



UNIVERSITAT^{DE}
BARCELONA

Naturalness, Amelioration, and Sex

Iñigo Valero Alzaga



Aquesta tesi doctoral està subjecta a la llicència **Reconeixement 4.0. Espanya de Creative Commons.**

Esta tesis doctoral está sujeta a la licencia **Reconocimiento 4.0. España de Creative Commons.**

This doctoral thesis is licensed under the **Creative Commons Attribution 4.0. Spain License.**



UNIVERSITAT DE
BARCELONA

PhD Dissertation

Naturalness, Amelioration, and Sex

Iñigo Valero Alzaga

Director: Dan López de Sa Medina

Codirector: María de Ponte Azkarate

Tutor: Genoveva Martí Campillo

PhD Program: Cognitive Science and Language

Branch: Analytic Philosophy

Facultat de Filosofia

March 2022

Abstract

This PhD Dissertation explores a variety of topics in relation with metaphysical naturalness and the assessment of our conceptual and classificatory practices. It is divided in three self-standing—albeit connected—papers.

In the first paper, “Radical Projectibilism: A Gradual Account of Natural Kindness”, I introduce an original account of naturalness according to which the naturalness of kinds is to be identified with their degree of projectibility. A distinctive characteristic of this proposal—which I label *Radical Projectibilism*—is that, by identifying naturalness with a gradual property such as projectibility, the notion of naturalness itself becomes one of degree. I argue that understanding naturalness in a gradual way not only appropriately counters the relevant notion of arbitrariness but, moreover, brings important advantages over dichotomic alternatives.

In the second paper, “Ameliorating at the Joints. A Permissive Normative Framework for Conceptual Engineering”, I defend an original normative framework for conceptual amelioration. I do so through a reply to a recent proposal by Mona Simion (2017) who articulates an “epistemic limiting procedure” for conceptual revision according to which it is not permissible to ameliorate concepts if the revision involves an “epistemic loss”. I argue, instead, for a more permissive alternative that allows for epistemic losses, provided they are duly compensated by the moral or pragmatic gains that motivate the revision.

Finally, in the third paper, “Is Sex Binary? Yes and No, and a Terminological Proposal”, I engage a specific case study involving sex categories. I argue that, appearances notwithstanding, the current dispute on the binary or non-binary nature of sex does not involve a factual disagreement. Moreover, I advance a terminological normative proposal according to which ‘binary’ should be replaced by ‘dimorphic’ in the characterization of sex. The main reason in favor of this move, I argue, is that ‘dimorphic’ preserves the important function that ‘binary’ presumably fulfils while avoiding the problems of its counterpart.

Resumen

Esta Tesis Doctoral explora una variedad de temas en relación con la naturalidad de clases de objetos y la evaluación de nuestras prácticas conceptuales y clasificatorias. Se compone de tres artículos independientes, aunque relacionados.

En el primer artículo, “Radical Projectibilism: A Gradual Account of Natural Kindness”, presento una caracterización original de la naturalidad según la cual la naturalidad de las clases debe identificarse con su grado de proyectabilidad. Una característica distintiva de esta propuesta es que, al identificar la naturalidad con una propiedad gradual como la proyectabilidad, la noción misma de naturalidad se convierte en una de *grado*. Sostengo que entender la naturalidad de manera gradual no solo contrarresta adecuadamente la noción relevante de arbitrariedad, sino que, además, aporta numerosas ventajas teóricas de las que alternativas dicotómicas no pueden dar cuenta.

En el segundo artículo, “Ameliorating at the Joints. A Permissive Normative Framework for Conceptual Engineering”, defiendo un marco normativo original para la revisión conceptual. Hago esto a través de una réplica a una propuesta reciente de Mona Simion (2017) quien articula un marco normativo según el cual no es permisible revisar conceptos si la revisión implica una “pérdida epistémica”. Frente a la propuesta de Simion, abogo por una alternativa más permisiva que tolere pérdidas epistémicas, siempre y cuando sean debidamente compensadas por los beneficios morales o pragmáticos que motivan la revisión.

Finalmente, en el tercer artículo, ¿“Is Sex Binary? Yes and No, and a Terminological Proposal”, abordo un caso específico que involucra las categorías correspondientes al sexo biológico. Argumento que, a pesar de las apariencias, el debate actual entorno a la naturaleza binaria o no binaria del sexo biológico no implica un desacuerdo con respecto a los hechos biológicos relevantes. Además, en este artículo presento una propuesta normativa terminológica según la cual el sexo no debería caracterizarse como 'binario' sino como 'dimórfico' (en inglés). Argumento que la razón principal a favor de este movimiento terminológico es que 'dimórfico' cumple con la función con la que el término 'binario' también cumple, al tiempo que evita los problemas asociados a este último término.

Acknowledgments

I am grateful to all the people who, in some way or another, have helped me through the process of writing this dissertation.

I would first like to thank my main supervisor, Dan López de Sa, whose infinite patience, dedication, and expertise have been essential in bringing about this work. Thanks also to my tutor, Genoveva Martí, and co-supervisor, María de Ponte, for their support and assistance.

I would also like to thank my colleagues from the University of Barcelona and Pompeu Fabra University, with a particular mention to the attendants to the weakly Social Metaphysics Seminar organized by Dan López de Sa, where I had the opportunity to present various parts of this work and receive valuable feedback. Thanks to Aarón Álvarez, Adam Sennet, Alba Lojo, Bianca Cepollaro, Enrico Terrone, Esa Díaz-León, Filippo Contesi, Giuliano Torrenco, Giulio Pietroiusti, Javier Suárez, John Horden, Josean Lozano, Lisa Vogt, Manolo García-Carpintero, Manolo Martínez, Manuel Gustavo Isaac, Marta Jorba, Pablo Magaña, Samuele Chilovi, and Teresa Marques.

