of these processes

³⁷ Various authors (Bueter, 2015; Elliott, 2009; Elliott & McKaughan, 2009) have made similar arguments on the basis of epistemological underdetermination. As this thesis focuses on ontological underdetermination, I will not elaborate on their arguments here.

³⁸ I only aim to show that epistemic reasons are closely entangled with value-judgments. This does not mean that non-epistemic reasons are free of epistemic standards. A quick look at the examples in table 1 shows that all non-epistemic reasons are likely a mix of both components. Matt Brown (2017) points out that this is important because it further enforces the claim that value-judgments are not merely subjective preferences and can have empirical content that bears on the decision at hand (see also Anderson, 2004). I will not focus on this here, as the argument only requires that epistemic reasons are entangled with value-judgments. If it turns out that value-judgments have epistemic content too, this further reinforces the conclusion that epistemic standards should not unconditionally trump value-judgments.

at work in the case of *jeanneae*. Thus, the explanation that is chosen depends on the way the data-types are operationalized (D_2) and even how the data is analysed (D_3). That is, the phrase ‘mtDNA-nDNA discordance’ in later stages of the research means different things depending on the particular evidence that supports it, and hence on the practical choices concerning the choice of data-type, operationalization and analysis earlier in the research. This severely influences the epistemic reasons that rely on these terms later in the research. Depending on the evidence that was sampled in the *jeanneae* case, mtDNA – nDNA discordance might point at sex-biased dispersal, lower viability of female hybrids, or incomplete lineage sorting (García-París & Jockusch, 1999). In the case of sex-biased dispersal and sex-biased hybrid viability, this would support a single-species hypothesis (as they assume gene-flow between the taxa), while in the case of incomplete lineage sorting it would support a two-species hypothesis. Other examples include the meaning of ‘gene-flow’ (D_{3b}), reciprocal monophyly (D_{4c}) and diagnosability (D_{4b}).

It seems then that the epistemic reasons listed in table 1 are entangled with the value-judgments that influence earlier decisions in the research process. This puts pressure on the distinction between epistemic and non-epistemic reasons: if the former are influenced by value-judgments, it is not clear what distinguishes them from the latter. This, in turn, puts pressure on the Epistemic Priority View, which proposes a strategy for regulating values in science that relies on this wedge. More precisely, the problem for the Epistemic Priority View is that prioritizing epistemic standards might inadvertently and unnoticeably lead us to arbitrarily prioritizing the values embedded in epistemic reasons over the values embedded in non-epistemic reasons.

It should be noted that the problem I am pointing to here is not the traditional problem of determining which normative statements are truth-attaining (i.e. epistemic standards) and which normative statements are not (i.e. value-judgments). As mentioned in chapter 3, I assume for the sake of argument that there is an unproblematic distinction between truth-attaining epistemic standards and value judgments. The point here is that even if we accept that distinction, it is not always clear which of these epistemic standards and value-judgments are at work when scientific decisions are made. This is because value-judgments from earlier decisions are embedded in epistemic reasons later in the research. Thus, for any decision in the internal stages of the research process, it might be that the clash between a value-

judgment and an epistemic standard conceals the clash between two value-judgments. This means that prioritizing the epistemic standard might inadvertently lead to arbitrarily prioritizing one set of values over another. In short, the problem is not that we cannot identify epistemic standards and value-judgments, but rather that we cannot identify which standards and value-judgments are at work in any decision.

4.2. Radical pluralism and ontological underdetermination

The previous subsection argued that the wedge between epistemic and non-epistemic reasons is not helpful because value-judgments and epistemic standards are strongly entangled in the internal stages of taxonomy. The proponent of the Epistemic Priority View might want to respond here that taxonomists could simply dissect the epistemic reasons in table 1, and distinguish between their epistemic and non-epistemic components. This would allow them to appropriately prioritize epistemic standards over value-judgments without arbitrarily prioritizing some values over others. Matt Brown (2017, p. 10) calls this position ‘epistemic reductionism’, because it relies on reducing complex reasons to their purely epistemic and non-epistemic components. Such epistemic reductionism would allow the Epistemic Priority View to regulate value-influences on scientific decisions despite the entanglement of the two components.

This subsection argues that while epistemic reductionism might be theoretically feasible, it is not practically useful as a basis for evaluating and regulating values in taxonomy. In order to make this argument, it is helpful to return to the discussion of species ontology in the first two chapters of this thesis. More precisely, it is worth considering how the difference between moderate and radical pluralism affects the usefulness of the Epistemic Priority View and its wedge between epistemic and non-epistemic reasons. I argue that while this wedge might be philosophically fruitful for domains that are monist or only moderately pluralist, it is not helpful for research, like taxonomy, that deals with a part of the world that is radically pluralist. Briefly put, the reason for this is that classification of pluralist domains requires making ontologically underdetermined, value-laden decisions. When these are low in number, as in the case of moderately pluralist domains, they are likely to be easy to keep track of. As the number of ontologically underdetermined decisions in various stages of research increases, this becomes increasingly difficult. Given that classification in radically pluralist domains requires very many ontologically underdetermined decisions, the entanglement of epistemic

and non-epistemic reasons is likely to be far too difficult to disentangle for each step in the research process.

First, imagine that the world was such that there are only three legitimate classifications of organisms into species, namely, phylogenetic species, interbreeding species, and ecological species (e.g. Ereshefsky, 2001, 2010a). Assuming that we want a classification that assigns each organism to a single species, accepting a group in our all-purpose classification means that two other equally real groups are disregarded. Thus, values would enter taxonomy through the choice between these three kinds of groups. It seems plausible that the Epistemic Priority View and its epistemic reductionism would provide a helpful framework for evaluating these value-influences. As delimiting interbreeding, ecological and phylogenetic species all require different methodological approaches, values would play a limited role in choosing the method of species delimitation. Values would also influence scientific decisions if there are multiple proposed classifications as a result of their picking out different types of groups. We might, for example, adopt an interbreeding group rather than a phylogenetic group because it will prove to be more useful for conservation purposes.

The point here is that for all these different roles of values, it seems possible to distinguish between the epistemic reasons weighing in on a decision (e.g. the hypothesis that this is an interbreeding species best accounts for the observation that these organisms produce viable offspring when mating with each other) and the usefulness of the group (e.g. we should recognize the interbreeding species because it is more useful for species conservation than the alternatives). Moreover, the epistemic reason does not seem to imply any value-judgments. For example, assuming that there is only one good classification of organisms into biological species, any good method that tracks reproductive isolation will lead to the same classification. Hence, value-judgments due to ontological underdetermination cannot be hiding in the way we test for reproductive compatibility.³⁹

The situation becomes considerably more complex if, instead of moderate pluralism, we assume radical pluralism concerning species classifications. Radical pluralism is the claim that there are innumerable equally real groups on the species level. On this view, species

³⁹ It should be noted that value-judgments could still enter in other ways; for example, the choice between these tests might be epistemologically underdetermined because it is very difficult to test for reproductive compatibility. However, I focus here solely on the implications of the AOU.

classification is ontologically underdetermined in a much stronger sense than on the moderately pluralist view: not only are there multiple different species concepts, but these species concepts themselves vastly underdetermine the different ways in which they can be specified. As a result, there are innumerable classifications of organisms into groups of the kinds that taxonomists recognise as species. This has important implications for the usefulness of the wedge-strategy.

Consider again the use of the wedge-strategy for mapping the role that values play in taxonomic research. In the case of moderate species pluralism, this role was limited to the choice between three different types of groups, and directly related methodological choices. On the radically pluralist view, the choice of concept is only one in a long range of ontologically underdetermined decisions that influence the outcome. The four decisions in table 1 provide a non-exhaustive selection from this range of choices. Other examples include decisions about sampling locations and density, the choice between discovery-based and hypothesis based approaches, and the choice of the starting hypothesis. In addition, it is worth pointing out that the four coarse-grained decisions in table 1 often contain further sub-decisions, such as the choice of assumptions and parameter values in a particular model of evolution, the substrate to use for typing DNA, or which outgroups to use for analysis. Again, these fine-grained decisions are often ontologically underdetermined and require value-judgments to bridge the gap.

The point here is that these value-judgments co-determine the outcome of every fine-grained decision. As the outcomes of these decisions are used in later stages of the research, these value-judgments become entangled with the epistemic reasons relevant to these later decisions. Because there are so many ontologically underdetermined decisions, the complex of value-judgments and epistemic standards becomes increasingly complicated as it relies on outcomes from previous research, and as the research continues. In order to decide what role value-judgments should play in a particular decision, we cannot simply investigate (as was the case for moderate pluralism) whether their only role was to choose between three types of legitimate classification. Rather, we would have to disentangle the epistemic and non-epistemic components of the long range of ontologically underdetermined scientific decisions leading up to the current decision. Only then would it be possible to properly prioritize the epistemic over the non-epistemic.

Proponents of the Epistemic Priority View might once again point out that these arguments only show that epistemic reductionism is very difficult to apply, and not that it is theoretically untenable. However, assuming that the primary desideratum for a view on values in science is to guide scientists and philosophers in evaluating and regulating values in scientific practice, this objection has very little bite. Even if such epistemic reductionism were in principle possible, it is hard to see how it could serve to guide the role of values in any particular example of taxonomic research. Given that the taxonomic research process is riddled with ontologically underdetermined decisions influencing other decisions further down the line, such a task seems even too complex to undertake for a single species hypothesis, let alone the millions of species already recognized and yet to be discovered. As Matt Brown (2017, p. 10) writes, ‘epistemic reductionism is not workable in practice, and it provides no normative guidance to scientists, nor is it useful to philosophers of science and others trying to understand or evaluate episodes of scientific *practice* in detail’. I conclude then that the second wedge assumed by the Epistemic Priority View is not tenable due to practical difficulties following from ontological underdetermination.

Let us take stock of the arguments in this chapter so far. I have argued that the Value-Free Ideal and the Epistemic Priority View rely on two wedges: the distinction between the internal and external stages, and the distinction between epistemic and non-epistemic reasons. I have argued that the AOU raises problems for both assumptions. First, values that enter the external stages continue to influence the internal stages; and second, distinguishing epistemic and non-epistemic reasons in the internal stages is difficult as the two are tightly entangled. It follows that any good view on values on science should not banish values from the internal stages of science (failure of the first wedge), and should not rely on prioritizing epistemic standards over value-judgments (failure of the second wedge). Given that the Aims view rejects these two wedges and the Value-Free Ideal and Epistemic Priority View depend on them, it follows that the AOU supports the former and discredits the latter two.

5. Objection: A weakened Epistemic Priority View

The arguments in the previous sections have shown that unconditionally prioritizing epistemic standards over value-judgments is not a viable strategy because the two factors are entangled in many complex ways. Daniel Steel, a major proponent of the Epistemic Priority View, concedes these points, but argues that they do not invalidate his views. He points out that the

Epistemic Priority View does not require that evidence is generated independently from value-judgments in a 'pristinely value-free environment' (Steel, 2017, p. 18). In line with the arguments in this chapter, he claims that values determine what counts as data in the first place, and how observations are converted into data. This suggests that the evaluation of hypotheses inevitably requires relying on the value-judgments implied by the epistemic standards and data. According to Steel (2017, p. 19), the entanglement of epistemic standards and values is compatible with the claim that 'some basic epistemic criteria of adequate science are non-negotiable'. These minimal, non-negotiable criteria, he claims, set constraints on what scientists are permitted to do without irredeemably corrupting their research.

Emphasizing their basic and minimal character, Steel chooses to call this view an 'Epistemic *Constraint View*' rather than an 'Epistemic Priority View'. On this constraint view, the epistemic does not unconditionally trump value-judgments. Rather, they mark minimal epistemic standards that any acceptable research should meet. If values somehow override these standards, Steel holds, corruption of science is inevitable and values have played an illegitimate role. The idea of this view is that even if values and standards are inevitably entangled, overriding these absolute minimal criteria can be avoided by ensuring that they are heeded in every decision throughout the research process.

This constraint view is considerably weaker than the Epistemic Priority View as presented in Douglas' (2009, 2013, 2016) work and Steel's earlier work (Steel, 2010; Steel & Whyte, 2012). While the Epistemic Priority View holds that values can only determine those decisions that are not determined by epistemic standards, the weakened version holds that these standards merely set minimal constraints within which values can act. The difference is significant, as it leaves open the possibility that scientists can opt for an epistemically weaker option as long as it falls within the constraints of the minimal criteria. While the constraint view still requires some prioritizing of the epistemic over the non-epistemic, this priority is not absolute or unconditional. Epistemic standards are only prioritized in the sense that whenever values entirely override these standards, they are considered to play an illegitimate role; but within the constraints of these standards, the weakened Epistemic Priority View allows that value-judgments sometimes trump epistemic standards. Steel (2017) calls this the 'qualified' priority of epistemic standards, and points out that such a non-absolutist sense of prioritization is by far the most common one:

It is not unusual to speak of one's priorities, but this would be absurd if the only priority were an absolute one—since in that case there could only be a single priority and not several priorities. Non-absolutist priorities, I claim, are the ordinary kind. For example, I prioritize my children's well-being, but this does not mean I will stop at nothing to advance it. (2017, p. 20)

Steel claims that proponents of the Epistemic Priority View can accept the strong entanglement of values and still hold on to a weakened version of epistemic priority. One might question whether this response successfully fends off the arguments from the previous sections. Even a weakened priority might still rely on the ability to distinguish these epistemic constraints from value-judgments in earlier stages, and it is not clear whether their minimal character is sufficient to assuage this worry. However, I will not follow this line of objection, and assume for the sake of argument that this entanglement does not pose problems for Steel's position. Instead, I want to argue that this weakened Epistemic Priority View merely restates the importance of the world dependency principle. As this principle is accepted by proponents of the Aims View, this effectively turns the Epistemic Priority View into a variant of the Aims View.

In order to show that the weakened Epistemic Priority View converges with the world-dependency criterion, it is helpful to look at Steel's examples of qualified priority. Steel (2017, p. 19) writes that 'scientists are not permitted to fudge their data or rig their experiments in the service of a greater good'. While these are the only examples that he gives in his response to the entanglement-objection and his discussion of qualified epistemic priority, other examples are scattered throughout the same and other papers. These examples mainly rely on the severity principle, which holds that a good scientific test should put the tested hypothesis at risk of being disproven (Steel & Whyte, 2012). In addition to this he refers to the manipulation and hiding of data such as in the Vioxx case (Steel, 2017, p. 4), the testing of drugs against other drugs that are known to not work very well, and the usage of inappropriately low or high dosages when comparing drugs (Steel & Whyte, 2012, p. 169). Douglas (2009, pp. 87, 100) refers to similar examples, and likewise stresses the importance of the severity principle.

It is clear that these examples concern cases where values play a clearly illegitimate role. It is also clear that by adopting the weakened Epistemic Priority View and its basic epistemic

standards we could avoid such cases. However, these examples fail to support a substantial sense of priority of the epistemic. This is because all these examples concern cases that breach the world-dependency criterion. That is, in all these cases, the evaluation of hypotheses was not dependent on what the world is like. Instead, scientists used value-judgments in the role that is usually played by data, and accepted or rejected hypotheses merely on the basis of the outcome they desired. The point here is that the world-dependency criterion is accepted by proponents of the Aims View and Epistemic Priority View alike. To argue against the profound role for values defended by the Aims View, proponents of the Epistemic Priority View have to show that values should be restricted beyond the restrictions already set in place by the world-dependency criterion. If they fail to do this, as I argue they do, their position is compatible with the Aims View. Thus, by weakening the Epistemic Priority View to a qualified version of epistemic priority that only prohibits wishful thinking and dogmatism, the view is rendered widely accepted but trivial. Brown (2017, p. 15) hints at a similar conclusion in his arguments against the priority of the epistemic:

Note that I have not argued that scientists may accept and advocate claims on the basis of values alone, without consideration for evidence or empirical inquiry. Nor have I argued that epistemic considerations are unimportant, or less important than values considerations. It is no doubt that epistemic constraints are a centrally important part of scientific inquiry, and if this is all that EPT [i.e. Epistemic Priority View] required, I doubt anyone could deny it.

Another way of stating this is that the difference between the Aims View and the Epistemic Priority View does not lie in the acceptance of the world-dependency principle. Rather, the Aims View argues that there is a role for values in science somewhere between the tie-breaker role and the world-dependency principle, while the Epistemic Priority View holds that values should not play any role beyond that of tie-breaker. Showing that the world-dependency principle constitutes a non-negotiable principle in science is not relevant to this debate.

I conclude then that while some sense of epistemic priority might be feasible even if we assume the entanglement of values and epistemic standards, this weakened version of the Epistemic Priority View reduces to the widely accepted world-dependency criterion. This means that even if the view expresses a legitimate concern regarding the role of values in science, it is trivial in this debate and fails to provide the substantially different view on values

in science that its proponents set out to provide. More specifically, it fails to show that values in the internal stages should be restricted in a more substantial sense than proponents of the Aims View already accept.

6. A Positive Argument from Ontological Underdetermination

This chapter has argued that distinguishing between epistemic considerations and value-judgments to prioritize the former over the latter is not a helpful strategy for thinking about values in science. More precisely, it has argued that the Value-Free Ideal fails because the distinction between values in the internal and external stages is untenable, and that the Epistemic Priority View fails because epistemic and non-epistemic reasons are tightly entangled. Given that these views directly oppose the Aims View, it can be argued that an argument in favour of the Aims View results. That is, if value-influences should not be restricted to the external stages (*contra* the Value-Free Ideal), it follows that they should influence the internal stages (as the Aims View claims). Similarly, if separating epistemic and non-epistemic reasons is unfeasible (*contra* the Epistemic Priority View), it follows that values and epistemic standards should be considered on equal footing in a joint necessity view (as the Aims View claims).