I am also grateful to my colleagues and teachers from the University of Exeter, where I had the opportunity to spend three months as a visiting researcher. Thanks to Adam Toon, Andrea Raimondi, Fabrizia Giulia, John Dupré, Lasse Bergmann, Lucy Osler, Sabina Leonelli, and Staffan Müller-Wille.

Finally, I would like to thank Miriam Masri, my best friend and life companion. Thank you for your support, wise counsel, and loud laughter.

I would like to dedicate this work to Alexander Iturria, Núria Gran, and Rubén Hervás. Thank you for the shelter.

Contents

Abstract.....	3
Resumen	5
Acknowledgments	7
Introduction	11
Radical Projectibilism: A Gradual Account of Natural Kindness.....	19
1 Introduction	19
2 Two Desiderata for a Natural Kind Theory	20
3 Failures to Preserve the Science Constraint: Towards Inclusiveness	25
3.1 Natural Kind Essentialism	25
3.2 Homeostatic Property Clusters	26
3.3 Bare Property Clusters	27
3.4 Beyond Bare Property Clusters.....	29
4 Radical Projectibilism	30
5 The Challenge of Projectibility-Based Accounts of Naturalness: Meeting the Contrast Desideratum	32
6 Putting the Gradual View to Work: Scientific Conceptual Progress	37
7 Necessity Worries	39
7.1 Superheavy Element 117	41
7.2 Magnus’s Polymorphism	44
8 Conclusion.....	49
Ameliorating at the Joints: A Permissive Normative Framework for Conceptual Engineering.....	51
1 Introduction	51
2 Simion’s Proposal	52
3 Epistemic Loss: Three Readings	54
3.1 Epistemic Loss as Ceasing to Represent a Real Kind.....	54
3.2 Epistemic Loss as Ceasing to Represent a Natural Kind.....	55
3.3 Epistemic Loss as Less Natural	58
4 An Alternative Normative Framework	59
5 Two Case Studies: MARRIAGE and SPECIES	61
5.1 The Revision of MARRIAGE	62
5.2 Conceptual Choice: The Concept of SPECIES.....	67
6 Conclusion.....	69
Is Sex Binary? Yes and No, and a Terminological Proposal	71

1	Introduction	71
2	The dispute	72
2.1	The relevant biological facts	72
2.2	The contenders	74
2.3	The tension around ‘binary’	81
3	A normative proposal: doing without ‘binary’	82
3.1	The harm of conveying EXHAUSTIVITY: the importance of visibility for stigmatized identities	83
3.2	Drop ‘binary’, use ‘dimorphic’ instead.....	84
4	Conclusion.....	86
	Conclusion	87
	References	89

Introduction

We humans have an inescapable tendency to categorize and organize the things we find in the world. We group particulars together whenever we deem them similar in relevant ways, and label the resulting groupings accordingly. This complex activity, which generally involves the interaction of kinds, concepts, and terms, is crucial, as it allows us to conceptualize reality, communicate, make inferences and predictions, build explanations, formulate theories, and so on. More mundanely, it is also through this classificatory practice that we are able to segregate people, decide which organisms deserve environmental protection, do gardening, play chess, and count the planets of the Solar System, among countless other things. As William James vividly puts it:

Kinds, and sameness of kind—what colossally useful *denkmittel* for finding our way among the many! The manyness might conceivably have been absolute. Experiences might have all been singulars, no one of them occurring twice. In such a world logic would have had no application; for kind and sameness of kind are logic's only instruments. Once we know that whatever is of a kind is also of that kind's kind, we can travel through the universe as if with seven-league boots. (James 1907: 179)

Now, as philosophers, we may ask whether there are better or worse ways of carrying out this classificatory activity. This question, though, does not seem to have a univocal answer, as our classificatory practices are multifaceted activities that can be assessed according to various different criteria.

As a first approximation, we could try to distinguish between *epistemic* and *non-epistemic* values or criteria. Although this distinction is commonly applied to discussions of underdetermination and theory choice, transferring it to the domain of classificatory choice requires certain adaptations. Indeed, in the context of theory choice, epistemic values are identified, roughly, as those whose instantiation makes a theory more likely to be true. As such, values such as simplicity, explanatory power and breadth of scope are considered *epistemic* in virtue of being *truth-conducive* (Steel 2010).

This criterion for differentiating epistemic and non-epistemic values, however, does not work—at least not straightforwardly—for classificatory choice. For one thing,

categories¹ and classificatory systems—lacking propositional content—do not have truth conditions and, as such, cannot be assessed in terms of their truth or falsity. Accordingly, one may wonder whether there is any epistemic standard for classificatory choice that could play a role similar to the one that truth plays for theory choice.

A good candidate to play this role, I submit, is the notion of *naturalness* (or *natural kinds*). Although it is difficult to say anything about natural kinds without committing to a specific and thus potentially disputable view, we can at least point out that natural kinds are often metaphorically described as *carving nature at its joints*, and that its paradigmatic instances are generally taken to be found among well-established scientific categories, in particular among those of chemistry and physics. We can also add, perhaps less neutrally, that any philosophically significant notion of naturalness has *arbitrariness* as its counterpart, and that it does not contrast, instead, with *artificial* or *social*. Indeed, the fact that certain entities are the result of human activities does not mark a significant difference of the sort that we are trying to capture with the philosophical notion of naturalness. Think, for instance, of those elements of the periodic table which do not occur “naturally” in our world, but instead need to be synthetically produced. Clearly, their artificiality notwithstanding, these elements are as good candidates for paradigmatic natural kinds as their “naturally” occurring neighbors in the periodic table. Natural kinds, therefore, are best conceived, not as contrasting with social or artificial categories such as *money* and *tennessine*, but rather, as contrasting with arbitrary groupings and random collections such as *discovered-on-a-Tuesday* and *heavier-than-1kg*.