To close the chapter, I want to briefly suggest that the considerations in this and the previous chapter also suggest a positive argument for the Aims View. To see this, it will be useful to return to the AOU. This argument shows that in any stage of science, the aim of truth underdetermines scientific decisions. Thus, it implies that scientists do more than merely producing true representations of the world. Very simply put, they also determine which of very many fine-grained causal relations are relevant to the particular problem they are trying to solve, and which of these can be glossed over in the classifications they come up with. It follows from this that values should guide methodological decisions, and the selection and interpretation of evidence. For example, when taxonomists set out to delimit a particular species, they do not only aim to provide a true representation of particular causal processes and a clustering of properties. They also attempt to pick out those processes that are relevant to the particular aims they have and the particular problems and questions they set out from. The point of the AOU is that neither of these two aspects (or aims) of science is more or less important. Species classification is as much a matter of selecting some causal relations and

glossing over others, as it is a matter of representing such relations. While it may be true that it is characteristic of science to solve problems by producing true representations of the world, it is also undeniable that such representations will be of little help in answering questions or solving problems without the appropriate selecting and glossing over. In addition to this, this chapter has shown that these two aspects are tightly entangled with each other. Just like decisions that are determined by epistemic considerations, ontologically underdetermined decisions occur throughout the research process and co-determine the outcomes of research. Assuming that both aspects are crucial and part of all scientific decisions, it is easy to see where the problem with prioritizing the epistemic lies. A successful solution to a scientific problem requires that both aspects are addressed to the degree required by that particular problem. As both aspects are important and necessary, there is no reason why either of these should be prioritized consistently and unconditionally. Moreover, as both aspects are entangled and are thus part of all scientific decisions, there is no reason why either of these should be considered more important in particular parts of the research process.

When an outcome is epistemically respectable but does not meet the relevant set of value-judgments, scientists may choose to reconsider their methodological choices, background assumptions, and the evidence they rely on. Similarly, when an outcome is in line with the preferred set of values but not epistemically respectable, scientists may choose to reconsider various steps in the research process or to reconsider the particular value-judgments they started from. Which of these considerations takes precedence in any particular decision depends on the particular details of the research project, and cannot be captured in a simple rule like the one proposed by proponents of the Epistemic Priority View. While such a straightforward rule would be easier and more appealing, it unfortunately cannot account for the complex entanglement of values (non-epistemic aims) and epistemic considerations (epistemic aims) in science.

This joint-necessity perspective helps in clarifying why a profound role for values in science does not necessarily lead to the corruption of science. As discussed in chapter 3, wishful thinking occurs whenever value-judgments are taken as reasons to believe claims about what the world is like. Proponents of the Value-Free Ideal and the Epistemic Priority View then point out that values logically cannot play such a role, and that they should, therefore, be banned from particular parts of science in order to avoid that they do. The AOU suggests that the

wedge-strategy for regulating values in science risks throwing out the baby with the bathwater. If, as the AOU suggests, the internal stages of taxonomy are as much concerned with generating true theories as with selecting the relevant part of this truth, it follows that there is another role that values can play in these internal stages. The use of the word 'relevant' already indicates that this role concerns normative considerations such as 'any acceptable species classification should be useful for conservation purposes'. The point is that the content of at least some taxonomic claims is not purely descriptive, but also normative. It follows that normative considerations such as value-judgments can legitimately weigh in on scientific decisions. Playing this role, values can shape scientific outcomes in a profound way and trump epistemic reasons without leading to wishful thinking (see Intemann, 2005 for a similar argument).

Chapter 5: Classificatory Norms and Normative Underdetermination

The two previous chapters made a case for the profound value-ladenness of taxonomy on the basis of radical pluralism. I argued that value-judgments and epistemic standards should be on equal footing when scientific decisions are being made. I called this the ‘Aims View’, because it implies that values should not in the first place be evaluated on the basis of the role they play, but rather on the basis of their agreement with the epistemic and non-epistemic aims of science. To clarify and strengthen this view further, this chapter compares it to an influential and popular view in debates about natural kinds that connects the legitimacy of scientific classifications to both the world and the purposes of that classification (e.g. Boyd, 2000; Ereshefsky, 2001; Slater, 2017). Because this view, which I will call the ‘classificatory norms view’ (CNV henceforth), connects the legitimacy of scientific classification to the aims of science, one might think that it is in line with the Aims View. However, the evaluation in this chapter will reveal that the Aims View and the CNV are significantly different. I argue that while the CNV may identify an important and legitimate way in which values should influence taxonomy, it does not go far enough. More precisely, I show that the role for values captured by the CNV is relatively superficial, and I argue that values should play the more profound role implied by the Aims View.

To make this argument, the first section presents the CNV and clarifies how the goals of taxonomy are claimed to influence classification. The second section compares the CNV to the Aims View, and distinguishes between two interpretations of the CNV: a weak interpretation compatible with the Aims View, and a strong interpretation that mitigates the role of values in taxonomy. The third section argues that the strong interpretation relies on the assumption that taxonomy is driven by a small number of broad, generally accepted goals. The fourth section argues that this assumption is not tenable. More precisely, it is argued that taxonomists are guided by many diverse aims which in turn are operationalized in many different ways. Section five, finally, argues that it follows from this that the strong interpretation of the CNV fails to mitigate the role of values in taxonomy. This suggests that

values play a profound role, in line with the Aims View. The chapter ends with a brief summary in section five.

1. Classificatory Norms and Radical Pluralism

In the two first chapters of this thesis, I argued that radical pluralism obtains if we assume the principles of species individuation used by the three major approaches to species classification. That is, I argued that if species classification is only determined by genealogical relations, cohesion-promoting causal factors, or similarity relations, then there are innumerable legitimate ways of classifying organisms into species. In order to present the classificatory norms view, it is helpful to return to this radical pluralism, and, more particularly, to an important argument against it.

Various authors argue that, even if the ontological argument for pluralism holds, such radical pluralism about the legitimacy of classification is not desirable (Ereshefsky, 2001; Ghiselin, 1987; Hull, 1997; Magnus, 2012). One important reason for this is that radical pluralism is strongly at odds with common practice in biology, where monism or moderate pluralism are the rule. If radical pluralism about the legitimacy of species classification really were to hold, it would follow that there are no good scientific reasons for recognizing only one or few classification(s). However, taxonomists regularly do recognize just one or a few classification(s). Moreover, they typically offer reasons for their choices. These practices and reasons are not comprehensible on the radically pluralist view. This suggests that radical pluralism fails to capture the considerations that scientists rely on to determine the legitimacy of classifications. A view on the legitimacy of classification that is so far removed from scientific practice is, as Hacking (2007, p. 229) writes, at risk of being merely ‘scholastic’ talk, or part of an ‘inbred set of degenerating problems that have increasingly little to do with the issues that arise in a larger context’.

In short, these authors reject radical pluralism about the legitimacy of classifications because it does not seem to have the resources to make sense of the practice of scientific classification. This means that they have to assume that there are other factors that, in addition to genealogical relations and cohesion-promoting mechanisms, determine the legitimacy of species classification. Such factors could explain the fact that scientists typically, and presumably for good reasons, only adopt one or a few distinct classifications. And indeed, this

is precisely what many authors have claimed in debates about species, and more generally, natural kinds. Kitcher (1984b, p. 309), for example, writes that the legitimacy of species classification is determined by the world and by what is biologically interesting. Similarly, Dupré (1993, p. 51) argues that the legitimacy of classifications depends on them being useful for ‘some significant purpose’, and Boyd (1999, p. 148) connects the legitimacy of natural kinds to the ‘inductive and explanatory aims’ of a discipline. Ereshefsky (2001, chapter 5) and Slater (2017) develop this position more extensively, and argue that the gap between radical and moderate pluralism about classifications is bridged by generally accepted norms of classification.⁴⁰ Their claim is that these norms, together with the world, limit the set of possible classifications and provide a criterion of legitimacy.

Of course, these conventions cannot just be any set of preferences. This would bring us back to the radical form of pluralism in which there is no way to arbitrate between multiple classifications beyond the constraints of the world. Instead, Ereshefsky and Slater single out a particular set of conventions that, together with the world, should determine the legitimacy of classification. More precisely, they claim that the generally accepted goals of the relevant scientific discipline determine the legitimacy of classification. Slater (2017, p. 9) writes that

[t]he moderate pluralist contends not merely that different classification systems are possible – this is obviously true – but that among those different possibilities, a number of them are legitimate by the lights of shared higher-level aims. Classificatory choice operates within a limited space of legitimate possibilities. The question of moderation or radicalness of classificatory pluralism for a given domain turns on the size of this space – the degree to which the world and our norms constrain our classificatory activities – at any level we care to countenance.

In other words, Slater accepts that there may be innumerable ontologically valid classifications of any part of the world, but argues that only a few of these are suitable given the purposes we have for constructing these classifications in the first place. The claim then is that radical species pluralism turns into moderate pluralism when we take into account the requirements of the goals of taxonomy.

⁴⁰ Many others, such as Chakravartty (2011), Magnus (2012) and Khalidi (2013), make similar remarks. I focus on Slater, Ereshefsky and Boyd as they have developed this position in most detail.

Boyd's HPC view assumes a similar position. He argues for a 'bicameral' view on natural kinds, on which the legitimacy of kinds depends as much on their accommodation to the inductive and explanatory aims of a discipline as on their accommodation to the causal structure of the world. He acknowledges that radical pluralism poses a threat if we think of natural kinds as solely determined by the world. He (Boyd, 2010, p. 222) writes that

if we think of natural kinds as things somehow independent of linguistic and methodological practices, then there are lots of natural kinds out there, and it is difficult to see how the causal conception of reference fixing could explain how a natural kind term could ever have a unique referent.

However, this problem disappears, according to Boyd, once the role of our conceptual practices and their inductive and explanatory aims is taken into account. According to Boyd, natural kinds are only natural relative to a particular disciplinary matrix. Thus, natural kinds are not just any categories that pick out chunks of the causal structure of the world, but rather those categories that pick out the causal structure that is relevant to meeting the inductive and explanatory demands of a particular discipline. Indeed, Boyd regularly points out that it makes little sense to talk about natural kinds solely in terms of the causal structure of the world, independently from the inductive and explanatory practices that are being accommodated to this causal structure.

Given that Slater, Ereshefsky and Boyd all refer to the goals of scientific disciplines to rein in radical pluralism, it is worth briefly considering how disciplines are individuated according to these authors. Ereshefsky does not address this question explicitly, but refers to biological taxonomy as the relevant discipline and to the goals of taxonomy as restricting the legitimacy of species classifications. Similarly, Slater does not explicitly address what counts as a discipline, but gives (evolutionary) biology and astronomy as examples.

Boyd (1999, p. 148) too refers to traditional scientific disciplines such as evolutionary biology, psychology and fundamental physics as scientific disciplines, but explicitly points out that disciplinary matrices need not correspond to 'academic or practical disciplines otherwise understood'. Instead, a disciplinary matrix is any 'family of inductive and explanatory aims and practices, together with the conceptual resources and vocabulary within which they are implemented' (Boyd, 2000, p. 57). These matrices are similar to HPC kinds in that they are partly determined by the world, and partly by our scientific practices. They are determined by

us in that they are dependent on the explanatory and inductive aims of research.⁴¹ There is also a sense in which disciplinary matrices are determined by the world. Boyd argues that it is characteristic for disciplinary matrices that the satisfaction of any particular accommodation demand usually contributes to the satisfaction of many other demands. The reason for this is that there are relations between the various causal mechanisms to which classifications in a single discipline are accommodated. Thus, the causal structure of the world co-determines the individuation of disciplinary matrices.⁴²

This account of disciplinary matrices is compatible both with interpreting them as relatively broad, conventional scientific disciplines (e.g. evolutionary biology), and with interpreting them as narrowly characterized research projects (e.g. understanding the transmission of malaria). However, if disciplinary matrices were to be understood as narrow research projects, it is unclear how referring to these matrices could help in avoiding radical pluralism. If there are innumerable disciplinary matrices within biology, then there may also be a legitimate species classification for each of these matrices. Boyd's (1999) occasional use of 'project' to refer to disciplinary matrices may suggest such a pluralist interpretation. However, his examples of evolutionary and functional biology, psychology and fundamental physics as disciplinary matrices, as well as his remark that these matrices are co-determined by very many interrelated causal mechanisms suggest that this is not Boyd's intention. Most importantly, given that Boyd explicitly refers to evolutionary biology and functional biology as the relevant disciplinary matrices for species classification, we can rely on this broad understanding of disciplinary matrices for the investigation here.

In short, then, Boyd, Ereshefsky and Slater argue that the legitimacy of species classification is determined both by the world and by the aims of taxonomy. This way, radical ontological pluralism is reduced to a more moderate form of pluralism about the legitimacy of classification. Note that these authors do not argue for species monism. Boyd (1999, p. 160), for example, refers to Kitcher's moderate species pluralism and writes that 'different accommodation demands within a disciplinary matrix, will – given the complexity of the

⁴¹ In addition, they are also determined by conventional choices between various 'equipollent' classifications that meet the accommodation demands of the discipline equally well and do not reflect differences in the causal structure of the world (Boyd, 1999, p. 176). I come back to these value-influences below.

⁴² P.D. Magnus adopts a similar view, but, unlike Boyd, claims that the individuation of disciplinary matrices is fully determined by what the world is like.

biological world – require reference to different, and cross classifying, kinds in order to achieve accommodation’. Similarly, Slater (2017, p. 9) writes that ‘a realistic and interesting moderate pluralism’ is likely to follow from the constraining role of classificatory norms. Ereshefsky, finally, argues that the overriding aims of taxonomy, in combination with the world, justify three distinct legitimate classifications. As noted in the introduction, I will refer to the view that constrains radical pluralism by relying on classificatory norms as the ‘classificatory norms view’ (CNV).

Before I turn to comparing the CNV to the Aims View defended over the previous chapters, it is worth briefly considering how precisely classificatory norms play their constraining role. Given the relatively broad conception of disciplines, these goals are likely to be quite general and vague. Indeed, Ereshefsky argues that the overriding goal of taxonomy is to construct an inferentially strong classification. Slater refers to simplicity, stability and inferential strength as examples of general goals of a discipline. Such general goals are too vague to guide classificatory choices in a direct and precise way. For this reason, Slater and Ereshefsky point out, these goals have to be operationalised through low-level norms that spell out the best way to attain the general goals.⁴³ For example, Slater argues that the norm ‘favour intrinsic over extrinsic properties as a basis for classification’ may be seen as a lower-level operationalization of a more general norm like ‘favour stable classifications’, which in turn may be an operationalization of a more general goal like Ereshefsky’s ‘support inferences’. It is best then to think of classificatory norms in terms of a hierarchical set of nested norms, where the higher-level norms justify the use of lower-level norms, and the overall goals ultimately justify all other norms. In that sense, the overriding goals of a discipline are considered intrinsically valuable within that discipline, while the lower-level norms are only instrumentally valuable as a means of fulfilling the overriding goals of the discipline. Together with the world, these overriding goals and lower-level norms then guide classificatory decisions in science.

2. Classificatory Norms and the Aims View

The CNV is in several respects similar to the Aims View defended over the previous two chapters. Like the Aims View, the CNV holds that value-judgments should shape the outcome

⁴³ Ereshefsky calls these low-level norms ‘methodological rules’; Slater calls them ‘ground-level norms’.

of species classification because these are not fully determined by what the world is like. Moreover, like the Aims View, the CNV also holds that this underdetermination problem occurs because different classifications reflect different aspects of the causal structure of the world. Finally, like the Aims View, the CNV holds that the legitimacy of these value-judgments is connected to the aims of taxonomy. That is, the space of legitimate classifications is constrained in different ways depending on the particular aims that are generally accepted. Whether a classification is acceptable, then, depends both on the particular set of values and the causal structure of the world.

I argue that despite these significant similarities, the two views are importantly different. This becomes apparent if we look in more detail at the role that values play according to proponents of the CNV. As shown in the passage quoted above (p. 129), Slater clarifies this role using the image of a space of ontologically legitimate classifications from which classificatory norms select the useful ones. Similarly, Ereshefsky first argues for ontological species pluralism, and then points out that radical pluralism (which he calls ‘anything goes pluralism’) can only be avoided by relying on values to determine which of many ontologically good classifications we should recognise. Thus, the CNV suggests a picture in which there are many ontologically good classifications ready to be found in the world, and the role of values lies in choosing one of these classifications. This picture is perhaps most clearly reflected in Boyd’s notion of accommodation, which refers to the tuning of classificatory practices to those causal structures that are relevant to the particular phenomena that we want to explain or make generalizations about.

This role for value-judgments in classification is aptly expressed by Reydon’s notion of the ‘zooming-in model’ of natural kinds. The CNV suggests that while values may determine what causal structures we ‘zoom in’ on, they play no further role in determining the content of these kinds. For example, values may decide whether we zoom in on interbreeding relations (Biological Species Concept), genealogical relations (Phylogenetic Species Concept) or cohesion-promoting causal processes (Evolutionary Species Concept), but they play no role in deciding which groups of organisms count as biological, phylogenetic or evolutionary species. In their discussion of the individuation of evolutionary groups, Barker and Velasco (2013, p. 972) make a similar remark about Ereshefsky’s and Kitcher’s views. They write that ‘we are asked to relinquish objectivism about the *application* of evolutionary groups concepts [...]

[but] whether a collection of entities satisfies a precisely understood definition of the morphological species concept is not something our conventions influence.'

The point here is that this 'zooming in' role for values is less profound than the role for values expressed by the Aims View. Indeed, the role for values highlighted by the CNV seems perfectly compatible with the Epistemic Priority View. Slater, Boyd and Ereshefsky hold that on the one hand there are mind-independent classifications in which values play no role, and on the other hand there are value-judgments that select which of these classifications we need. This means that values only come in after epistemic considerations have played their role. These values can then act as a tie-breaker to choose between many good ways of carving up the world into natural kinds. The Aims View, on the other hand, is stronger. It claims that values are involved in all parts of the taxonomic research process, and not merely in choosing between various natural classifications. That is, on top of allowing for a 'zooming in' role for values, proponents of the Aims View claim that values are involved in the process of determining what counts as a natural kind in the first place.

It should be noted that even though the CNV points to a more superficial role for values than the Aims View, these two views need not be incompatible. Interpreted in a weak sense, the CNV only implies that classificatory norms constitute one of the ways in which values influence the internal stages of taxonomy. This weak interpretation is fully compatible with the Aims View, which holds that value-judgments profoundly affect the internal stages of taxonomy and should guide all scientific choices together with epistemic standards. Interpreted in a stronger sense, on the other hand, proponents of the CNV can be taken to claim that classificatory norms constitute the *only* way in which value-judgments affect ontologically underdetermined decisions in the internal stages of taxonomy. This strong reading is incompatible with the Aims view, which holds that values play a profound role beyond the mere zooming in on the world.

Most authors who acknowledge the role of classificatory norms do not distinguish between these weak and strong interpretations of the CNV. Dupré, Slater and Ereshefsky only argue that classificatory norms play an important role in fending off the threat of radical pluralism, but do not explicitly claim that classificatory norms exhaust the role of values in the internal stages. Others, like Kitcher (2001, p. 48), adopt the stronger position, and explicitly claim that

values have no further role to play after they have determined what kind of group to recognise.

Boyd, too, seems to adopt this strong interpretation in his discussion of conventional elements in classification.⁴⁴ He argues that there are two unproblematic ways in which conventions may shape natural classifications. First, value-judgments play a crucial role in the shape of explanatory or inductive aims that determine which parts of the causal structure of the world to 'zoom in' on. This is the role also highlighted by other proponents of the CNV. Second, Boyd (1999, p. 176) holds that value-judgments also play a role in choosing between various 'equifertile' conceptual schemes, i.e. schemes that meet the accommodation demands of a discipline equally well. For example, the choice between various currencies to express the differences between the national debts of multiple countries unproblematically relies on conventions, as all currencies are equifertile here.

Boyd denies that values should play any role that extends beyond scientific aims and choices between equifertile classifications. Most importantly, he denies that there should be any conventional elements that play a role in 'the establishment of causal powers and relations' (Boyd, 1989, p. 22). He writes that 'what is arbitrary or conventional [...] does not reflect a distinctive feature of causal structure' (Boyd, 1989, p. 24). Thus, values should not be involved in determining the causal profile of the kinds we recognise. For example, while they can determine the decision to study phylogenetic groups and even how many Linnaean ranks to recognise, they cannot decide what groups are exclusive or monophyletic, and what kind of causal relations and properties these groups have. This, according to Boyd, should only be a matter of the causal structures of the world.

Let us take stock of the chapter so far. The first section introduced the view that values in taxonomy take the shape of classificatory norms. The second section then pointed out that while this position is somewhat similar to the Aims View, it is only compatible with that view provided we adopt a weak interpretation of the CNV. This weak interpretation holds that one way in which values enter taxonomy is in the shape of generally accepted classificatory norms.

⁴⁴ The term 'conventions' is, of course, not entirely synonymous with 'values'. However, as I define values in opposition to epistemic standards, rules of logical inference, and data, conventions arguably fall within this category. Moreover, Boyd seems to use the term 'conventions' in a rather broad way, since he clarifies the role of conventions in classification by opposing it to the role played by the causal structure of the world (Boyd, 1991, p. 140).

On a strong interpretation, which is not compatible with the Aims View, this position holds that values should *only* enter taxonomy through these overriding goals. The remainder of this chapter will argue against the strong interpretation. Briefly summarized the argument has the following structure. First (section 3), I clarify the conditions under which classificatory norms can be assumed to reduce radical ontological pluralism to a moderate pluralism about the legitimacy of classifications. Then (section 4) I show that these conditions do not apply in taxonomic practice; this means that even if we assume that there are generally accepted classificatory norms at work in taxonomy, they do not fend off radical pluralism. This in turn means that additional value-judgments are still needed for species classification. Section 5 then argues that these values play a profound role. More precisely, I argue that in addition to ‘zooming in’ on kinds, values are also involved in determining what counts as a kind in the first place. This role for values goes beyond that supported by the CNV, and is in line with the Aims View.

Even if these claims are compatible with Slater’s and Ereshefsky’s (assuming they adopt the weak reading), I believe it is important to highlight the role of these values in order to avoid mitigating the role of values in taxonomy. Moreover, these claims go against Boyd’s and Kitcher’s strong interpretation of the CNV, which holds that conventional elements should be restricted to classificatory norms.⁴⁵

3. Assumptions of the Classificatory Norms View

To make the argument against the strong interpretation, it will be helpful to return to the way in which classificatory norms are employed to fend off the threat of radical pluralism. While the weak interpretation merely holds that classificatory norms constrain the legitimacy of classification, the strong interpretation holds that they do this in such a way that moderate pluralism or monism results, and additional value-judgments are no longer needed for ontologically underdetermined decisions. This means that the strong interpretation must

⁴⁵ It is important to note that even though I argue against Boyd’s accommodation view here, my views are in many respects similar to his. Most importantly, by rejecting the accommodation view I do not mean to reject that the world plays a crucial role in determining the legitimacy of kinds. Thus, while I argue that values play a deeper role than what Boyd allows for, this should not be taken to imply that values can play any role whatsoever, or that, depending on our values, literally any group of organisms may qualify as a species. The difference between Boyd’s view and my view is thus one in degree. Despite being merely a difference in degree, I hold this difference is important because it has significant implications for the way we should regulate and evaluate values in classification.

assume that classificatory norms reduce radical ontological pluralism to a moderate pluralism concerning legitimate classifications. If this is not the case, and there are still innumerable potentially legitimate classifications after classificatory norms have been taken into account, then the strong interpretation fails. In such cases, classification is both ontologically and normatively underdetermined, and additional value-judgments are required to determine the legitimacy of species classifications.

In other words, then, the strong interpretation is only tenable if the world and classificatory norms together constrain the space of legitimate classifications to a moderate few. As discussed extensively in chapter 2, there are innumerable potential classifications that fall within the constraints set by the world. This means that the strong interpretation requires an account of the goals of taxonomy that shows how they reduce this space of ontologically legitimate classifications to a few that are legitimate given a particular set of goals. While it is beyond the scope of this paper to provide and evaluate such an account, we can derive at least two properties these goals must have in order to play the constraining role required by the strong interpretation.

First, these goals must be limited in number and generally accepted within the relevant discipline. This is because different classifications are legitimate depending on the goals that are adopted. If different scientists adopt different goals, then the number of legitimate classifications is likely to increase. A (radical) pluralism of goals thus results in a (radical) pluralism of legitimate classifications. This requirement fits well with Ereshefsky's claim that the goal of constructing inferentially strong classifications is the only overriding goal of taxonomy, and is universally accepted within the discipline. It is worth noting that while Ereshefsky's goal is a cognitive goal, this is not required for the CNV to work. In line with the arguments from the previous two chapters, the non-epistemic aims of science can play this constraining role too.

Second, like the general goals themselves, the ways in which these goals are operationalized must be generally accepted and limited in number. This is because different operationalizations lead to different classifications. If there are innumerable ways to operationalize the overriding goals of taxonomy, these goals cannot serve to constrain radical pluralism. For example, even if all taxonomists accept that the single overriding goal of their discipline is to construct inferentially strong classifications, radical pluralism still threatens if

there are many different sets of rules that lead to distinct classifications that are inferentially strong in different ways.

One may think that it is trivially true that there are innumerable possible goals that taxonomists could adopt. Without a further meta-norm to arbitrate between these norms, this leads us back to anything goes pluralism. And it is not clear where such a meta-norm could be found, as there are again innumerable possible meta-norms to choose from, and so on *ad infinitum*. However, the obvious fact that there are many *possible* norms does not show that the *actual* goals of science do not constrain the legitimacy of classifications (Slater, 2017, p. 9). Thus, the question is whether *in fact* all taxonomists share a few general (cognitive) goals that guide research and limit the number of legitimate classifications. I turn to that question in the next section.

4. Normative Underdetermination

Ereshefsky (2001, chapter 5) investigates important texts by leading taxonomists, and argues that taxonomy is guided by the overriding cognitive goal of supporting inferences, and a number of low-level norms derived from that goal. In this section, I argue, *contra* Ereshefsky, that a further look at the discipline of taxonomy suggests that it has many intrinsically valued goals, and even more low-level norms through which these are operationalised.⁴⁶

Let us start by considering the overriding goals of taxonomy. I have already pointed out in section 1 that it is important to the CNV that disciplines are individuated relatively broadly, because there will be as many legitimate species classifications as there are disciplines using the species concept. One reason then to think that the goals of taxonomy fail to rein in radical pluralism, is that taxonomy is defined in various ways. Taxonomy is sometimes defined as the scientific discipline involved with assigning names to groups of organisms, identifying groups of organisms and ordering these groups in a system of classification, or with discovering, identifying and naming species and reconstructing their history (Ereshefsky, 2007; Wheeler, 2008). If we assume that different definitions of taxonomy come with different sets of aims

⁴⁶ Ereshefsky (2001) and Ereshefsky and Reydon (2014) also offer an account of the norms that *should* guide (species) classification. These accounts single out a set of epistemic standards such as internal coherence, empirical testability and progressiveness as important classificatory norms. As my interest here lies in the aims and norms that actually guide species classification, I will not go into these normative accounts in detail.

for the discipline (e.g. reconstructing phylogenetic history vs. constructing a clear system of naming), pluralism might enter taxonomy through the different ways in which it is defined.

It is also worth pointing out that taxonomy is strongly entangled with other scientific disciplines. Gotelli (2004) uses his research on North American ants to illustrate the impact of taxonomy on ecological research. He emphasizes the importance of usable taxonomic keys, current nomenclature not hindered by synonymy, species occurrence records, and phylogenies. Similarly, Isaac et al. (2004) and Halme (2015) argue that a reliable and stable taxonomy is crucial for ecology. Others (Braby & Williams, 2016; Khuroo, Dar, Khan, & Malik, 2007; Mace, 2004) mainly emphasise the importance of taxonomy for conservation biology, which requires a reliable and as-complete-as-possible inventory of life on earth.

The close connection between taxonomy and other disciplines of evolutionary biology is further illustrated by the fact that a large share of recent taxonomic work is published as part of papers exploring hypotheses from these disciplines. Indeed, this is increasingly seen as a necessary aspect of taxonomic research. Halme (2015, p. 1834) writes that

[b]uilding one's resume' strategically is becoming more and more of a standard among academics and publishing solely descriptive taxonomy has become a difficult pathway to scientific positions. Many skillful taxonomists already work in close collaboration with systematists and evolutionary biologists or they are working on evolutionary hypotheses themselves, which allows publishing in higher-impact journals and improving citation rates.

The point here is that the close relation between taxonomy and other disciplines suggests that the goals of taxonomy are not entirely distinct from the goals of these closely related disciplines. Often taxonomic research is part of broader research projects, or taxonomists are motivated by the needs of other disciplines. This adds the goals of these other disciplines and research projects to the goals of taxonomy. If it turns out that, in line with my claims about taxonomy, these other disciplines have a wide range of cognitive and non-epistemic aims, this connection between taxonomy and the other disciplines suggests further proliferation of the aims of taxonomy.

However, let us set aside these worries for now and assume, for the sake of argument, that taxonomy is a well-delineated and clearly distinct discipline. I argue that even then it is hard

to maintain that taxonomists are guided by a few, generally accepted goals. Even if we assume, like Ereshefsky, that the goals of this discipline are purely cognitive, it is clear that they are broader than merely supporting inferences. MacLeod (2013; see also Love, 2009) argues that a narrow focus on the projectibility of natural kinds based on their causal basis has tended to obscure the various other epistemic roles that classification plays in science. One such role of species classification, which is also mentioned in several of the texts that Ereshefsky cites, is to store and retrieve information in an efficient way. Nelson and Platnick (1981, p. 9), for example, write that ‘classifications obviously perform an essential function in information storage and retrieval. They allow us to deal with tremendous amounts of data by subsuming a great deal of information into single words’. Love and MacLeod point to other epistemic goals of classification, such as facilitating modelling and experimentation, understanding, and explanation.

In addition to this, and in line with my arguments over the past two chapters, there is no reason to think that the goals of taxonomy are purely cognitive. Various authors have recently argued that non-epistemic goals are no less important than epistemic goals in guiding science (Elliott, 2013; Elliott & McKaughan, 2014; Potochnik, 2015). Species classification does not seem to be an exception here. This is clearly illustrated by taxonomists’ selection of research topics. If the main goal of taxonomy really were the maximization of inferential potential, then taxonomists should spread their research organisms evenly over the tree of life. However, it is clear that this is not the case in practice. While some clades, particularly within metazoans, are very commonly studied, other clades receive no or very little attention. By comparing the representation of species in the scientific literature with the representation of species on the internet, Wilson et al. (2007) identify economic incentives, biomedical incentives, practical limitations, and even charisma as the main factors that influence topic choice. These practical interests thus at least partially trump the goal of constructing a classification that is maximally inferentially strong.

In other words, the goals of taxonomy are not merely cognitive, but also pragmatic, social, environmental, etc. In one sense, this follows inevitably from the argument from ontological underdetermination (AOU). Given that there are many classifications that correctly represent the world, it should be no surprise that goals directly related to truth do not fully determine the choice between these classifications. There is also a sense in which this is a trivial point.

Taxonomists nearly always sacrifice epistemic goals for pragmatic goals when they accept species hypotheses despite having only collected a minute proportion of all the possible evidence that bears on the hypothesis. Collecting more data and conducting more research would likely lead to stronger classifications, but this is practically unfeasible given the limited resources of taxonomists and the large proportion of biodiversity yet to be discovered and described.

However, pragmatic concerns guide taxonomy not only in the form of practical constraints, but also more directly as overriding goals. Most obviously, taxonomists frequently emphasize enabling and facilitating biodiversity conservation as an important and practical goal of taxonomy. This goal is particularly pressing given the ongoing human-induced mass-extinction. Khuroo et al. (2007, p. 260) express this when they write that 'with growing demand for credible taxonomic information, taxonomists must facilitate, not roadblock, in stopping what may be the irreversible loss of biodiversity.' Frankham et al. (2012) put this into practice when they argue that a Biological Species Concept should be preferred over a diagnosability-based Phylogenetic Species Concept because this is important for the conservation of fragmented groups. Others (Mace, 2004) have pointed at the role of species classification in conservation legislation. As methodological debates about the future of taxonomy, almost without exception, point at the importance of the discipline for conservation, further examples are not hard to come by.

Biodiversity conservation is not the only practical goal that guides taxonomists. The recent Plenary meeting of the Linnean society of London titled 'Who Needs Taxonomists?' (see Linnean Society, 2014) provides a wide range of examples of taxonomic research directly aimed at practical applications. Tammy Horton discussed the importance of taxonomic research for the petroleum industry in deep water. Most countries have regulations that companies must meet before deep sea wells can be drilled. Meeting these regulations typically requires broad taxonomic knowledge of fauna and flora. Given the limited knowledge of the diversity of life in deep water, the petroleum industry typically needs taxonomists to do research in the particular areas they wish to explore. Geoff Boxshall discussed the importance of taxonomic work on caligid sealice, a parasite, for the salmon farming industry. Alistair Griffiths stressed the role of taxonomic research in managing plant tests, and informing policy and law about plant trade. Monique Simmonds, finally, discussed a wide range of cases in

which the ability to identify plants is crucial for various commercial activities, and for the regulation of these activities; examples include the trade in sandalwood essential oils, star anise, but also mining projects in Guinea and the impact of climate change on the distribution of coffee species.⁴⁷

These examples comprise just a small selection of the wide range of practical applications that directly motivate taxonomic research. The importance of these practical goals is further reflected in the guidelines of important funding sources for taxonomy such as the ‘Biotechnology and Biological Sciences Research Council’ and ‘Natural Environment Research Council’, which explicitly require taxonomic research to have direct practical applications, in particular food security, industrial biotechnology, health research, and more generally, wealth creation (see Linnaean Society, 2014). Taxonomists Kim and Byrne (2006, p. 799) express this source of goals for taxonomy powerfully when they write that ‘[t]axonomy should be reinvigorated and reinvented through collaborative, interdisciplinary research that brings taxonomic insights to bear on topics important to twenty-first century society (e.g., food security, invasive species, and ecosystem services).’ They argue that the mere description of biodiversity and the construction of an inferentially strong system form an overly narrow view on the goals of taxonomy. Instead, taxonomic research should accommodate ‘environmental and societal issues’ (p799).