On this approach, then, the *epistemic* success of our classificatory practices depends on whether or not they track natural kinds. Of course, this is still somewhat vague, as the notion of naturalness is itself subject to important disagreements. Still, the current natural kind literature, although diverse, seems unified enough when it comes to ascribing natural kinds a certain *epistemic value* (Kendig & Grey 2019). Accordingly, identifying naturalness as the epistemic standard for our classificatory practices seems to be in line with this common attitude towards natural kinds.

Now, assuming that naturalness is indeed the best indicator of our categories’ epistemic success, it is crucial to acknowledge that our classificatory and conceptual

¹ Given the preliminary nature of this introduction, I will be using ‘category’ rather loosely, to refer indiscriminately to kind-tracking terms and concepts.

practices can also be assessed from other evaluative perspectives. We might distinguish two of them.

First, we can assess our classificatory practices from a *pragmatic* perspective by focusing on how *useful* they are for achieving a given *purpose*. Granted, this way of framing this evaluative dimension is very generic, as virtually anything can count as a *purpose*. For instance, it is plausible to think of certain classificatory systems as mainly having an *epistemic* purpose (Khalidi 2013: 62), in which case their *pragmatic* success would amount to an *epistemic* success of the sort we identified above. Notice, therefore, that the pragmatic evaluative dimension should not be considered strictly *non-epistemic*. Rather, the epistemic and pragmatic dimensions must be conceived as standing in an orthogonal relation. That is, although the epistemic and pragmatic successes of certain categories might sometimes coincide, they will often not.

Indeed, formulating the pragmatic evaluative dimension in these general terms allows us to also acknowledge other, more specific, pragmatic purposes that often drive our classificatory practices, and which are clearly not epistemic. As we already suggested above, these purposes can include everyday activities such as cooking and gardening, but also more consequential goals such as segregating the sexes in competitive sports or the environmental protection of endangered species. It is interesting to acknowledge, though, that for classificatory and conceptual systems to fulfil such pragmatic purposes, they are often required to exhibit a certain degree of naturalness. In this sense, notice that completely arbitrary categories such as *discovered-on-a-Tuesday* or *heavier-than-my-head*—which enter into very few generalizations, and uninteresting ones at that—will hardly contribute to fulfilling any purpose at all, epistemic or otherwise. Richard Boyd (1999: 161) expresses this idea succinctly: “Gruified gardening would be as unsuccessful as gruified mineralogy.”

Finally, having identified the epistemic and pragmatic evaluative dimensions, we can further single out an *ethical* dimension by acknowledging that moral and political considerations might also play a crucial role in the assessment of our conceptual and classificatory practices. Indeed, by mediating our relationship with the world, our conceptual and classificatory systems have the potential to impact our everyday lives in very significant ways. Moreover, certain classificatory and conceptual systems might contribute to the perpetuation of unjust and harmful social structures, or simply have morally undesirable effects. In light of this, a comprehensive assessment of our classificatory practices must also include an ethical perspective.

Now, the way I have introduced these different evaluative dimensions has been intentionally loose. Doing otherwise would have betrayed the preliminary purpose of this introduction by committing me too soon to certain specific views that I am not yet in a position to defend. That being said, it is worth noting that the literature's discussion of these evaluative dimensions has often been carried out in a compartmentalized way. That is, while some philosophers have worked at the level of *kinds*, and have accordingly focused on the notion of *naturalness*, others have instead explored conceptual revision from an ethical perspective, without always acknowledging that characterizing naturalness in one way or another might have an impact on how conceptual revision is approached. Of course, this is not problematic in itself, as these issues are interesting and complex enough in themselves to constitute self-standing philosophical topics. As I said before, though, there are certain tangential points involving how these different dimensions interact with one another that are worth discussing.

This Ph.D. dissertation is situated at this crossroads, as it aims to explore, from a more holistic perspective, some of the interactions that might occur between the different dimensions along which our categories can be assessed. In this sense, one of the main concerns of the present work involves identifying and addressing the *tension* that may arise between the various normative forces that drive our classificatory and conceptual practices, without losing sight of the fact that alternative characterizations of naturalness might significantly condition how this tension is dealt with.

As an illustration of the sort of tension I have in mind, consider the case of the sex categories 'female' and 'male'. Given their social, political, and legal significance, and given the fact that they are commonly employed in scientific contexts such as biology or medicine, sex categories are particularly susceptible to being caught in a tension between the different evaluative dimensions I have identified above. For the sake of exposition, consider the following simplified scenario.

Imagine a world where medical doctors and biologists all agree on the definition of 'male' and 'female'. Even more remarkably, in this world philosophers have also come to agree that these two categories pick out natural kinds. At the same time, however, the usage of these sex categories is denounced by social activists for being exclusionary, in a way that is morally and politically problematic. In light of this, these social activists suggest a non-exclusionary revision of these categories, thereby solving the political and moral problem. Although many subsequently endorse the suggested revision, some scientists and philosophers complain that the revised categories do not track natural kinds,

and thus that the original ones should be preserved. While all this is going on, certain sports institutions, in order to, as they put it, “preserve fairness in competition”, segregate the sexes using their own criteria, which, interestingly, match neither the standard biological definition of the sexes, nor the revision suggested by the social activists.

Clearly, this scenario involves a normative tension that needs to be addressed. What are we, as classifiers, to do? Should we keep using the joint-carving classificatory system despite its morally undesirable consequences? Or should we instead revise our sex categories to avoid this problematic result, even if we end up with categories that no longer track natural kinds? Should we rather be pluralists about sex and accept various alternative definitions?

Before going any further, let me introduce another twist to this scenario. In this alternative case, everything remains the same except for the fact that philosophers characterize naturalness in *gradual* terms. That is, on their view, the contrast between categories such as *discovered-on-a-Tuesday* and *gold* is not a sharp contrast between a non-natural and a natural kind, but instead corresponds to the two extremes of a spectrum. As such, our imagined philosophers’ assessment of the situation varies: here, the difference between the original exclusionary categories and the revised non-exclusionary categories is not that the former track natural kinds while the latter don’t. Rather, the difference is—“merely”—that the revised ones are *less natural* than the original ones.

Now, does this alternative characterization of naturalness affect how the aforementioned normative tension is dealt with? One may wonder, for instance, whether the scientists and philosophers who objected to the non-exclusionary sex categories for not picking out natural kinds would react in the same way in this second scenario, where the revised categories do exhibit a certain degree of naturalness.²

Although these scenarios have been simplified for the sake of illustration, it is easy to see that they preserve a certain degree of similarity with real ongoing debates. As such, they serve to depict the sort of normative tension I am concerned with. Moreover,

² To insist even more on this idea, imagine yet another case in which the favored characterization of natural kinds is *very restrictive*, such that it results in no sex category (nor the original exclusionary ones, nor the revised ones) being considered natural. In this alternative scenario, philosophers consider only categories from chemistry and physics to pick out natural kinds. How would this characterization of natural kinds affect the normative tension involved in this case? Indeed, given that according to this characterization there would no longer be any transition from a natural kind to a non-natural kind, it is plausible to think that at least some of the normative tension might disappear.

they also serve to illustrate how specific characterizations of naturalness may condition how this tension is dealt with.

Now, to address these and other related questions, this Ph.D. dissertation will be divided into three self-standing—albeit connected—papers.

In the first paper, “Radical Projectibilism: A Gradual Account of Natural Kindness”, I engage with the natural kind literature and put forward an original account of natural kindness based on the notion of *projectibility*. Although this notion is often invoked by characterizations of natural kinds, the originality of my proposal derives from identifying the naturalness of kinds *exclusively* with their degree of projectibility. The most important consequence of this move, I argue, is that, given that projectibility is a *gradual* property, the notion of naturalness itself becomes one of degree. I argue that rather than being a shortcoming of my proposal, understanding naturalness gradually has significant advantages over dichotomic accounts.

Indeed, this gradual notion of naturalness will play a crucial role in the second paper, “Ameliorating at the Joints: A Permissive Framework for Conceptual Engineering”, where I introduce a novel normative framework for conceptual amelioration. I do so in the form of a reply to the proposal of Mona Simion (2017), who puts forward a rather restrictive procedure for conceptual amelioration. According to Simion’s “Epistemic Limiting Procedure”, a conceptual revision is *permissible* only insofar as it does not result in an “epistemic loss”. I argue, instead, in favor of a more permissive normative framework that allows for epistemic losses, provided they are duly compensated by the corresponding pragmatic or moral benefits that motivate the revision.

That being said, one of the most important contributions of this second paper is that it develops in more detail the notion of “epistemic loss”, which I take to be underdeveloped in Simion’s work. After considering various alternatives, I argue that the best way to understand this notion is in terms of a *loss of naturalness*. More precisely, I argue that a conceptual epistemic loss occurs when, as a result of revision, a concept tracks a *less natural* kind than the pre-revised concept did. In this way I argue that epistemic losses do not systematically prevent amelioration, but “merely” generate reasons against it.

Notice that the gradual account of naturalness I put forward in the first paper plays a crucial role in this normative framework, as it renders epistemic losses themselves a matter of degree. Correspondingly, the strength of the normative reasons against revision generated by epistemic losses can vary significantly depending on the magnitude of the

epistemic loss. In other words, the smaller the epistemic loss involved in the conceptual revision, the easier it is for it to be compensated by potential moral or pragmatic gains. Crucially, I take it that introducing this nuance in relation to epistemic losses constitutes an important move, as it renders epistemic losses—or some of them, at least—more *tolerable*. Indeed, notice that one might be more reluctant to accept conceptual revisions—no matter the potential gain—if they involve a dramatic transition from a natural-kind-tracking concept to a non-natural-kind-tracking concept. In contrast, one might be more willing to accept a revision from a natural-kind-tracking concept to a slightly less natural-kind-tracking concept, if the potential gains are worth it. As I show in more detail in this second paper, this normative framework accommodates our classificatory practices significantly better than Simion’s more restrictive alternative.

Finally, in my third paper, “Is Sex Binary? Yes and No, and a Terminological Proposal”, I address a specific case study involving sex categories. More to the point, in this paper I survey the current dispute on the binary or non-binary nature of biological sex and argue that, contrary to appearances, this dispute does not involve any disagreement concerning relevant *facts* about human sexual biology. Exposing this absence of factual disagreement, however, does not dissolve this dispute. Quite to the contrary, I show that there remains a significant unsettled issue concerning the use of the relevant terms and concepts. In this spirit, a part of this paper also involves a more specific normative proposal regarding the use of the term ‘binary’ in characterizing sex. I argue that, on at least one plausible interpretation, ‘binary’ conveys an epistemically and morally problematic idea about human sex. I suggest, in light of this, to avoid its usage and to characterize sex as ‘dimorphic’ instead. The strongest reason in favor of this terminological move, I argue, is that ‘dimorphic’ preserves the important function that ‘binary’ presumably fulfils, while avoiding the problems of its counterpart.