While these various practical goals undoubtedly overlap with Ereshefsky’s aim of supporting new inferences, it is also clear that these aims sometimes clash and thus favour different classification. The clash between what one might call the ‘internal’ goals of taxonomy – constructing a system of classification that is maximally informative, compact, and allows for various inferences – and the external goals – being maximally catered to the needs of other biological disciplines – is vividly illustrated by recent debates about the current state and future of taxonomy. Some authors argue that because of the importance of an inventory of biodiversity for conservation biology, taxonomists should aim to discover and describe new species as fast as possible, even if this comes at the cost of a lower degree of evidential support and a shallower understanding of the evolutionary history, biogeography and current evolutionary characteristics of species (see Costello, Wilson, & Houlding, 2013; Joppa,

⁴⁷ Brief summaries of their points are provided in the report of the meeting (Linnaean Society, 2014), and the slides of the talks can be accessed on <https://www.linnean.org/meetings-and-events/events/who-needs-taxonomists>.

Roberts, & Pimm, 2011). This view, which Sluys (2013, p. 1096) calls the 'end user mode' of taxonomy, holds that providing usable classifications for other biological disciplines and practical applications is the main goal of taxonomy. Others argue that the main goal of taxonomy is to provide classifications that reflect a deep understanding of both the evolutionary history of taxa and the current evolutionary processes in these taxa. They argue that by reducing taxonomy to the 'superficial' practice of discovering and describing new species as fast as possible, we sacrifice its ability to provide a deeper understanding of the organic world (Wheeler, 2008, p. 3). Carvalho et al. (2014, p. 323) write that those who argue for such a superficial taxonomy

utterly fail to grasp the relevance of taxonomy in terms of its higher purpose for comparative biology. Mastering the minutiae of morphological and taxonomic description, elucidating historical relationships (phylogenetic, biogeographic), and producing predictive classifications are tasks that require time, specific training, and institutional commitment, but that ultimately will contribute much more to our understanding of the biota.

This trade-off between a taxonomy aimed at the end-users and a taxonomy aimed at the internal goals is nicely illustrated in the debates about DNA-barcoding (DeSalle, Egan, & Siddall, 2005; Hebert & Gregory, 2005; Tautz, Arctander, Minelli, Thomas, & Vogler, 2003; Waugh, 2007; Will, Mishler, & Wheeler, 2005). As a method of species discovery and delimitation, DNA barcoding is the method that delimits species by using a short and standardized gene region, most commonly the mitochondrial cytochrome c oxidase subunit I gene (abbreviated as COI or COX1). This method has been proposed as a fast and easy method of delimiting comparable species. Because it requires no further knowledge about the subject organisms, species delimitation using DNA barcoding can be done by non-specialists with very little training. In addition, barcoding is relatively cheap and fast. For these reasons, some authors have argued that it may offer a solution to the taxonomic impediment and provide conservation biologists with much needed inventories of biodiversity (e.g. Tautz et al., 2003). Waugh (2007, p. 195), for example, writes that even though DNA barcoding 'is not a thorough taxonomic description nor is it a tool for phylogenetic reconstruction', it 'may help speed up the work of taxonomists'. He (2007, p. 196) goes on to say that '[g]iven the enormity of the

task of identifying the world's biota and the many other potential tasks for which it might be employed, it would seem imprudent to ignore the promise of DNA barcoding'.

However, many taxonomists have reacted strongly against the use of this technique for species delimitation, pointing out that classifications entirely based on COI completely miss out on the complexity and biological context of groupings in the organic world. As Sluys (2013, p. 1097) writes, 'a DNA barcode merely points to a phenetic difference in a small stretch of DNA but does not provide more in-depth information on organismal divergence, either genetic (e.g. nuclear) or morphological'. Similarly, Wheeler (2014, p. 314) writes:

There are those who regard DNA-based phylogenetic analyses and DNA barcoding as the new incarnation of taxonomy or single path to its future [...], but I am not among them. Identifying species, conducting cladistic analyses, and retrieving information are all important parts of taxonomy, of course, but taxonomy is much more. Ultimately, taxonomy's exploration of homologous characters, species, and clades results in a detailed, comprehensive overview of the origins, diversification and 3.8 billion-yr-history of life summarized in a phylogenetic classification and associated Linnaean names. Understanding even a single taxonomic character, often dismissed as 'merely descriptive', involves a rich context of theory, background assumptions, and testable hypotheses.

The explicit debate about the relative importance of the internal and external aims of taxonomy clearly shows that taxonomists do not all agree about the overriding aims of their discipline.

Three important conclusions follow from the discussion in this section: first, taxonomy has many practical and cognitive goals; second, these goals sometimes clash, thus leading to different classifications; and third, there is no general agreement among taxonomists on which of these goals should trump other goals. These conclusions have direct implications on the ability of classificatory norms to keep radical pluralism at bay. These norms only limit the space of legitimate classifications to a moderate number if there are only a few, generally accepted goals. If there are many different goals, and if different taxonomists work with different goals in mind or balance them in different ways, radical pluralism re-enters through these multiple goals. It is worth noting here that this argument is crucially dependent on the arguments for radical species pluralism in chapter 2. If there were only a few ontologically

legitimate classifications, then not even an infinite number of distinct goals would lead to radical pluralism. Thus, it is only because there are innumerable different ontologically legitimate classifications that it is possible there are different optimal classifications for different goals.

There is at this point a potential objection one might raise against this conclusion and avert radical pluralism. This objection holds that the wide variety of goals can be rephrased in terms of a few very general goals, such as 'support inferences and explanations' and 'simplicity' or even 'usefulness'. This is consistent with the taxonomic literature, as taxonomists often explicitly state the goals of their discipline in such general terms (see Ereshefsky, 2001). However, regardless of whether this accurately describes the goals of taxonomy, this move fails to reduce pluralism. In order to see this, it is helpful to turn to the second tenet of classificatory norms, namely, the low-level norms that operationalise the high-level goals. If there are many different sets of low-level norms that are equally suitable to attain the high-level goals of a discipline, pluralism obtains even if there is general agreement about one or a few overriding goals. Phrasing the high-level goals of taxonomy in very general, vague terms makes it likely that this is the case. That is, goals like 'support inferences' or 'allow for practically useful classifications' can be interpreted and operationalised in innumerable ways. Depending on how we do this, different low-level norms and different classifications result, thus leading to radical pluralism.

This is confirmed by the enormous methodological diversity in the field of species delimitation (Camargo & Sites, 2013; Sites & Marshall, 2004). Even among taxonomists who subscribe to the same conception of species, and consequently also to similar goals of taxonomy, there is a dazzling variety of methods being used for species delimitation. Depending on which of these methods is adopted, different classifications result. Satler et al. (2013), for example, applied seven commonly used model-based methods of species delimitation to the same group of trapdoor spiders, which variously yielded between three and eighteen different species. The differences between these methods lie, among other things, in the choice between validation approaches and discovery approaches, in the kinds of evidence they use, and in the various assumptions of the model of evolution they adopt, e.g. concerning panmixia, the neutrality of the investigated genes, and gene-flow (Camargo & Sites, 2013; Leavitt et al., 2015; Sites & Marshall, 2004). Even if we assume that these methods are all

different operationalisations of the same high-level goals, it is undeniable that they lead to different results, and hence form a source of pluralism within taxonomy. The sheer number of different methods currently used by taxonomists suggests that this pluralism is of the radical rather than moderate kind.

I conclude then that even if we assume that taxonomists share at least some general goals, radical pluralism still holds because taxonomy also has multiple other practical and cognitive goals which are realised through multiple sets of low-level norms. Depending on which of these goals taxonomists adopt, how these goals are balanced, and how they are operationalized through low-level norms, different classifications result. In short, species classification is not only ontologically underdetermined, but also *normatively underdetermined*. Additional value-judgments are needed for selecting and balancing the various goals of a particular taxonomic project, and operationalizing these goals in a way that meets this balance and other normative constraints (e.g. do not harm the study organisms) in all steps of the research process.

5. Combining Normative and Ontological underdetermination

The previous section argued that even if classificatory norms play a substantial role in taxonomy, species classification is still underdetermined in an ontological and normative sense, and hence radical pluralism still applies. Taxonomic research projects have many varied aims that are operationalized in many ways. The multiplicity of aims, operationalizations, and various ways of balancing these goals implies that recognising species involves making value-judgments that go beyond the generally accepted aims of taxonomy. It follows from this that in addition to classificatory norms, further value-judgments are required for species classification.

An important implication of these arguments is that the CNV fails to reject radical pluralism in favour of a more moderate variant. Even if classificatory norms constrain classification, there will still be many more ontologically and normatively legitimate classifications than scientists actually recognise. Note that this does not mean that scientists may not have good scientific reasons when they opt for one particular classification rather than recognising very many. It only means that any good explanation of moderate pluralism (or monism) in scientific practice must refer to more than only the world and generally accepted classificatory norms. More

precisely, such explanations must be able to account for the influence of further value-judgments determining ontologically and normatively underdetermined decisions.⁴⁸

One might worry here that even if radical pluralism still holds, these arguments do not yet show that the CNV really fails. This is because the strong reading of the CNV does not merely hold that values only play a role in taxonomy in the shape of classificatory norms, but also that these values play a relatively superficial role. More precisely, the main reason why I objected to the CNV was that it sees values as merely ‘selecting’ classifications, rather than shaping them in a substantial sense. While the previous sections have shown that there should be more values in taxonomy than just generally accepted classificatory norms, it remains to be shown that these values should play a role that is more profound than just zooming in on the world.

However, I argue that support for a more profound role for these values follows directly from the arguments for normative and ontological underdetermination. More precisely, if we take into consideration the combination of these two forms of underdetermination, it is clear that the CNV’s picture of a set of generally accepted norms recognising one classification from a pool of ready-made, ontologically legitimate options is too simple. If taxonomy only had one or a few goals, then perhaps this picture could work. That is, if taxonomists constructed species classifications with one precisely defined goal in mind, then perhaps one best classification could be selected from innumerable ontologically good ones. However, instead of one set of generally accepted goals, taxonomy has many goals, which can be operationalized and balanced in many ways. In combination with the radical pluralism entailed by the ontological argument in chapter 2, it follows that classification can turn out in many different ways. Depending on which particular goals are adopted and how they are balanced and operationalized, the many ontologically underdetermined decisions scattered throughout the research process may be pushed in different directions. This, in turn, leads to the recognition of different groups of organisms as species.⁴⁹

⁴⁸ I come back to this point in the concluding remarks to this thesis.

⁴⁹ Note that these claims do not require that taxonomists make these decisions with explicit value-judgments in mind. It is likely that many of these decisions are incorporated in commonly used models, methods of analysis and data-collection practices. As Winsberg (2012, p. 130) puts it, value-judgments may be hiding in the ‘nooks and crannies’ of scientific methods. In addition, these value-judgments will often be determined by practical constraints, such as the cost of particular methods or familiarity with a method. All this is compatible with both the normative and descriptive claims of the AOU. The descriptive version holds that values inevitably play a role,

It follows that species are not simply found in nature by zooming in on the causal structure guided by the explanatory interests of taxonomy. Rather, species are shaped by the particular answers to ontologically underdetermined decisions. Each of these decisions is in turn determined by a set of value-judgments consisting of the multiple aims of taxonomy, practical constraints, and decisions about operationalizing and balancing these factors. Whatever species we end up with are as much the result of these value-judgments as they are the result of the causal structure of the world. Thus, instead of conceptualizing species in terms of a 'zooming-in' model of natural kinds, it is better to adopt what Reydon (2015a) calls a 'co-creation' model. Species are not merely found through zooming in on the appropriate causal structures, but through the co-creation of taxonomists and the world. For example, instead of assuming that biological species, phylogenetic species, or evolutionary species are simply found in the world, it is more fitting to say that they are constructed in the interaction between these value-judgments and the world.

be it through explicit decisions or by hiding in the nooks and crannies of our models. The normative version is stronger; it holds that given this role for values, scientists should try to consider the practical, moral and social consequences of their decisions and take these into account when making their decisions.

Concluding remarks

This thesis has argued that, due to radical pluralism and ontological underdetermination, value-judgments do and should play a profound role in constructing species classifications. This claim goes against what is arguably the dominant view about values in science, that is, the Epistemic Priority View; and goes beyond the most popular view on the role of values in classification, that is, the Classificatory Norms View.

In these concluding remarks, I want to briefly point to some implications of these claims. These are, at the same time, avenues of further research. Section 1 briefly discusses the epistemological questions concerning species delimitation raised in this thesis. Section 2 points out three important questions concerning values in science raised by this thesis. Section 3, finally, considers the implications of my arguments for debates about natural kinds.

1. The Epistemology of Species Delimitation

The role of values in species delimitation highlighted in this thesis raises questions about the objectivity of species classifications. Unlike classificatory norms, these value-judgments are not shared among taxonomists, and thus often lead different taxonomists to construct different classifications for the same organisms. Hence, these value-judgments seem to make species delimitation dependent on the particular point of view of the taxonomists. This potential threat to the objectivity of species classification in turn raises questions about the stability, repeatability and comparability of species classifications and the reliability of the methods of species delimitation. If the outcome of species delimitation is dependent on particular value-judgments of the taxonomists, different taxonomists investigating the same organisms may come to different results (not repeatable and stable), and different species are expected to be widely different kinds of groups (not comparable). Moreover, if different methods lead to different outcomes, it is not clear why any of these methods should be considered reliable.

These worries are particularly interesting for two reasons. First, many taxonomists claim that the development of new methods and reliance on new kinds of evidence over the last decade has strongly reduced problems of subjectivity, reliability, incomparability and stability in

species classification. The argument for the value-ladenness of taxonomy on the basis of ontological (and normative) underdetermination suggests that it is not clear how these new methods and types of evidence would decrease these problems. Indeed, given that these innumerable new methods all use different kinds of evidence and adopt different models of evolution, statistical assumptions and background theory, one could argue that the proliferation of methods has only increased the value-ladenness due to ontological and normative underdetermination. The question then is how these new methods and new types of evidence have increased the objectivity of species delimitation, and, more generally, why current methods of species delimitation should be considered reliable. While mostly unexplored in the philosophy of biology, these epistemological questions are in line with an increasing theoretical interest in the taxonomic literature in methods of species delimitation (Sites and Marshall 2004; Camargo and Sites 2013; Leavitt et al. 2015). Any good answer to these questions must be able to account for the influence of value-judgments as discussed in this thesis.

The second reason why these epistemological worries are particularly pressing is that they raise doubts about common explanations for the epistemic success of species classification. Species classifications and species numbers are used by various branches of biology and even outside science with considerable success. This success is typically explained by referring to the relation between species classifications and the features of the world they represent. If the arguments in this thesis hold, species classifications reflect value-judgments as well as features of the world. This suggests that we cannot straightforwardly refer to the relation between classifications and the world to explain their success. Instead, we have to explain how classification can be so successful despite the role of values, and we have to investigate whether and how values contribute to this success.

These epistemological questions suggest a fruitful new way in which philosophical investigations about species may progress.⁵⁰ Importantly, these questions can and must be tackled even if metaphysical discussions about species remain without solution. The immense variety of ways in which the various approaches to species classification can be

⁵⁰ This comment takes its cue from David Hull (1997, 372–373), who wrote in relation to species that '[u]nfortunately, philosophers have not addressed this question [about operationalization] in any detail [...]. For any student of science who is getting a bit bored with the Covering Law Model of Scientific Explanation or incommensurability, the topic of operationalizing concepts is wide open'.

operationalized obtains regardless of whether one believes that there is a fact of the matter about species classification. Thus, value-judgments play a role in species classification independently from what view on species ontology might eventually prove to be true. This is important as a resolution for metaphysical debates about species classification does not seem to be in sight, despite decades of debate and a massive literature on the ‘species problem’.

2. Values in Science

The previous chapters have argued that values play a profound role in species classification, and that the role of these values should be evaluated on equal footing with that of epistemic standards. This means that, apart from blatant cases of wishful thinking, the regulation of values should not focus on *how* values influence science, but on *which* values are doing this. Very simply put, the Aims View defended in this thesis holds that the relevant question about values in science is whether they are the right values. This claim raises two important and difficult questions.

First, and most obviously, the question is which values are the right ones, and how we can determine this. There already is much debate on these questions, mostly focusing on democratic procedures of selecting the relevant values (M. B. Brown, 2009; Kitcher, 2001, 2011). Unfortunately, these questions have proven to be very difficult, and consensus about adequate answers does not seem near. Second, even if the question about selecting values is ignored, it is not obviously clear how precisely these values should be implemented in scientific practice. I have argued that values should stand on equal footing with epistemic considerations in all scientific decisions. While this may be theoretically feasible, it is hard to see how this could be done without incurring high costs in the efficiency, trust in science and feasibility of science. Again, there already is debate on the various ways in which value-judgments could be implemented in science, and no clear answer to this question has at this point emerged (Betz, 2013; Douglas, 2005; Elliott, 2011b; John, 2015; Wilholt, 2013).