Radical Projectibilism: A Gradual Account of Natural Kindness

1 Introduction

The goal of this paper is to engage in the natural kind debate, and to put forward an original proposal according to which the naturalness of kinds is to be identified with their degree of projectibility.

This view, which I label *Radical Projectibilism*, is congenial to a tradition of natural kind theories that has ascribed a central role to projectibility in the characterization of natural kinds. Radical Projectibilism, however, departs from other views in singling out no condition for naturalness other than projectibility itself. As such, where other theories have often taken projectibility to be necessary yet insufficient for naturalness, I propose, instead, to identify naturalness with projectibility alone. This move follows a similar proposal by Sören Häggqvist (2005), although it goes further, I contend, by addressing important objections, as well as by emphasizing two important dimensions of projectibility that are not considered by Häggqvist: *graduality* and *abundance*.

Identifying naturalness with a gradual property such as projectibility, I argue, constitutes a significant departure from a tradition of natural kinds that has focused on drawing a demarcatory line between natural and non-natural kinds. Far from being a shortcoming of Radical Projectibilism, I will show that understanding naturalness in a gradual way is the most appropriate way to counter the relevant notion of arbitrariness and, moreover, brings significant advantages over dichotomic approaches.

The paper is structured as follows. In section 2, I identify two desiderata that have constrained natural kind theories and which underpin Radical Projectibilism too. I will call these desiderata the *contrast desideratum* and the *science constraint*, respectively. The first of these states that a natural kind theory ought to explain the intuitive contrast between blatantly arbitrary categories (e.g. *discovered on a Tuesday*) and those that seem to, following the classic metaphor, carve nature at its joints (e.g. *water*). The second desideratum states that a natural kind theory ought not to exclude scientifically legitimate categories. Having introduced these desiderata, in section 3 I give an overview of some of the most important natural kind theories and highlight that, while these theories have

generally succeeded in meeting the contrast desideratum, all of them have, in some way or another, violated the science constraint. Then, in section 4, I introduce Radical Projectibilism and focus on its two most distinctive features: *graduality* and *abundance*. I argue that the abundance of projectibility constitutes an advantage of Radical Projectibilism *vis à vis* alternative accounts of naturalness, insofar as it makes the theory extremely inclusive and thus, unlikely to violate the science constraint. Interestingly, though, the abundance of projectibility, which is so useful for meeting one of the desiderata, is the source of an important challenge for Radical Projectibilism. For the abundance of projectibility would seem to prevent Radical Projectibilism from meeting the contrast desideratum, as most categories can be said to be at least slightly projectible. I introduce this challenge in section 5, where I argue, not only that Radical Projectibilism meets the contrast desideratum, but more significantly that, by identifying naturalness with a gradual property, Radical Projectibilism meets this desideratum in a more appropriate way than its dichotomic rivals do. I then turn, in section 6, to presenting an additional advantage of Radical Projectibilism by showing how understanding naturalness gradually can be particularly well suited to accounting for scientific taxonomic progress. In section 7, I defend Radical Projectibilism from views that consider projectibility to be unnecessary for naturalness. More precisely, I discuss two counterexamples from Magnus (2012) and Spencer (2015) respectively, who argue that some scientifically legitimate categories are projectibly weak. In section 8, I present my conclusions.

2 Two Desiderata for a Natural Kind Theory

One, if not *the* central motivation of natural kind theory is to explain the intuitive contrast that exists between categories that seem clearly arbitrary³ and those that, following the classic metaphor, carve nature at its joints. Indeed, some groupings seem to correspond to specific anthropocentric concerns or perspectives (e.g. *pet*), while others have often

³ I take the notion of “arbitrariness” to be the best candidate for appropriately contrasting with the philosophically relevant notion of naturalness. Other potential alternatives such as *social* or *artificial*, in contrast, do not seem to be apt. Indeed, the fact that certain entities are the result of human activities does not seem to mark a significant difference. What the notion of naturalness is supposed to capture, instead, is the fact that certain groupings seem to reflect objective differences in the world (social or otherwise). This, in turn, contrasts with those groupings that are the result of anthropocentric interests, or which are simply random collections.

been assumed to correspond to kinds that pre-exist our classificatory practices or, at least, that are constrained by the way the world is, rather than by our particular interests. Trying for now not to make any strong commitment, we can identify the first desideratum that a satisfactory account of natural kinds should meet. Let us call this the *contrast desideratum*.

Contrast desideratum: A natural kind theory should explain the intuitive contrast between kinds such as *discovered-on-a-Tuesday* and *four-legged animal*, and kinds such as *water*, *tiger*, and *electron*.

In trying to account for this contrast, natural kind theories have often taken *projectibility* to play a central role. Although this notion will be further fleshed out below, the basic and common idea is that alleged natural categories seem to be particularly projectible, meaning that they exhibit a distinctive capacity to support many inductive generalizations (Mill 1843/1974). A paradigmatic natural kind such as *tiger*, for instance, can figure in numerous generalizations regarding its behavior, morphology, lineage, etc. As such, upon observing a member of this kind we will be able to *project* onto it many as yet unobserved properties. We will be able to predict, for instance, that it will likely engage in predatory behavior, that it can run as fast as 65 km/h, or that it is a carnivore. Similarly, projections can also be made in the other direction. That is, from particulars to the kind. When zoologists, for instance, observe for the first time a morphological feature or behavior of a not very well-known species, they will often rightly assume that their discovery is not restricted to the observed organism but can, instead, be *projected* to all the members of its kind.

In contrast, some categories do not seem to have that sort of inductive power. For instance, there is little we can predict or project by knowing that a particular is a member of the kind *discovered-on-a-Tuesday*: there do not seem to be many things unifying the members of this kind, beyond their membership of the kind itself.

On the face of it, it seems clear that some sort of *contrast* needs to be articulated and that the notion of projectibility can be a good starting point. As said above, this idea is far from original, as many have considered projectibility to be central to the characterization of natural kinds. To see this, consider some relevant quotes from Richard Boyd, P.D. Magnus, Muhammad Khalidi, and, more recently, Anjan Chakravartty:

It is a truism that the philosophical theory of natural kinds is about how classificatory schemes come to contribute to the epistemic reliability of inductive and explanatory practices. Quine was right in “Natural Kinds” (1969) that the theory of natural kinds is about how schemes of classification contribute to the formulation and identification of projectible hypotheses (in the sense of Goodman 1973). The naturalness of natural kinds consists in their aptness for *induction* and *explanation*. (Boyd 1999: 146)

A central assumption about natural kinds—the canonical assumption—is that you can make inductive inferences about them. So natural kind terms are the projectible predicates required for successful inductive inference. (Magnus 2012: 10)

Projectability is the most widely agreed upon characteristic of natural kinds, and may in fact be the very reason for positing natural kinds in the first place. (Khalidi 2013: 18)

I have belabored this point about the central role of inductive success in theorizing about kinds because it is crucial to what follows, but the idea that there is some such connection has not generally been controversial. Exploring conceptual linkages between kinds and induction is a theme that appears from antiquity to the present. (Chakravartty forthcoming: 5)

While projectibility, as can be seen, has played a central role in the discussion of natural kinds, it has often been considered insufficient for characterizing naturalness. Indeed, most natural kind accounts do not identify naturalness simply with projectibility, but instead impose further conditions that kinds need to fulfil in order to count as natural. This attitude, often implicitly assumed, is explicitly endorsed by Khalidi, who says: “I only claim that projectibility is a necessary condition on natural kind terms, not that it is a sufficient condition” (2018: 1381).

One of the reasons why natural kind theorists have considered projectibility to be insufficient for characterizing natural kinds is, I contend, the fear of being *overly inclusive*. For projectibility is arguably *abundant*, in the sense that most of the categories we employ, both within and outside of scientific discourse, exhibit a certain degree of projectibility. If you are not convinced about this abundance, notice that basic categories from ordinary language (e.g. *stone*), and even apparently arbitrary categories (e.g. *things heavier than my head*), allow for certain projections, useless as they might be.

As such, identifying natural kinds with projectible kinds could be considered to violate the contrast desideratum, as categories on both sides of it are at least minimally projectible. Magnus voices this concern when he suggests that one problem with identifying natural kinds with those kinds that support inductive inferences is that non-natural kinds such as *jade* also support many inductive generalizations. He says:

So it is typical to say that **jade** is not a natural kind. The problem is that there *are* general facts about jade. Both varieties are fairly hard minerals, which makes them *inedible* and *suitable for making stone tools*. These and many other predicates are projectible for jade *simpliciter*. (Magnus 2012: 12)

Similarly, Alexander Bird (2009) suggests that natural kinds cannot be characterized only by appealing to their inductive power; for many arguably non-natural kinds permit inductive inferences. He says:

Mill and others have allied kinds to induction—natural kind classifications are those that permit induction. But that again does not, as it stands, distinguish natural kinds from the sharing of some natural property—we can make inductions concerning positively charged objects, objects 1kg in mass, and white objects, respectively. (Bird 2009: 502)

Although I will ultimately defend a projectibility-based approach, there is a sense in which this “over-inclusiveness fear” is well-founded. Indeed, as I will argue below, the abundance of projectibility is the source of an important challenge for a projectibility-based approach to natural kinds, as it is not immediately clear how such an account would meet the contrast desideratum.

Before confronting the challenge for projectibility-based accounts, though, it is important that we discuss another element that has played a key role in the development of natural kind theories. A common assumption throughout the discussion of natural kinds has been that scientific inquiry is particularly well-suited to carve nature at its joints and that, in this sense, scientific categories are particularly good candidates for natural kinds. Although the more precise nature of the relation between natural kinds and scientific categories can take different forms, most authors have, implicitly or explicitly, endorsed this assumption.⁴ Laura Franklin-Hall, for instance, says:

⁴ Brian Ellis (2001) can be considered an exception to this attitude, as he is willing to concede that biological categories are not natural kinds. He says: “If evolution occurs in the gradual way that Darwin supposed, or

Most contemporary accounts of natural kinds presume that we learn about the identity of the natural kinds [...] by inspecting the categories and classifications in use in mature sciences. For this policy to be truth-conducive the scientific categories and the natural kinds must to some degree line up with one another. (Franklin-Hall 2015: 932)

In a similar spirit, Khalidi suggests:

Natural kinds correspond to those categories that enable us to gain knowledge about reality. Since science is the enterprise dedicated to acquiring knowledge about the world, natural kinds are identified by the various branches of science. (Khalidi 2013: xi-xii)

As a final illustration, consider also Marc Ereshefsky and Thomas Reydon, who claim:

Any adequate account of natural kinds should capture the kinds offered by our best scientific theories. [...] We assume that the kinds of science have been and are epistemically superior, on the average, to those posited by ordinary language or intuition (except in cases where ordinary kinds are found to be scientific ones). (Ereshefsky and Reydon 2015: 972-973)

This attitude towards scientific categories is widespread among natural kind theorists. On the face of it, we can articulate the second desideratum for a natural kind theory as follows. Let us call it the *science constraint*:

Science constraint: an account of natural kinds should not exclude legitimate scientific categories.⁵

As we shall see, the science constraint has played a decisive role throughout the development of natural kind theories. For, in attempting to articulate the contrast desideratum, most natural kind theories have been accused of violating this constraint in some way or another. That being so, the science constraint desideratum is responsible for

if small changes in genetic constitution can be brought about artificially, then the distinctions between adjacent species—living, dead or yet to be created—must ultimately be arbitrary” (Ellis 2001: 169).