As the Aims View can only be helpful for guiding values in science if there is a way of answering these two questions, one might think that that the failure to do this so far provides a strong argument against the Aims View. However, it should be noted that any view on values in science that allows values to affect the external stages, needs to answer these questions. This means that proponents of the Value-Free Ideal and of the Epistemic Priority View too must

solve these problems before they can regulate values in science. One could argue that because the Aims View allows values free rein in a far larger part of science than the other two views, these problems are more damning for the former than for the latter. However, this objection has little bite. The problem with these questions does not in the first place lie in their application to the internal stages, but in the general principles required for answering them. Indeed, it seems likely that any good way of answering these questions for the external stages will contribute greatly to answering the same questions for the internal stages too.

A final question about values in science raised by this thesis concerns the nature and particularities of classification as opposed to other scientific activities when it comes to evaluating and regulating values. The lion's share of the philosophical debate about values in science has focused on questions about theory choice in the biomedical sciences. While an impressive range of case studies involving theory choice have been discussed in much detail by philosophers of science and ethicists, very little attention has been paid to the role of value-judgments in scientific classification.⁵¹ This thesis has argued that value-judgments may play a particularly prominent role in species classification through ontological underdetermination. However, this is just one way in which values may play a role, namely ontological underdetermination, based on the investigation of one particular case, namely species classification. Thus, this thesis constitutes just one small step towards a better understanding of values in scientific classification.

3. Values and Natural Kinds

Finally, it is interesting to consider briefly the implications of the arguments in this thesis for debates about natural kinds. As discussed in chapters 1, 4 and 5, the entanglement of value-judgments and natural kinds constitutes an important part of the epistemic views on natural kinds. The core idea of these views is that good scientific classification consists in attuning our classification to those parts of the world that are relevant to explaining the phenomena of interest. This fits well with what Reydon calls a 'zooming-in model' of natural kinds, which views scientific classification as zooming in on the joints of nature relevant to our interests. While such practices have an important human component – after all, our interests determine

⁵¹ There are, of course, notable exceptions, such as Dupré (2007a), Ludwig (2015), Slater (2017), and other proponents of the classificatory norms view.

which phenomena to study – natural kinds are also in an important sense discovered, since the causal structures they latch onto are mind-independent features of the world. Simply put, these views hold that science may pick up on many different kinds in the world depending on what our interests are, but the precise properties of whatever kinds we choose to recognise is entirely determined by the world.

The arguments in chapters 4 and 5 suggest that this ‘zooming in’ model is not appropriate for species classification. While classificatory norms may help us to zoom in on various parts of the organic world, it is unclear whether the kinds we are zooming in on are as mind-independent as the epistemic views on natural kinds assume. If the arguments in this thesis are correct, then there are innumerable ways of zooming in on any part of the causal structure of the organic world. How precisely this is done depends on the mind-independent world, but also on the epistemic and non-epistemic aims of taxonomy, practical and ethical constraints, and how all these considerations are balanced and operationalized. Thus, the resulting kind is as much the workmanship of these value-judgments as it is part of the mind-independent structure of the world. In the words of Reydon (2015a), then, a ‘co-creation’ model of species-kinds seems more appropriate than a ‘zooming-in’ model.

This raises the question to what extent this conclusion can be extended beyond the case of species to natural kinds in general. While answering this question is far beyond the scope of this thesis, it is worth pointing out that, at least at first sight, the arguments used here can be transferred to any domain that combines causal complexity (like the organic world) with a multiplicity of practical and theoretical goals (like taxonomy). This question is important because of the epistemological questions pointed at in section 1 of these concluding remarks. If natural kinds are value-laden in a profound sense, one might ask in what sense scientific classifications are objective, and how the value-ladenness of natural kinds is compatible with their epistemic usefulness. Any suitable answer to these questions must then be able to account for the role that value-judgments play.

Bibliography

- Alexandrova, A. (2016). Can the science of well-being be objective? *The British Journal for the Philosophy of Science*. <https://doi.org/10.1093/bjps/axw027>
- Anderson, E. (1995a). Feminist epistemology: An interpretation and a defense. *Hypatia*, 10(3), 50–84. <https://doi.org/10.1111/j.1527-2001.1995.tb00737.x>
- Anderson, E. (1995b). Knowledge, human interests, and objectivity in feminist epistemology. *Philosophical Topics*, 23(2), 27–58.
- Anderson, E. (2004). Uses of value judgments in science: A general argument, with lessons from a case study of feminist research on divorce. *Hypatia*, 19(1), 1–24.
- Andersson, J. O. (2005). Lateral gene transfer in eukaryotes. *Cellular and Molecular Life Sciences CMLS*, 62(11), 1182–1197. <https://doi.org/10.1007/s00018-005-4539-z>
- Andersson, J. O. (2009). Gene Transfer and Diversification of Microbial Eukaryotes. *Annual Review of Microbiology*, 63(1), 177–193. <https://doi.org/10.1146/annurev.micro.091208.073203>
- Attenborough, R. (2015). What are species and why does it matter? Anopheline taxonomy and the transmission of malaria. In A. Behie & O. Marc (Eds.), *Taxonomic Tapestries* (pp. 129–155). ANU Press.
- Bapteste, E., & Boucher, Y. (2009). Epistemological impacts of horizontal gene transfer on classification in microbiology. *Methods in Molecular Biology*, 532, 55–72. https://doi.org/10.1007/978-1-60327-853-9_4
- Bapteste, E., O'Malley, M. A., Beiko, R. G., Ereshefsky, M., Gogarten, J. P., Franklin-Hall, L., ... Martin, W. (2009). Prokaryotic evolution and the tree of life are two different things. *Biology Direct*, 4, 34. <https://doi.org/10.1186/1745-6150-4-34>
- Barbará, T., Martinelli, G., Fay, M. F., Mayo, S. J., & Lexer, C. (2007). Population differentiation and species cohesion in two closely related plants adapted to neotropical high-altitude “inselbergs”, *Alcantarea imperialis* and *Alcantarea geniculata* (Bromeliaceae). *Molecular Ecology*, 16(10), 1981–1992. <https://doi.org/10.1111/j.1365-294X.2007.03272.x>
- Barker, M. J. (2007). The empirical inadequacy of species cohesion by gene flow. *Philosophy of Science*, 74(5), 654–665. <https://doi.org/10.1086/525611>
- Barker, M. J., & Velasco, J. D. (2013). Deep conventionalism about evolutionary groups. *Philosophy of Science*, 80(5), 971–982. <https://doi.org/10.1086/673923>
- Barker, M. J., & Wilson, R. A. (2010). Cohesion, gene flow, and the nature of species. *The Journal of Philosophy*, 107(2), 61–79. <https://doi.org/10.5840/jphil201010721>
- Baum, D. A. (2009). Species as ranked taxa. *Systematic Biology*, 74–86. <https://doi.org/10.1093/sysbio/syp011>
- Baum, D. A., & Donoghue, M. J. (1995). Choosing among alternative “phylogenetic” species concepts. *Systematic Botany*, 20(4), 560. <https://doi.org/10.2307/2419810>
- Baum, D. A., & Shaw, K. (1995). Genealogical perspectives on the species problem. In P. C. Hoch & A. G. Stephenson (Eds.), *Experimental and Molecular Approaches to Plant Biosystematics*, (pp. 289–303). St. Louis, MO: Missouri Botanical Garden.
- Beatty, J. (1995). The evolutionary contingency thesis. In E. Sober (Ed.), *Conceptual Issues in Evolutionary Biology* (pp. 45–81). Cambridge, MA: MIT Press.

- Bernardo, J. (2011). critical appraisal of the meaning and diagnosability of cryptic evolutionary diversity, and its implications for conservation in the face of climate change. In T. R. Hodkinson, M. B. Jones, S. Waldren, & J. A. N. Parnell (Eds.), *Climate Change, Ecology and Systematics* (pp. 380–438). Cambridge University Press.
- Betz, G. (2013). In defence of the value free ideal. *European Journal for Philosophy of Science*, 3(2), 207–220. <https://doi.org/10.1007/s13194-012-0062-x>
- Biddle, J. (2007). Lessons from the Vioxx debacle: What the privatization of science can teach us about social epistemology. *Social Epistemology*, 21(1), 21–39. <https://doi.org/10.1080/02691720601125472>
- Biddle, J. (2013). State of the field: Transient underdetermination and values in science. *Studies in History and Philosophy of Science Part A*, 44(1), 124–133. <https://doi.org/10.1016/j.shpsa.2012.09.003>
- Boenigk, J., Ereshefsky, M., Hoef-Emden, K., Mallet, J., & Bass, D. (2012). Concepts in protistology: species definitions and boundaries. *European Journal of Protistology*, 48(2), 96–102. <https://doi.org/10.1016/j.ejop.2011.11.004>
- Boyd, R. (1989). What realism implies and what it does not. *Dialectica*, 43(1–2), 5–29. <https://doi.org/10.1111/j.1746-8361.1989.tb00928.x>
- Boyd, R. (1991). Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies*, 61(1–2), 127–148. <https://doi.org/10.1007/BF00385837>
- Boyd, R. (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 141–85). MIT Press.
- Boyd, R. (2000). Kinds as the “workmanship of men”: Realism, constructivism, and natural kinds. In J. Nida-Rümelin (Ed.), *Rationalität, Realismus, Revision: Vorträge des 3. Internationalen Kongresses der Gesellschaft für Analytische Philosophie*. Berlin: Walter de Gruyter.
- Boyd, R. (2010). Realism, natural kinds, and philosophical methods. In H. Beebe & N. Sabbarton-Leary (Eds.), *The Semantics and Metaphysics of Natural Kinds* (pp. 212–234). Routledge.
- Braby, M. F., & Williams, M. R. (2016). Biosystematics and conservation biology: critical scientific disciplines for the management of insect biological diversity. *Austral Entomology*, 55(1), 1–17. <https://doi.org/10.1111/aen.12158>
- Brigandt, I. (2003). Species pluralism does not imply species eliminativism. *Philosophy of Science*, 70(5), 1305–1316. <https://doi.org/10.1086/377409>
- Brigandt, I. (2009). Natural kinds in evolution and systematics: Metaphysical and epistemological considerations. *Acta Biotheoretica*, 57(1–2), 77–97. <https://doi.org/10.1007/s10441-008-9056-7>
- Brigandt, I. (2015). Social values influence the adequacy conditions of scientific theories: Beyond inductive risk. *Canadian Journal of Philosophy*, 45(3), 326–356. <https://doi.org/10.1080/00455091.2015.1079004>
- Brown, M. B. (2009). *Science in Democracy: Expertise, Institutions, and Representation*. MIT Press.
- Brown, M. J. (2012). The source and status of values for socially responsible science. *Philosophical Studies*, 163(1), 67–76. <https://doi.org/10.1007/s11098-012-0070-x>
- Brown, M. J. (2013). Values in science beyond underdetermination and inductive risk. *Philosophy of Science*, 80(5), 829–839. <https://doi.org/10.1086/673720>

- Brown, M. J. (2015). John Dewey's pragmatist alternative to the belief-acceptance dichotomy. *Studies in History and Philosophy of Science Part A*, 53, 62–70. <https://doi.org/10.1016/j.shpsa.2015.05.012>
- Brown, M. J. (2017). Values in science: Against epistemic priority. In D. Steel & K. C. Elliott (Eds.), *Current Controversies in Values and Science (Hardback) - Routledge* (pp. 64–78).
- Bueter, A. (2015). The irreducibility of value-freedom to theory assessment. *Studies in History and Philosophy of Science Part A*, 49, 18–26. <https://doi.org/10.1016/j.shpsa.2014.10.006>
- Busack, S. D. (1986). Biochemical and morphological differentiation in Spanish and Moroccan populations of *Discoglossus* and the description of a new species from southern Spain (Amphibia, Anura, Discoglossidae). *Annals of Carnegie Museum*, 55, 53–61.
- Camargo, A., & Sites, J. (2013). Species delimitation: A decade after the renaissance. In I. Pavlinov (Ed.), *The Species Problem - Ongoing Issues* (pp. 225–247). Rijeka: InTech.
- Capula, M., & Carti, M. (1993). Morphometric variation and divergence in the West Mediterranean *Discoglossus* (Amphibia: Discoglossidae). *Journal of Zoology*, 231(1), 141–156. <https://doi.org/10.1111/j.1469-7998.1993.tb05358.x>
- Carvalho, M. R. de, Ebach, M. C., Williams, D. M., Nihei, S. S., Trefaut Rodrigues, M., Grant, T., ... Nelson, G. (2014). Does counting species count as taxonomy? On misrepresenting systematics, yet again. *Cladistics*, 30(3), 322–329. <https://doi.org/10.1111/cla.12045>
- Carvalho, S. B., Brito, J. C., Pressey, R. L., Crespo, E., & Possingham, H. P. (2010). Simulating the effects of using different types of species distribution data in reserve selection. *Biological Conservation*, 143(2), 426–438. <https://doi.org/10.1016/j.biocon.2009.11.010>
- Chakravartty, A. (2011). Scientific realism and ontological relativity. *Monist*, 94(2), 157–180. <https://doi.org/10.5840/monist20119428>
- Costello, M. J., Wilson, S., & Houlding, B. (2013). More taxonomists describing significantly fewer species per unit effort may indicate that most species have been discovered. *Systematic Biology*, 62(4), 616–624. <https://doi.org/10.1093/sysbio/syt024>
- Cotterill, F. P. D., Taylor, P. J., Gippoliti, S., Bishop, J. M., & Groves, C. P. (2014). Why one century of phenetics is enough: response to “Are there really twice as many bovid species as we thought?” *Systematic Biology*, syu003. <https://doi.org/10.1093/sysbio/syu003>
- Coyne, J. A., & Orr, A. D. (2010). Speciation: A catalogue and critique of species concepts. In A. Rosenberg & R. Arp (Eds.), *Philosophy of Biology: An Anthology* (pp. 272–292). West Sussex: Wiley-Blackwell.
- Coyne, J. A., & Orr, H. A. (2004). *Speciation*. New York, NY: W.H. Freeman.
- Cracraft, J. (1983). Species concepts and speciation analysis. In R. F. Johnston (Ed.), *Current Ornithology*. Boston, MA: Springer US.
- Craver, C. F. (2009). Mechanisms and natural kinds. *Philosophical Psychology*, 22(5), 575–594. <https://doi.org/10.1080/09515080903238930>
- Crespi, B., & Nosil, P. (2013). Conflictual speciation: species formation via genomic conflict. *Trends in Ecology & Evolution*, 28(1), 48–57. <https://doi.org/10.1016/j.tree.2012.08.015>
- Crisp, A., Boschetti, C., Perry, M., Tunnacliffe, A., & Micklem, G. (2015). Expression of multiple horizontally acquired genes is a hallmark of both vertebrate and

- invertebrate genomes. *Genome Biology*, 16(1), 50. <https://doi.org/10.1186/s13059-015-0607-3>
- de Melo-Martín, I., & Intemann, K. (2012). Interpreting evidence. *Perspectives in Biology and Medicine*, 55(1), 59–70. <https://doi.org/10.1353/pbm.2012.0007>
- de Melo-Martín, I., & Intemann, K. (2016). The risk of using inductive risk to challenge the value-free ideal. *Philosophy of Science*, 83(4), 500–520. <https://doi.org/10.1086/687259>
- De Queiroz, K. (1999). The general lineage concept of species, species criteria, and the process of speciation: a conceptual unification and terminological recommendations. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 49–89). Cambridge, MA: MIT Press.
- De Queiroz, K. (2005a). A unified concept of species and its consequences for the future of taxonomy. *Proceedings of the California Academy of Sciences*, ((suppl. 1)), 196–215.
- De Queiroz, K. (2005b). Different species problems and their resolution. *BioEssays*, 27(12), 1263–1269. <https://doi.org/10.1002/bies.20325>
- De Queiroz, K. (2007). Species concepts and species delimitation. *Systematic Biology*, 56(6), 879–886. <https://doi.org/10.1080/10635150701701083>
- Degnan, J. H., & Rosenberg, N. A. (2006). Discordance of Species Trees with Their Most Likely Gene Trees. *PLoS Genetics*, 2(5), e68. <https://doi.org/10.1371/journal.pgen.0020068>
- Degnan, J. H., & Rosenberg, N. A. (2009). Gene tree discordance, phylogenetic inference and the multispecies coalescent. *Trends in Ecology & Evolution*, 24(6), 332–340. <https://doi.org/10.1016/j.tree.2009.01.009>
- DeSalle, R., Egan, M. G., & Siddall, M. (2005). The unholy trinity: Taxonomy, species delimitation and DNA barcoding. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1462), 1905–1916. <https://doi.org/10.1098/rstb.2005.1722>
- Devitt, M. (2008). Resurrecting biological essentialism. *Philosophy of Science*, 75(3), 344–382. <https://doi.org/10.1086/596056>
- Dewey, J. (1929). *The Quest for Certainty: A Study of the Relation of Knowledge and Action*. Minton, Balch.
- Doolittle, W. F. (2010). The attempt on the life of the Tree of Life: Science, philosophy and politics. *Biology & Philosophy*, 25(4), 455–473. <https://doi.org/10.1007/s10539-010-9210-x>
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67(4), 559–579.
- Douglas, H. (2005). Inserting the public into science. In S. Maasen & P. Weingart (Eds.), *Democratization of Expertise?* (pp. 153–169). Springer, Dordrecht. https://doi.org/10.1007/1-4020-3754-6_9
- Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press.
- Douglas, H. (2013). The value of cognitive values. *Philosophy of Science*, 80(5), 796–806. <https://doi.org/10.1086/673716>
- Douglas, H. (2016). Values in science. In P. Humphreys (Ed.), *The Oxford Handbook of Philosophy of Science*.
- Dowling, T. E., & Secor, C. L. (1997). The role of hybridization and introgression in the diversification of animals. *Annual Review of Ecology and Systematics*, 28, 593–619.

- Drezen, J.-M., Gauthier, J., Josse, T., Bézier, A., Herniou, E., & Huguet, E. (2017). Foreign DNA acquisition by invertebrate genomes. *Journal of Invertebrate Pathology*, 147(Supplement C), 157–168. <https://doi.org/10.1016/j.jip.2016.09.004>
- Dubois, A. (1998). Lists of European species of amphibians and reptiles: will we soon be reaching “stability”? *Amphibia-Reptilia*, 19(1), 1–28. <https://doi.org/10.1163/156853898X00304>
- Dupré, J. (1981). Natural Kinds and biological taxa. *The Philosophical Review*, 90(1), 66–90. <https://doi.org/10.2307/2184373>
- Dupré, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- Dupré, J. (1996). Promiscuous realism: Reply to Wilson. *The British Journal for the Philosophy of Science*, 47(3), 441–444. <https://doi.org/10.1093/bjps/47.3.441>
- Dupré, J. (1999). On the impossibility of a monistic account of species. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 3–22). Cambridge, MA: Bradford Books.
- Dupré, J. (2001). In Defence of classification. *Studies in History and Philosophy of Science Part C*, 32(2), 203–219. [https://doi.org/10.1016/S1369-8486\(01\)00003-6](https://doi.org/10.1016/S1369-8486(01)00003-6)
- Dupré, J. (2002). Is “natural kind” a natural kind term? *The Monist*, 85(1), 29–49.
- Dupré, J. (2007a). Fact and value. In H. Kincaid, J. Dupré, & A. Wylie (Eds.), *Value-Free Science: Ideals and Illusions?* (pp. 27–41). Oxford University Press.
- Dupré, J. (2007b). Fact and Value. Oxford University Press. Retrieved from <https://ore.exeter.ac.uk/repository/handle/10871/14370>
- Eldridge, N., & Cracraft, J. (1980). *Phylogenetic Patterns and the Evolutionary Process. Method and Theory in Comparative Biology*. New York, NY: Columbia University Press.
- Elliott, K. C. (2009). The ethical significance of language in the environmental sciences: Case studies from pollution research. *Ethics, Place & Environment*, 12(2), 157–173. <https://doi.org/10.1080/13668790902863382>
- Elliott, K. C. (2011a). Direct and indirect roles for values in science. *Philosophy of Science*, 78(2), 303–324. <https://doi.org/10.1086/659222>
- Elliott, K. C. (2011b). *Is a Little Pollution Good for You?: Incorporating Societal Values in Environmental Research*. Oxford: Oxford University Press.
- Elliott, K. C. (2012). Epistemic and methodological iteration in scientific research. *Studies in History and Philosophy of Science Part A*, 43(2), 376–382. <https://doi.org/10.1016/j.shpsa.2011.12.034>
- Elliott, K. C. (2013). Douglas on values: From indirect roles to multiple goals. *Studies in History and Philosophy of Science Part A*, 44(3), 375–383. <https://doi.org/10.1016/j.shpsa.2013.06.003>
- Elliott, K. C., & McKaughan, D. J. (2009). How values in scientific discovery and pursuit alter theory appraisal. *Philosophy of Science*, 76(5), 598–611. <https://doi.org/10.1086/605807>
- Elliott, K. C., & McKaughan, D. J. (2014). Nonepistemic values and the multiple goals of science. *Philosophy of Science*, 81(1), 1–21. <https://doi.org/10.1086/674345>
- Elliott, K. C., & Willmes, D. (2013). Cognitive attitudes and values in science. *Philosophy of Science*, 80(5), 807–817. <https://doi.org/10.1086/673719>
- Ereshefsky, M. (1991). Species, higher taxa, and the units of evolution. *Philosophy of Science*, 58(1), 84–101.
- Ereshefsky, M. (1992). Eliminative pluralism. *Philosophy of Science*, 59(4), 671–690.

- Ereshefsky, M. (1998). Species pluralism and anti-realism. *Philosophy of Science*, 65(1), 103–120.
- Ereshefsky, M. (2001). *The Poverty of the Linnaean Hierarchy: A Philosophical Study of Biological Taxonomy*. Cambridge, UK: Cambridge University Press.
- Ereshefsky, M. (2007). Systematics and taxonomy. In S. Sahotra & A. Plutynski (Eds.), *The Blackwell companion to the Philosophy of Biology* (pp. 99–118). Oxford: Blackwell. <https://doi.org/10.1002/9780470696590.ch6>
- Ereshefsky, M. (2010a). Darwin's solution to the species problem. *Synthese*, 175(3), 405–425. <https://doi.org/10.1007/s11229-009-9538-4>
- Ereshefsky, M. (2010b). Microbiology and the species problem. *Biology & Philosophy*, 25(4), 553–568. <https://doi.org/10.1007/s10539-010-9211-9>
- Ereshefsky, M. (2010c). What's wrong with the new biological essentialism. *Philosophy of Science*, 77(5), 674–685. <https://doi.org/10.1086/656545>
- Ereshefsky, M. (2011). Mystery of mysteries: Darwin and the species problem. *Cladistics*, 27(1), 67–79. <https://doi.org/10.1111/j.1096-0031.2010.00311.x>
- Ereshefsky, M. (2016). Species. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2016). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/sum2016/entries/species/>
- Ereshefsky, M., & Reydon, T. A. C. (2014). Scientific kinds. *Philosophical Studies*, 172(4), 969–986. <https://doi.org/10.1007/s11098-014-0301-4>
- Faurby, S., Eiserhardt, W. L., & Svenning, J.-C. (2016). Strong effects of variation in taxonomic opinion on diversification analyses. *Methods in Ecology and Evolution*, 7(1), 4–13. <https://doi.org/10.1111/2041-210X.12449>
- Fleishman, E., Noss, R. F., & Noon, B. R. (2006). Utility and limitations of species richness metrics for conservation planning. *Ecological Indicators*, 6(3), 543–553. <https://doi.org/10.1016/j.ecolind.2005.07.005>
- Frankham, R. (2015). Genetic rescue of small inbred populations: Meta-analysis reveals large and consistent benefits of gene flow. *Molecular Ecology*, 24(11), 2610–2618. <https://doi.org/10.1111/mec.13139>
- Frankham, R., Ballou, J. D., Dudash, M. R., Eldridge, M. D. B., Fenster, C. B., Lacy, R. C., ... Ryder, O. A. (2012). Implications of different species concepts for conserving biodiversity. *Biological Conservation*, 153, 25–31. <https://doi.org/10.1016/j.biocon.2012.04.034>
- Frankham, R., Lacy, R. C., Ballou, J. D., Dudash, M. R., Eldridge, M. D. B., Fenster, C. B., ... Ryder, O. A. (2014). Species concepts for conservation - Reply to Russello and Amato. *Biological Conservation*, 170, 334–335.
- Franklin-Hall, L. R. (2015). Natural kinds as categorical bottlenecks. *Philosophical Studies*, 172(4), 925–948. <https://doi.org/10.1007/s11098-014-0326-8>
- Fritz, B., Vences, M., & Glaw, F. (1994). Comparative DNA content in Discoglossus (Amphibia, Anura, Discoglossidae). *Zoologischer Anzeiger*, (233), 135–145.
- Fromhage, L., Vences, M., & Veith, M. (2004). Testing alternative vicariance scenarios in Western Mediterranean discoglossid frogs. *Molecular Phylogenetics and Evolution*, 31(1), 308–322. <https://doi.org/10.1016/j.ympev.2003.07.009>
- Fujita, M. K., Leaché, A. D., Burbrink, F. T., McGuire, J. A., & Moritz, C. (2012). Coalescent-based species delimitation in an integrative taxonomy. *Trends in Ecology & Evolution*, 27(9), 480–488. <https://doi.org/10.1016/j.tree.2012.04.012>

- García-París, M., & Jockusch, E. L. (1999). A mitochondrial DNA perspective on the evolution of Iberian *Discoglossus* (Amphibia: Anura). *Journal of Zoology*, 248(02), 209–218. <https://doi.org/null>
- García-París, M., & Jockusch, E. L. (1999). A mitochondrial DNA perspective on the evolution of Iberian *Discoglossus* (Amphibia: Anura). *Journal of Zoology*, 248(2), 209–218. <https://doi.org/10.1111/j.1469-7998.1999.tb01197.x>
- Ghiselin, M. T. (1974). A radical solution to the species problem. *Systematic Biology*, 23(4), 536–544. <https://doi.org/10.1093/sysbio/23.4.536>
- Ghiselin, M. T. (1987). Species concepts, individuality, and objectivity. *Biology and Philosophy*, 2(2), 127–143. <https://doi.org/10.1007/BF00057958>
- Ghiselin, M. T. (1999). Natural kinds and supraorganismal individuals. In D. L. Medin & S. Atran (Eds.), *Folkbiology* (pp. 447–460). Cambridge: MIT Press.
- Giangrande, A. (2003). Biodiversity, conservation, and the “Taxonomic impediment.” *Aquatic Conservation: Marine and Freshwater Ecosystems*, 13(5), 451–459. <https://doi.org/10.1002/aqc.584>
- Giere, R. N. (2003). A new program for philosophy of science? *Philosophy of Science*, 70(1), 15–21. <https://doi.org/10.1086/367865>
- Glaw, F., & Vences, M. (1991). Bioacoustic differentiation in Painted frogs (*Discoglossus*). *Amphibia-Reptilia*, 12(4), 385–394. <https://doi.org/10.1163/156853891X00031>
- Goodman, N. (1975). Words, works, worlds. *Erkenntnis*, 9(1), 57–73. <https://doi.org/10.1007/BF00223133>
- Gotelli, N. J. (2004). A taxonomic wish–list for community ecology. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 359(1444), 585–597. <https://doi.org/10.1098/rstb.2003.1443>
- Graham, L. R. (1987). *Science, Philosophy, and Human Behavior in the Soviet Union*. New York: Columbia University Press.
- Griffiths, P. (1999). Squaring the circle: Natural kinds with historical essences. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 209–228). Cambridge, MA: MIT Press.
- Gross, B. L., & Rieseberg, L. H. (2005). The ecological genetics of homoploid hybrid speciation. *Journal of Heredity*, 96(3), 241–252. <https://doi.org/10.1093/jhered/esi026>
- Groves, C. P. (2013). The nature of species: A rejoinder to Zachos et al. *Mammalian Biology - Zeitschrift Für Säugetierkunde*, 78(1), 7–9. <https://doi.org/10.1016/j.mambio.2012.09.009>
- Groves, C. P., & Grubb, P. (2011). *Ungulate Taxonomy*. Baltimore: JHU Press.
- Gutiérrez-Rodríguez, J., Salvi, D., Geffen, E., Gafny, S., & Martínez-Solano, Í. (2014). Isolation and characterisation of novel polymorphic microsatellite loci in Iberian painted frogs (*Discoglossus galganoi* and *D. jeanneae*), with data on cross-species amplification in *Discoglossus* and *Latonia* (Alytidae). *The Herpetological Journal*, 24(4), 261–265.
- Haber, M. H. (2012). Multilevel Lineages and Multidimensional Trees: The Levels of Lineage and Phylogeny Reconstruction. *Philosophy of Science*, 79(5), 609–623. <https://doi.org/10.1086/667849>
- Haber, M. H. (2016). The individuality thesis (3 ways). *Biology & Philosophy*, 31(6), 913–930. <https://doi.org/10.1007/s10539-016-9548-9>
- Hacking, I. (1991). A tradition of natural kinds. *Philosophical Studies*, 61(1–2), 109–126. <https://doi.org/10.1007/BF00385836>

- Hacking, I. (2007). Natural kinds: Rosy dawn, scholastic twilight. *Royal Institute of Philosophy Supplement*, 82(61), 203–239.
- Häggqvist, S. (2005). Kinds, projectibility and explanation. *Croatian Journal of Philosophy*, V(13), 71–87.
- Halme, P., Kuusela, S., & Juslén, A. (2015). Why taxonomists and ecologists are not, but should be, carpooling? *Biodiversity and Conservation*, 24(7), 1831–1836. <https://doi.org/10.1007/s10531-015-0899-3>
- Harris, J., & Froufe, E. (2005). Taxonomic inflation: species concept or historical geopolitical bias? *Trends in Ecology & Evolution*, 20(1), 6–7. <https://doi.org/10.1016/j.tree.2004.11.004>
- Hausdorf, B. (2011). Progress toward a general species concept. *Evolution*, 65(4), 923–931. <https://doi.org/10.1111/j.1558-5646.2011.01231.x>
- Hebert, P., & Gregory, T. R. (2005). The promise of DNA barcoding for taxonomy. *Systematic Biology*, 54(5), 852–859. <https://doi.org/10.1080/10635150500354886>
- Hedrick, P. W. (2013). Adaptive introgression in animals: Examples and comparison to new mutation and standing variation as sources of adaptive variation. *Molecular Ecology*, 22(18), 4606–4618. <https://doi.org/10.1111/mec.12415>
- Heller, R., Frandsen, P., Lorenzen, E. D., & Siegismund, H. R. (2013). Are there really twice as many bovid species as we thought? *Systematic Biology*, 62(3), 490–493. <https://doi.org/10.1093/sysbio/syt004>
- Hempel, C. G. (1965). Science and human values. In *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science* (pp. 81–96). New York: The Free Press.
- Hicks, D. J. (2014). A new direction for science and values. *Synthese*, 191(14), 3271–3295. <https://doi.org/10.1007/s11229-014-0447-9>
- Hoser, R. T. (1989). *Australian Reptiles and Frogs*. Sydney: Pierson & Company Pty Limited.
- Hoser, R. T. (1991). *Endangered Animals of Australia*. Sydney: Pierson & Company Pty Limited.
- Howard, D. (2006). Lost wanderers in the forest of knowledge: Some thoughts on the discovery-justification distinction. In J. Schickore & F. Steinle (Eds.), *Revisiting Discovery and Justification* (pp. 3–22). Berlin: Springer Science & Business Media. https://doi.org/10.1007/1-4020-4251-5_2
- Howard, D. (2009). Better red than dead—Putting an end to the social irrelevance of postwar philosophy of science. *Science & Education*, 18(2), 199–220. <https://doi.org/10.1007/s11191-007-9117-3>
- Hudson, R. R. (1992). Gene trees, species trees and the segregation of ancestral alleles. *Genetics*, 131(2), 509–513.
- Hull, D. L. (1965a). The effect of essentialism on taxonomy: Two thousand years of stasis (I). *The British Journal for the Philosophy of Science*, 15(60), 314–326.
- Hull, D. L. (1965b). The effect of essentialism on taxonomy: Two thousand years of stasis (II). *The British Journal for the Philosophy of Science*, 16(61), 1–18.
- Hull, D. L. (1976). Are species really individuals? *Systematic Biology*, 25(2), 174–191. <https://doi.org/10.2307/2412744>
- Hull, D. L. (1978). A matter of individuality. *Philosophy of Science*, 45(3), 335–360.
- Hull, D. L. (1981). Kitts and Kitts and Caplan on species. *Philosophy of Science*, 48(1), 141–152. <https://doi.org/10.1086/288985>
- Hull, D. L. (1987). Genealogical actors in ecological roles. *Biology and Philosophy*, 2(2), 168.