⁵ Notice that a more radical version of this constraint might have it that scientific categories—at least in the ultimate stage of inquiry—*perfectly* correspond to natural kinds and, as such, that an appropriate natural kind theory should not only include all legitimate scientific categories, but also exclude non-scientific categories (e.g. folk categories). The alternative presented in this paper is more permissive and, in this sense, only requires natural kind theories *not to exclude* scientific categories, while allowing that some non-scientific categories might count as natural.

a significant tendency towards inclusiveness, which has characterized the development of natural kind theories. The following section considers some of the most important proposals and focuses on their difficulties in respecting the science constraint.

3 Failures to Preserve the Science Constraint: Towards Inclusiveness

3.1 Natural Kind Essentialism

The most significant, and likely the most discussed violation of the science constraint comes from Natural Kind Essentialism (NKE hereafter). Given NKE has been widely discussed in the literature and that it has become a minority position (see Ellis 2001, 2008; Wilkerson 1988), I won't delve far into it here. It is important for our purposes, however, to emphasize that it is precisely the violation of the science constraint that has made NKE a marginal view among philosophers. Chakravartty puts this idea as follows:

The most obvious and compelling sources of resistance to an exclusive commitment to kinds with essences are the sciences themselves. The kinds of objects investigated by the sciences are sometimes describable in terms of essences, but often resist this sort of description. The traditional view that kinds are ontologically distinguished by essences has a storied past, but many of the kinds one theorizes about and experiments on today simply do not have any such things. (Chakravartty 2007: 157)

As has often been pointed out, the most notorious failure of NKE comes from its incapacity to accommodate biological categories. Indeed, the standard view among philosophers of biology is that biological categories do not fit in a strict essentialist framework, as there is no single genotypic or phenotypic property that would serve to individuate species (Dupré 1981, 1993; Khalidi 2013; Magnus 2012; Kitcher 1984). Ereshefsky expresses this idea vividly. He says:

Biologists have been hard-pressed to find a biological trait that occurs in all and only the members of a particular species. Evolutionary theory explains why. Suppose a genetically based trait were found in all the members of a species. Mutation, random drift, and recombination can cause that trait not to occur in a future member of that species. All it takes is the disappearance of a trait in a single member of a species to

show that it is not essential. The universality of a trait among the members of a species is quite fragile. (Ereshefsky 2007: 405)

In a similar vein, Samir Okasha claims that “the idea that species can somehow be ‘defined in terms of their DNA’ has no basis in biological fact” (2002: 197). Given the status of biological species as paradigmatic natural kinds and legitimate scientific categories, essentialism’s failure to accommodate them is likely the strongest instance of a violation of the science constraint that we can think of.⁶

3.2 *Homeostatic Property Clusters*

The limitations of NKE in the biological domain constitute the main motivation for Boyd’s (1991, 1999) account of natural kinds as Homeostatic Property Clusters (HPC).

According to the HPC view, natural kinds are clusters of properties whose stable co-occurrence is maintained by homeostatic mechanisms. That is, mechanisms responsible for preserving the properties of the cluster in a state of equilibrium. HPC theory thus departs from NKE in dropping many of its most controversial requirements, and by explaining the inductive potential of natural categories without positing essences. With this in mind, Boyd’s proposal can be read as an attempt to provide a more flexible and inclusive framework that is able to accommodate biological categories, and thus able to preserve the science constraint.

This flexibility makes HPC more inclusive than its essentialist predecessor, and arguably a better alternative for accommodating biological categories. Indeed, not only do species lack any good candidate to play the role of an essence, but also, the properties that biological kinds share are often the result of various mechanisms involving environmental pressures, interbreeding, developmental processes, and genetic descent, among other factors. HPC theory thus provides a compelling alternative account of the non-accidental co-occurrence of properties that ground our inductive practices involving biological categories.

Just like its predecessor, however, HPC has been accused of violating the science constraint and excluding legitimate scientific categories. Several authors (Ereshefsky & Reydon 2015; Khalidi 2013; Slater 2015) have argued in this direction and have

⁶ See Devitt (2008) for a divergent minority view, according to which biological species do fit the essentialist picture.

suggested, with varying degrees of emphasis, that not all natural kinds constitute homeostatic property clusters. Some of these critics concede that HPC accommodates biological kinds, while failing to accommodate other kinds whose equilibrium does not seem to be sustained by homeostatic mechanisms (e.g. chemical elements, fundamental physical particles); but some others go as far as to insist that HPC does not even fit all biological species (Ereshefsky & Matthen 2005).

Be that as it may, most theorists agree on the idea that, in some way or another, HPC is still too restrictive, as it cannot accommodate the vast heterogeneity of scientifically legitimate categories.

Wary of the difficulties of providing a general theory of natural kinds that is able to encompass this heterogeneity, a recent trend in natural kind theory focuses on the epistemic utility characteristic of natural kinds and avoids making any metaphysical commitment as to what grounds this epistemic utility. More precisely, these views attempt to characterize the clustering of properties while remaining neutral about any specific metaphysical grounding for it. Let us consider these views, which, following Conix (2017), we may refer to as Bare Property Cluster accounts of natural kinds.

3.3 *Bare Property Clusters*

Bare Property Cluster (BPC) accounts of natural kinds constitute a significant departure from previous approaches to natural kinds insofar as they focus on the robust clustering of properties in virtue of which inductive inferences are reliable, without committing to any specific account of this clustering.