- Hull, D. L. (1997). The ideal species concept-and why we can't get it. In M. Claridge, H. Dawah, & M. R. Wilson (Eds.), *Species: the units of biodiversity*. (pp. 357–380). London: Chapman & Hall.
- Hull, D. L. (1999). On the plurality of species: Questioning the party line. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 23–48). Cambridge, MA: MIT Press.
- Intemann, K. (2001). Science and values: Are value judgments always irrelevant to the justification of scientific claims? *Philosophy of Science*, *68*(3), S506–S518.
- Intemann, K. (2005). Feminism, underdetermination, and values in science. *Philosophy of Science*, *72*(5), 1001–1012. <https://doi.org/10.1086/508956>
- Intemann, K. (2015). Distinguishing between legitimate and illegitimate values in climate modeling. *European Journal for Philosophy of Science*, *5*(2), 217–232. <https://doi.org/10.1007/s13194-014-0105-6>
- Intemann, K., & Melo-Martín, I. de. (2010). Social values and scientific evidence: the case of the HPV vaccines. *Biology & Philosophy*, *25*(2), 203–213. <https://doi.org/10.1007/s10539-009-9191-9>
- Isaac, N., Mallet, J., & Mace, G. M. (2004). Taxonomic inflation: Its influence on macroecology and conservation. *Trends in Ecology & Evolution*, *19*(9), 464–469. <https://doi.org/10.1016/j.tree.2004.06.004>
- Jardine, N., & Sibson, R. (1968). A model for taxonomy. *Mathematical Biosciences*, *2*(3), 465–482. [https://doi.org/10.1016/0025-5564\(68\)90030-8](https://doi.org/10.1016/0025-5564(68)90030-8)
- Jeffrey, R. C. (1956). Valuation and acceptance of scientific hypotheses. *Philosophy of Science*, *23*(3), 237–246. <https://doi.org/10.1086/287489>
- John, S. (2015). Inductive risk and the contexts of communication. *Synthese*, *192*(1), 79–96. <https://doi.org/10.1007/s11229-014-0554-7>
- Joppa, L. N., Roberts, D. L., & Pimm, S. L. (2011). The population ecology and social behaviour of taxonomists. *Trends in Ecology & Evolution*, *26*(11), 551–553. <https://doi.org/10.1016/j.tree.2011.07.010>
- Kaiser, H., Crother, B., Kelly, C., Luiselli, L., O'Shea, M., Ota, H., ... Wüster, W. (2013). Best practices: In the 21st century, taxonomic decisions in herpetology are acceptable only when supported by a body of evidence and published via peer-review. *Herpetological Review*, *44*(1), 8–23.
- Karl, S. A., & Bowen, B. W. (1999). Evolutionary significant units versus geopolitical taxonomy: Molecular systematics of an endangered sea turtle (genus *Chelonia*). *Conservation Biology*, *13*(5), 990–999. <https://doi.org/10.1046/j.1523-1739.1999.97352.x>
- Kendig, C. (2014). Towards a multidimensional metaconception of species. *Ratio*, *27*(2), 155–172. <https://doi.org/10.1111/rati.12037>
- Khalidi, M. A. (2013). *Natural Categories and Human Kinds: Classification in the Natural and Social Sciences*. Cambridge, UK: Cambridge University Press.
- Khalidi, M. A. (2015). Natural kinds as nodes in causal networks. *Synthese*, 1–18. <https://doi.org/10.1007/s11229-015-0841-y>
- Khuroo, A. A., Dar, G. H., Khan, Z. S., & Malik, A. H. (2007). Exploring an inherent interface between taxonomy and biodiversity: Current problems and future challenges. *Journal for Nature Conservation*, *15*(4), 256–261. <https://doi.org/10.1016/j.jnc.2007.07.003>
- Kim, K. C., & Byrne, L. B. (2006). Biodiversity loss and the taxonomic bottleneck: Emerging biodiversity science. *Ecological Research*, *21*(6), 794. <https://doi.org/10.1007/s11284-006-0035-7>

- Kitcher, P. (1984a). Against the monism of the moment: A reply to Elliott Sober. *Philosophy of Science*, 51(4), 616–630. <https://doi.org/10.1086/289208>
- Kitcher, P. (1984b). Species. *Philosophy of Science*, 51(2), 308–333.
- Kitcher, P. (1987). Ghostly whispers: Mayr, Ghiselin, and the “Philosophers” on the ontological status of species. *Biology and Philosophy*, 2(2), 184–192. <https://doi.org/10.1007/BF00057962>
- Kitcher, P. (2001). *Science, Truth, and Democracy*. Oxford University Press.
- Kitcher, P. (2011). *Science in a Democratic Society*. Prometheus Books.
- Kitts, D. B., & Kitts, D. J. (1979). Biological species as natural kinds. *Philosophy of Science*, 46(4), 613–622. <https://doi.org/10.1086/288898>
- Köhler, Jörn, Vieites, D. R., Bonett, R. M., García, F. H., Glaw, F., Steinke, D., & Vences, M. (2005). New amphibians and global conservation: A boost in species discoveries in a highly endangered vertebrate group. *BioScience*, 55(8), 693–696. [https://doi.org/10.1641/0006-3568\(2005\)055\[0693:NAAGCA\]2.0.CO;2](https://doi.org/10.1641/0006-3568(2005)055[0693:NAAGCA]2.0.CO;2)
- Köhler, Jörn, Vieites, D. R., Bonett, R. M., García, F. H., Glaw, F., Steinke, D., & Vences, M. (2005). New Amphibians and Global Conservation: A Boost in Species Discoveries in a Highly Endangered Vertebrate Group. *BioScience*, 55(8), 693. [https://doi.org/10.1641/0006-3568\(2005\)055\[0693:NAAGCA\]2.0.CO;2](https://doi.org/10.1641/0006-3568(2005)055[0693:NAAGCA]2.0.CO;2)
- Kornet, D. J. (1993). Permanent splits as speciation events: A formal reconstruction of the internodal species concept. *Journal of Theoretical Biology*, 164(4), 407–435. <https://doi.org/10.1006/jtbi.1993.1164>
- Kornet, D. J., & McAllister, J. W. (2005). The composite species concept: A rigorous basis for cladistic practice. In T. A. C. Reydon & L. Hemerik (Eds.), *Current Themes in Theoretical Biology* (pp. 95–127). Springer Netherlands. https://doi.org/10.1007/1-4020-2904-7_5
- Kourany, J. A. (2003). A philosophy of science for the twenty-first century. *Philosophy of Science*, 70(1), 1–14. <https://doi.org/10.1086/367864>
- Kourany, J. A. (2010). *Philosophy of science after feminism*. Oxford University Press, USA.
- Lacey, H. (2005). On the interplay of the cognitive and the social in scientific practices. *Philosophy of Science*, 72(5), 977–988. <https://doi.org/10.1086/508954>
- LaPorte, J. (2004). *Natural Kinds and Conceptual Change*. Cambridge, UK: Cambridge University Press.
- Laudan, L. (2004). The epistemic, the cognitive, and the social. In P. Machamer & G. Wolters (Eds.), *Science, Values, and Objectivity* (pp. 14–23). Pittsburgh, Pa: University of Pittsburgh Press.
- Lawrence, J. G., & Retchless, A. C. (2010). The myth of bacterial species and speciation. *Biology & Philosophy*, 25(4), 569–588. <https://doi.org/10.1007/s10539-010-9215-5>
- Leaders. (2007, May). Hail Linnaeus; Species inflation. *The Economist*, (383), 18.
- Leavitt, S. D., Moreau, C. S., & Lumbsch, H. T. (2015). The dynamic discipline of species delimitation: Progress toward effectively recognizing species boundaries in natural populations. In *Recent Advances in Lichenology* (pp. 11–44). Springer, New Delhi. https://doi.org/10.1007/978-81-322-2235-4_2
- Lewens, T. (2012a). Pheneticism reconsidered. *Biology & Philosophy*, 27(2), 159–177. <https://doi.org/10.1007/s10539-011-9302-2>
- Lewens, T. (2012b). Species, essence and explanation. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(4), 751–757. <https://doi.org/10.1016/j.shpsc.2012.09.013>

- Linnean Society. (2014). *Report on the Taxonomy & Systematics Plenary Meeting 11th September 2014 "Who Needs Taxonomists?"* Retrieved from www.linnean.org/meetings-and-events/events/who-needs-taxonomists
- Lipton, P. (1996). Inductive inference and its natural ground: An essay in naturalistic epistemology. *Philosophy and Phenomenological Research*, 56(2), 492–494.
- Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton, NJ: Princeton University Press.
- Longino, H. E. (1995). Gender, politics, and the theoretical virtues. *Synthese*, 104(3), 383–397. <https://doi.org/10.1007/BF01064506>
- Longino, H. E. (1996). Cognitive and non-cognitive values in science: Rethinking the dichotomy. In L. H. Nelson & J. Nelson (Eds.), *Feminism, Science, and the Philosophy of Science* (pp. 39–58). Springer Netherlands. https://doi.org/10.1007/978-94-009-1742-2_3
- Love, A. C. (2009). Typology reconfigured: From the metaphysics of essentialism to the epistemology of representation. *Acta Biotheoretica*, 57(1–2), 51–75. <https://doi.org/10.1007/s10441-008-9059-4>
- Ludwig, D. (forthcoming). Letting go of “natural kind”. Towards a multidimensional framework of non-arbitrary classification. *Philosophy of Science*, 0((-Not available-)). <https://doi.org/10.1086/694835>
- Ludwig, D. (2015). Ontological choices and the value-free ideal. *Erkenntnis*, 1–20. <https://doi.org/10.1007/s10670-015-9793-3>
- Ma, J.-X., Li, Y.-N., Vogl, C., Ehrendorfer, F., & Guo, Y.-P. (2010). Allopolyploid speciation and ongoing backcrossing between diploid progenitor and tetraploid progeny lineages in the *Achillea millefolium* species complex: Analyses of single-copy nuclear genes and genomic AFLP. *BMC Evolutionary Biology*, 10, 100. <https://doi.org/10.1186/1471-2148-10-100>
- Mace, G. M. (2004). The role of taxonomy in species conservation. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 359(1444), 711–719. <https://doi.org/10.1098/rstb.2003.1454>
- MacLeod, M. (2013). Limitations of natural kind talk in the life sciences: Homology and other cases. *Biological Theory*, 7(2), 109–120. <https://doi.org/10.1007/s13752-012-0079-6>
- MacLeod, M., & Reydon, T. A. C. (2013). Natural kinds in philosophy and in the life sciences: Scholastic twilight or new dawn? *Biological Theory*, 7(2), 89–99. <https://doi.org/10.1007/s13752-012-0080-0>
- Maddison, W. P. (1997). Gene Trees in Species Trees. *Systematic Biology*, 46(3), 523–536. <https://doi.org/10.1093/sysbio/46.3.523>
- Magnus, P. D. (2012). *Scientific Enquiry and Natural Kinds*. London: Palgrave Macmillan UK.
- Magnus, P. D. (2014, December 1). Epistemic categories and causal kinds. Retrieved from https://scholarsarchive.library.albany.edu/cas_philosophy_scholar/22
- Magnus, P. D. (2015). Taxonomy, ontology, and natural kinds. *Synthese*, 1–13. <https://doi.org/10.1007/s11229-015-0785-2>
- Mallet, J. (1995). A species definition for the modern synthesis. *Trends in Ecology & Evolution*, 10(7), 294–299. [https://doi.org/10.1016/0169-5347\(95\)90031-4](https://doi.org/10.1016/0169-5347(95)90031-4)
- Mallet, J. (2001). Species concepts. In C. W. Fox & J. B. Wolf (Eds.), *Encyclopedia of biodiversity* (pp. 367–373). Oxford: Oxford University Press.

- Mallet, J. (2004). Perspectives Poulton, Wallace and Jordan: How discoveries in *Papilio* butterflies led to a new species concept 100 years ago. *Systematics and Biodiversity*, 1(4), 441–452. <https://doi.org/10.1017/S1477200003001300>
- Mallet, J. (2005). Hybridization as an invasion of the genome. *Trends in Ecology & Evolution*, 20(5), 229–237. <https://doi.org/10.1016/j.tree.2005.02.010>
- Mallet, J. (2007). Hybrid speciation. *Nature*, 446(7133), 279–283. <https://doi.org/10.1038/nature05706>
- Mallet, J. (2008). Hybridization, ecological races and the nature of species: empirical evidence for the ease of speciation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1506), 2971–2986. <https://doi.org/10.1098/rstb.2008.0081>
- Mallet, J. (2010). Why was Darwin's view of species rejected by twentieth century biologists? *Biology & Philosophy*, 25(4), 497–527. <https://doi.org/10.1007/s10539-010-9213-7>
- Mallet, J., Beltrán, M., Neukirchen, W., & Linares, M. (2007). Natural hybridization in heliconiine butterflies: The species boundary as a continuum. *BMC Evolutionary Biology*, 7, 28. <https://doi.org/10.1186/1471-2148-7-28>
- Mallet, J., Besansky, N., & Hahn, M. W. (2015). How reticulated are species? *BioEssays*, 140–149. <https://doi.org/10.1002/bies.201500149>
- Martínez-Solano, I., Gonçalves, H. A., Arntzen, J. W., & García-París, M. (2004). Phylogenetic relationships and biogeography of midwife toads (Discoglossidae: Alytes). *Journal of Biogeography*, 31(4), 603–618. <https://doi.org/10.1046/j.1365-2699.2003.01033.x>
- Mayden, R. L. (1999). Consilience and a hierarchy of species concepts: Advances toward closure on the species puzzle. *Journal of Nematology*, 31(2), 95.
- Mayo, D. G. (1996). *Error and the Growth of Experimental Knowledge*. Chicago, IL: University of Chicago Press.
- Mayr, E. (1982). *The Growth of Biological Thought: Diversity, Evolution, and Inheritance*. Cambridge, MA: Harvard University Press.
- McMullin, E. (1982). Values in science. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1982(2), 3–28. <https://doi.org/10.1086/psaprocbienmeetp.1982.2.192409>
- Merriam, C. H. (1918). Review of the grizzly and big brown bears of North America (genus *Ursus*) with description of a new genus, *Vetularctos*. *North American Fauna*, 1–137.
- Meyer, N., Meierott, L., Schuwerk, H., & Angerer, O. (2005). *Beiträge zur Gattung Sorbus in Bayern*. Munich.
- Mishler, B. D. (1999). Getting rid of species. In R. N. Brandon (Ed.), *Species: New Interdisciplinary Essays* (pp. 307–315). Cambridge, MA: MIT Press.
- Mishler, B. D., & Brandon, R. N. (1987). Individuality, pluralism, and the phylogenetic species concept. *Biology and Philosophy*, 2(4), 397–414. <https://doi.org/10.1007/BF00127698>
- Mishler, B. D., & Donoghue, M. (1982). Species concepts: A case for pluralism. *Systematic Zoology*, 31, 491–503.
- Mishler, B. D., & Theriot, E. (2000). The phylogenetic species concept (Sensu Mishler and Theriot). In Q. D. Wheeler & N. I. Platnick (Eds.), *Species Concepts and Phylogenetic Theory: A Debate* (pp. 44–54). New York: Columbia University Press.
- Mitchell, S. (2004). The prescribed and proscribed values in science policy. In P. Machamer & G. Wolters (Eds.), *Science, values, and objectivity*. Pittsburgh, Pa: University of Pittsburgh Press.

- Mittermeier, R. A., Turner, W. R., Larsen, F. W., Brooks, T. M., & Gascon, C. (2011). Global biodiversity conservation: The critical role of hotspots. In F. E. Zachos & J. C. Habel (Eds.), *Biodiversity Hotspots* (pp. 3–22). Springer, Berlin, Heidelberg.
https://doi.org/10.1007/978-3-642-20992-5_1
- Monzón, J., Kays, R., & Dykhuizen, D. E. (2014). Assessment of coyote–wolf–dog admixture using ancestry-informative diagnostic SNPs. *Molecular Ecology*, *23*(1), 182–197.
<https://doi.org/10.1111/mec.12570>
- Mora, C., Tittensor, D. P., Adl, S., Simpson, A. G. B., & Worm, B. (2011). How many species are there on earth and in the ocean? *PLoS Biol*, *9*(8), e1001127.
<https://doi.org/10.1371/journal.pbio.1001127>
- Morrison III, W. R., Lohr, J. L., Duchon, P., Wilches, R., Trujillo, D., Mair, M., & Renner, S. S. (2009). The impact of taxonomic change on conservation: Does it kill, can it save, or is it just irrelevant? *Biological Conservation*, *142*(12), 3201–3206.
<https://doi.org/10.1016/j.biocon.2009.07.019>
- Morrison, M. (2014). Values and uncertainty in simulation models. *Erkenntnis*, *79*(5), 939–959. <https://doi.org/10.1007/s10670-013-9537-1>
- Naish, D. (2013). Taxonomic vandalism and the Raymond Hoser problem [Web log comment]. Retrieved February 9, 2017, from <https://blogs.scientificamerican.com/tetrapod-zoology/taxonomic-vandalism-and-hoser/>
- Nelson, G., & Platnick, N. I. (1981). *Systematics and biogeography: cladistics and vicariance*. New York, NY: Columbia University Press.
- Neurath, O. (1983). The lost wanderers of Descartes and the auxiliary motive. In R. S. Cohen & M. Neurath (Eds.), *Philosophical Papers 1913–1946* (pp. 1–12). Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-009-6995-7_1
- Nixon, K. C., & Wheeler, Q. D. (1990). An amplification of the phylogenetic species concept. *Cladistics*, *6*(3), 211–223. <https://doi.org/10.1111/j.1096-0031.1990.tb00541.x>
- Okasha, S. (2002). Darwinian metaphysics: Species and the question of essentialism. *Synthese*, *131*(2), 191–213. <https://doi.org/10.1023/A:1015731831011>
- O’Malley, M. A. (2010). Ernst Mayr, the tree of life, and philosophy of biology. *Biology & Philosophy*, *25*(4), 529–552. <https://doi.org/10.1007/s10539-010-9214-6>
- O’Malley, M. A. (2014). *Philosophy of Microbiology*. Cambridge, UK: Cambridge University Press.
- O’Malley, M. A., & Dupré, J. (2007). Size doesn’t matter: Towards a more inclusive philosophy of biology. *Biology & Philosophy*, *22*(2), 155–191.
<https://doi.org/10.1007/s10539-006-9031-0>
- Pabijan, M., Crottini, A., Reckwell, D., Irisarri, I., Hauswaldt, J. S., & Vences, M. (2012). A multigene species tree for Western Mediterranean painted frogs (*Discoglossus*). *Molecular Phylogenetics and Evolution*, *64*(3), 690–696.
<https://doi.org/10.1016/j.ympev.2012.05.017>
- Padial, J. M., Castroviejo-Fisher, S., Köhler, J., Vilà, C., Chaparro, J. C., & De la Riva, I. (2009). Deciphering the products of evolution at the species level: the need for an integrative taxonomy. *Zoologica Scripta*, *38*(4), 431–447. <https://doi.org/10.1111/j.1463-6409.2008.00381.x>
- Padial, J. M., Miralles, A., De la Riva, I., & Vences, M. (2010). The integrative future of taxonomy. In *Frontiers in Zoology*, *7*(1), 16.

- Pamilo, P., & Nei, M. (1988). Relationships between gene trees and species trees. *Molecular Biology and Evolution*, 5(5), 568–583.
<https://doi.org/10.1093/oxfordjournals.molbev.a040517>
- Parker, W. (2014). Values and uncertainties in climate prediction, revisited. *Studies in History and Philosophy of Science Part A*, 46(Supplement C), 24–30.
<https://doi.org/10.1016/j.shpsa.2013.11.003>
- Potochnik, A. (2015). The diverse aims of science. *Studies in History and Philosophy of Science Part A*, 53, 71–80. <https://doi.org/10.1016/j.shpsa.2015.05.008>
- Poulton, E. B. (1903). What is a species? *Proceedings of the Entomological Society of London*, lxxvii–cxvi.
- Putnam, H. (2002). *The Collapse of the Fact/value Dichotomy and Other Essays*. Cambridge, MA: Harvard University Press.
- Queiroz, K., & Donoghue, M. J. (1988). Phylogenetic systematics and the species problem. *Cladistics*, 4(4), 317–338. <https://doi.org/10.1111/j.1096-0031.1988.tb00518.x>
- Real, R., Barbosa, A., & Garci-Paris, M. (2005). Distinguishing the distributions of two cryptic frogs (Anura: Discoglossidae) using molecular data and environmental modeling. *Canadian Journal of Zoology*, 83(4), 536–545.
- Reydon, T. A. C. (2003). Discussion: Species are individuals—Or are they? *Philosophy of Science*, 70(1), 49–56. <https://doi.org/10.1086/367868>
- Reydon, T. A. C. (2005). On the nature of the species problem and the four meanings of “species.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(1), 135–158.
<https://doi.org/10.1016/j.shpsc.2004.12.004>
- Reydon, T. A. C. (2006). Generalizations and kinds in natural science: The case of species. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 37(2), 230–255.
<https://doi.org/10.1016/j.shpsc.2006.03.003>
- Reydon, T. A. C. (2009a). Do the life sciences need natural kinds? *Croatian Journal of Philosophy*, IX(26), 167–190.
- Reydon, T. A. C. (2009b). Gene names as proper names of individuals: An assessment. *The British Journal for the Philosophy of Science*, 60(2), 409–432.
<https://doi.org/10.1093/bjps/axp002>
- Reydon, T. A. C. (2015a). From a zooming-in model to a co-creation model: Toward a more dynamic account of classification and kinds. In C. Kendig (Ed.), *Natural Kinds and Classification in Scientific Practice* (pp. 59–73). Routledge.
- Reydon, T. A. C. (2015b). Natural kinds no longer are what they never were. *Metascience*, 24(2), 259–264.
- Ricard, G., McEwan, N. R., Dutilh, B. E., Jouany, J.-P., Macheboeuf, D., Mitsumori, M., ... Huynen, M. A. (2006). Horizontal gene transfer from Bacteria to rumen Ciliates indicates adaptation to their anaerobic, carbohydrates-rich environment. *BMC Genomics*, 7, 22. <https://doi.org/10.1186/1471-2164-7-22>
- Richards, T. A., Dacks, J. B., Jenkinson, J. M., Thornton, C. R., & Talbot, N. J. (2006). Evolution of filamentous plant pathogens: Gene exchange across eukaryotic kingdoms. *Current Biology*, 16(18), 1857–1864. <https://doi.org/10.1016/j.cub.2006.07.052>
- Ridley, M. (1989). The cladistic solution to the species problem. *Biology and Philosophy*, 4(1), 1–16. <https://doi.org/10.1007/BF00144036>

- Rieppel, O. (2010). Species monophyly. *Journal of Zoological Systematics and Evolutionary Research*, 48(1), 1–8. <https://doi.org/10.1111/j.1439-0469.2009.00545.x>
- Rieseberg, L. H., & Burke, J. M. (2001). The biological reality of species: Gene flow, selection, and collective evolution. *Taxon*, 50(1), 47–67. <https://doi.org/10.2307/1224511>
- Romero, D., Olivero, J., Márquez, A. L., Báez, J. C., & Real, R. (2014). Uncertainty in distribution forecasts caused by taxonomic ambiguity under climate change scenarios: a case study with two newt species in mainland Spain. *Journal of Biogeography*, 41(1), 111–121. <https://doi.org/10.1111/jbi.12189>
- Rooney, P. (1992). On values in science: Is the epistemic/non-epistemic distinction useful? *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1992, 13–22.
- Rosenberg, A. (1985). *The Structure of Biological Science*. Cambridge, UK: Cambridge University Press.
- Rosenberg, A. (1994). *Instrumental Biology, Or The Disunity of Science*. Chicago: University of Chicago Press.
- Rosenberg, N. A. (2002). The probability of topological concordance of gene trees and species trees. *Theoretical Population Biology*, 61(2), 225–247. <https://doi.org/10.1006/tpbi.2001.1568>
- Rosselló-Mora, R., & Amann, R. (2001). The species concept for prokaryotes. *FEMS Microbiology Reviews*, 25(1), 39–67.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20(1), 1–6.
- Ruedas, L., & Brown, R. M. (2007, May 17). Hail, Linneaus. Message posted to <http://www.economist.com/node/20013964/comments>.
- Ruse, M. (1987). Biological species: Natural kinds, individuals, or what? *The British Journal for the Philosophy of Science*, 38(2), 225–242.
- Sangster, G. (2009). Increasing numbers of bird species result from taxonomic progress, not taxonomic inflation. *Proceedings of the Royal Society B: Biological Sciences*, 276(1670), 3185–3191. <https://doi.org/10.1098/rspb.2009.0582>
- Sangster, G. (2014). The application of species criteria in avian taxonomy and its implications for the debate over species concepts: Application of species criteria in practice. *Biological Reviews*, 89(1), 199–214. <https://doi.org/10.1111/brv.12051>
- Satler, J. D., Carstens, B. C., & Hedin, M. (2013). Multilocus species delimitation in a complex of morphologically conserved trapdoor spiders (Mygalomorphae, Antrodiaetidae, Aliatypus). *Systematic Biology*, 62(6), 805–823. <https://doi.org/10.1093/sysbio/syt041>
- Schlick-Steiner, B. C., Steiner, F. M., Seifert, B., Stauffer, C., Christian, E., & Crozier, R. H. (2010). Integrative taxonomy: A multisource approach to exploring biodiversity. *Annual Review of Entomology*, 55(1), 421–438. <https://doi.org/10.1146/annurev-ento-112408-085432>
- Sillero, N., Brito, J. C., Skidmore, A. K., & Toxopeus, A. G. (2009). Biogeographical patterns derived from remote sensing variables: the amphibians and reptiles of the Iberian Peninsula. *Amphibia-Reptilia*, 30(2), 185–206. <https://doi.org/10.1163/156853809788201207>
- Simpson, G. (1945). The principles of classification and a classification of mammals. *Bulletin of the American Museum of Natural History*, 85, xvi+350.

- Sites, J. W., & Marshall, J. C. (2004). Operational criteria for delimiting species. *Annual Review of Ecology, Evolution, and Systematics*, 35, 199–227.
- Slater, M. H. (2014). Natural kindness. *The British Journal for the Philosophy of Science*, 66(2), 375–411. <https://doi.org/10.1093/bjps/axt033>
- 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 Part A*, 61, 1–10. <https://doi.org/10.1016/j.shpsa.2017.01.001>
- Sluys, R. (2013). The unappreciated, fundamentally analytical nature of taxonomy and the implications for the inventory of biodiversity. *Biodiversity and Conservation*, 22(4), 1095–1105. <https://doi.org/10.1007/s10531-013-0472-x>
- Sneath, P. H. A., & Sokal, R. R. (1973). *Numerical Taxonomy: The Principles and Practice of Numerical Classification*. San Fransisco, CA: W.H. Freeman.
- Sober, Elliott. (1980). Evolution, population thinking, and essentialism. *Philosophy of Science East Lansing, Mich.*, 47(3), 350–383.
- Sober, Elliott. (1984). Sets, species, and evolution: Comments on Philip Kitcher’s “Species.” *Philosophy of Science*, 51(2), 334–341.
- Sober, Elliott. (2000). *Philosophy of Biology (2nd edition)*. Boulder, CO: Westview Press.
- Soucy, S. M., Huang, J., & Gogarten, J. P. (2015). Horizontal gene transfer: Building the web of life. *Nature Reviews Genetics*, 16(8), 472–482. <https://doi.org/10.1038/nrg3962>
- Stanford, P. K. (1995). For pluralism and against realism about species. *Philosophy of Science*, 62(1), 70–91.
- Steel, D. (2010). Epistemic values and the argument from inductive risk. *Philosophy of Science*, 77(1), 14–34. <https://doi.org/10.1086/650206>
- Steel, D. (2013). Acceptance, values, and inductive risk. *Philosophy of Science*, 80(5), 818–828. <https://doi.org/10.1086/673936>
- Steel, D. (2017). Qualified epistemic priority: Comparing two approaches to values in science. In D. Steel & K. C. Elliott (Eds.), *Current controversies in values and science* (pp. 49–63). New York: Routledge.
- Steel, D., & Whyte, K. P. (2012). Environmental justice, values, and scientific expertise. *Kennedy Institute of Ethics Journal*, 22(2), 163–182. <https://doi.org/10.1353/ken.2012.0010>
- Steele, K. (2012). The scientist qua policy advisor makes value judgments. *Philosophy of Science*, 79(5), 893–904. <https://doi.org/10.1086/667842>
- Stegenga, J. (2016). Population pluralism and natural selection. *The British Journal for the Philosophy of Science*, 67(1), 1–29. <https://doi.org/10.1093/bjps/axu003>
- Steppan, S. (2007, May 17). Hail, Linneaus. Message posted to <http://www.economist.com/node/20013964/comments>.
- Tattersall, I. (2007). Madagascar’s Lemurs: Cryptic diversity or taxonomic inflation? *Evolutionary Anthropology: Issues, News, and Reviews*, 16(1), 12–23. <https://doi.org/10.1002/evan.20126>
- Tautz, D., Arctander, P., Minelli, A., Thomas, R. H., & Vogler, A. P. (2003). A plea for DNA taxonomy. *Trends in Ecology & Evolution*, 18(2), 70–74. [https://doi.org/10.1016/S0169-5347\(02\)00041-1](https://doi.org/10.1016/S0169-5347(02)00041-1)
- Timmis, J. N., Ayliffe, M. A., Huang, C. Y., & Martin, W. (2004). Endosymbiotic gene transfer: Organelle genomes forge eukaryotic chromosomes. *Nature Reviews Genetics*, 5(2), 123–135. <https://doi.org/10.1038/nrg1271>

- Trochet, A., Moulherat, S., Calvez, O., Stevens, V. M., Clobert, J., & Schmeller, D. S. (2014). A database of life-history traits of European amphibians. *Biodiversity Data Journal*, 2, e4123. <https://doi.org/10.3897/BDJ.2.e4123>
- Valen, L. V. (1976). Ecological species, multispecies, and oaks. *Taxon*, 25(2/3), 233–239. <https://doi.org/10.2307/1219444>
- Velasco, J. D. (2009). When monophyly is not enough: Exclusivity as the key to defining a phylogenetic species concept. *Biology & Philosophy*, 24(4), 473–486. <https://doi.org/10.1007/s10539-009-9151-4>
- Velasco, J. D. (2010). Species, genes, and the tree of life. *The British Journal for the Philosophy of Science*, 61(3), 599–619. <https://doi.org/10.1093/bjps/axp051>
- Velo-Antón, G., Garci-Paris, M., & Martínez-Solano, I. (2008). β -fibrinogen intron 7 variation in *Discoglossus* (Anura: Discoglossidae): Implications for the taxonomic assessment of morphologically cryptic species. *Amphibia-Reptilia*, 29(4), 523–533. <https://doi.org/10.1163/156853808786230460>
- Vences, M., & Glaw, F. (1996). Further investigations on *Discoglossus* bioacoustics. *Amphibia-Reptilia*, 17(4), 333–340. <https://doi.org/10.1163/156853896X00063>
- Vences, M., Pous, P. de, Nicolas, V., Díaz-Rodríguez, J., Donaire, D., Hugemann, K., ... Tejedo, M. (2014). New insights on phylogeography and distribution of painted frogs (*Discoglossus*) in northern Africa and the Iberian Peninsula. *Amphibia-Reptilia*, 35(3), 305–320. <https://doi.org/10.1163/15685381-00002954>
- Venema, D. (2013). Evolution Basics: Species Trees, Gene Trees and Incomplete Lineage Sorting [Web log comment] [Biologos]. Retrieved from <http://biologos.org/blogs/dennis-venema-letters-to-the-duchess/evolution-basics-species-trees-gene-trees-and-incomplete-lineage-sorting>
- vom Saal, F. S., & Hughes, C. (2005). An extensive new literature concerning low-dose effects of bisphenol A shows the need for a new risk assessment. *Environmental Health Perspectives*, 113(8), 926–933. <https://doi.org/10.1289/ehp.7713>
- Waugh, J. (2007). DNA barcoding in animal species: Progress, potential and pitfalls. *BioEssays*, 29(2), 188–197. <https://doi.org/10.1002/bies.20529>
- Wells, R. W., & Wellington, R. (1984). A synopsis of the class Reptilia in Australia. *Australian Journal of Herpetology*, 1, 73–129.
- Wells, R. W., & Wellington, R. (1985). A classification of the Amphibia and Reptilia of Australia. *Australian Journal of Herpetology, Suppl. Ser. 1*, 1–61.
- Wheeler, Q. D. (2008). *The New Taxonomy*. Florida: CRC Press.
- Wheeler, Q. D. (2014). Are reports of the death of taxonomy an exaggeration? *New Phytologist*, 201(2), 370–371. <https://doi.org/10.1111/nph.12612>
- Wiley, E. O. (1978). The Evolutionary Species Concept reconsidered. *Systematic Biology*, 27(1), 17–26. <https://doi.org/10.2307/2412809>
- Wilholt, T. (2009). Bias and values in scientific research. *Studies in History and Philosophy of Science Part A*, 40(1), 92–101. <https://doi.org/10.1016/j.shpsa.2008.12.005>
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64(2), 233–253. <https://doi.org/10.1093/bjps/axs007>
- Will, K., Mishler, B., & Wheeler, Q. D. (2005). The perils of DNA barcoding and the need for integrative taxonomy. *Systematic Biology*, 54(5), 844–851. <https://doi.org/10.1080/10635150500354878>

- Williams, D., Wüster, W., & Fry, B. G. (2006). The good, the bad and the ugly: Australian snake taxonomists and a history of the taxonomy of Australia's venomous snakes. *Toxicon*, 48(7), 919–930. <https://doi.org/10.1016/j.toxicon.2006.07.016>
- Willis, S. C. (2017). One species or four? Yes!...and, no. Or, arbitrary assignment of lineages to species obscures the diversification processes of Neotropical fishes. *PLOS ONE*, 12(2), e0172349. <https://doi.org/10.1371/journal.pone.0172349>
- Wilson, J. R., Procheş, Ş., Braschler, B., Dixon, E. S., & Richardson, D. M. (2007). The (bio)diversity of science reflects the interests of society. *Frontiers in Ecology and the Environment*, 5(8), 409–414. <https://doi.org/10.1890/060077.1>
- Wilson, R. A. (1999). Realism, essence, and kind: Resuscitating species essentialism? In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays* (pp. 187–207). Cambridge, MA: MIT Press.
- Wilson, R. A., Barker, M. J., & Brigandt, I. (2007). When traditional essentialism fails: Biological natural kinds. *Philosophical Topics*, 35(1/2), 189–215.
- Winsberg, E. (2012). Values and uncertainties in the predictions of global climate models. *Kennedy Institute of Ethics Journal*, 22(2), 111–137. <https://doi.org/10.1353/ken.2012.0008>
- Wüster, W., Bush, B., Keogh, J. S., O'Shea, M., & Shine, R. (2001). Taxonomic contributions in the “amateur” literature: Comments on recent descriptions of new genera and species by Raymond Hoser. *Litteratura Serpentina*, 21(3), 67–91.
- Zachos, F. E. (2014). Taxonomic inflation, the Phylogenetic Species Concept and lineages in the Tree of Life – A cautionary comment on species splitting. *Journal of Zoological Systematics and Evolutionary Research*, n/a-n/a. <https://doi.org/10.1111/jzs.12088>
- Zachos, F. E., Apollonio, M., Bärmann, E. V., Festa-Bianchet, M., Göhlich, U., Habel, J. C., ... Suchentrunk, F. (2013). Species inflation and taxonomic artefacts—A critical comment on recent trends in mammalian classification. *Mammalian Biology - Zeitschrift Für Säugetierkunde*, 78(1), 1–6. <https://doi.org/10.1016/j.mambio.2012.07.083>
- Zachos, F. E., & Habel, J. C. (2011). *Biodiversity Hotspots: Distribution and Protection of Conservation Priority Areas*. Heidelberg: Springer Science & Business Media.
- Zangari, F., Cimmaruta, R., & Nascetti, G. (2006). Genetic relationships of the western Mediterranean painted frogs based on allozymes and mitochondrial markers: Evolutionary and taxonomic inferences (Amphibia, Anura, Discoglossidae). *Biological Journal of the Linnean Society*, 87(4), 515–536. <https://doi.org/10.1111/j.1095-8312.2006.00585.x>