This departure is motivated by past failures on the part of previous natural kind theories, which, as we have seen, always seem to violate the science constraint in some way or another. Indeed, BPC defenders believe that no general grounding claim will be able to account for all natural kinds, and thus that the only way for a notion of natural kinds to be appropriately inclusive is for it to remain neutral regarding the metaphysical grounding of this robust clustering.

Matthew Slater (2015), for instance, claims that natural kind theories have focused too much on the “grounding claim” and should instead turn their attention to the epistemic usefulness of categories, without committing to any specific metaphysical story about essences or homeostatic mechanisms. He says:

I disagree with the claim that the epistemic value of natural kinds is contingent on the existence of some concrete ground—some essence, mechanism, or feature of the causal structure of the world—that Kornblith believes ‘binds together the properties which we use to identify kinds’. Call this the grounding claim; its broad acceptance seems a likely explanation for the emphasis on the metaphysics of natural kinds to the comparative neglect of the precise epistemic role they play. (Slater 2015: 384)

Slater argues directly against the HPC view and puts forwards an original proposal (*Stable Property Clusters*) that attempts to articulate more systematically how to understand the *stability* in virtue of which clusters of properties can support inductive generalizations and inferences.

To convey the relevant notion of *stability*, Slater presents the picture of a clique of friends with three members: Peg, Ralph, and Quinn. These individuals form a *stable* clique and like hanging out together. As such, spotting any of these three in the mall is generally a good indicator that the others will be there as well. This is, very roughly, the sense of stability that Slater wants to capture; the instantiation of a property of the cluster reliably indicates the presence of the other co-occurring properties of the cluster.

Similarly, Chakravartty (2007) suggests the metaphor of “sociability” to refer to all the ways in which properties enter into systematic relations and thus ground our inductive practices. Crucially, the metaphor of sociability is intended to be neutral about, yet compatible with, more specific grounding accounts of these systematic patterns of sociability. Chakravartty says:

Properties, or property instances, are not the sorts of things that come randomly distributed across space-time. They are systematically “sociable” in various ways. They “like” each other’s company. The highest degree of sociability is evidenced by essence kinds, where specific sets of properties are always found together. In other cases, lesser degrees of sociability are evidenced by the somewhat looser associations that make up cluster kinds. In either case, it is the fact that members of kinds share properties, to whatever degree, that underwrites the inductive generalizations and predictions to which these categories lend themselves. This is a reflection of the striking, poetic fact that some collections of property instances like each other’s company and others do not. It is this fact that one captures with talk about natural kinds, and this feature of reality surely has a place in the ontology of scientific realism. (Chakravartty 2007: 189)

Interestingly, despite their attempts at inclusivity, it could be argued that even certain BPC accounts end up being too restrictive and violate the science constraint. In this line of thought, Manolo Martínez (2020) has suggested that Slater's SPC account could fail to include what he calls "synergic kinds". Let us flesh this out.

3.4 *Beyond Bare Property Clusters*

Martínez argues that some kinds (i.e. synergic kinds) ground inductive inference not, as in the case of co-occurring property clusters, because the instantiation of a property of the cluster is indicative of the instantiation of other (co-occurrent) properties of the cluster, but instead because "the joint instantiation of all or many of those properties [...] plays the explanatory role for which the natural kind is recruited" (2020: 1935). To illustrate the point and convey more vividly what is different about synergic kinds, Martínez makes use of Slater's clique example, as presented above, while incorporating some modifications.

Martínez suggests that we think of Peg, Ralph and Quinn not as a clique of friends that like each other's company, but instead as a rather tense love triangle. In this case, spotting only one of the three at the mall is not indicative of the presence of either of the other two, while spotting two of them together is a reliable indicator that the third one is *not* going to be there. The idea that this picturesque metaphor is meant to convey is that, when it comes to synergic clusters, the instantiation of properties, individually, is not a reliable indicator of the instantiation of other properties of the cluster. Instead, it is the joint instantiation of properties that allows for reliable inferences. Martínez makes clear that this discussion is not otiose, as some scientific categories and inferences seem to have this synergic structure. More precisely, Martínez (2020: 1943-1944) suggests that this could be the case with some categories involving *epistatic* interactions (e.g. *fruit-fly wing*) and categories from *brain connectomics* (e.g. *human brain*). Martínez also suggests that invoking synergic kinds can be helpful for accommodating the much-discussed case of *polymorphic* species (2020: 1941-1942).

What this discussion reveals, I argue, is that even some BPC views such as Slater's, despite their attempted inclusiveness, seem susceptible to violating the science constraint. On the face of these successive failures, a more promising alternative, I suggest, is to focus exclusively on the inductive power of categories; that is, on

projectibility. Indeed, if a recurrent problem of natural kind theories is that they fail to be appropriately inclusive, identifying projectibility with naturalness appears to be a good solution. For not only is projectibility, as mentioned above, abundant among categories; it is also neutral with regard to specific metaphysical grounding claims. A projectibility-based account, then, will stand out from the rest because of its inclusivity and, as such, will have no trouble in subsuming both Slater's Stable Property Clusters and Martínez's synergic kinds (along with kinds with essences, HPCs, etc.). As we will see, however, this inclusivity will be the source of a potential problem for this approach that will need to be properly dealt with. In the following section I present the approach. Then I introduce the challenge.

4 Radical Projectibilism

Let me use the label "Radical Projectibilism" to refer to the view that identifies the naturalness of kinds with their projectibility, that is, with their capacity or disposition to enter into predictions and generalizations (independently of the metaphysical grounding that this inductive power may have).

As we have seen, most natural kind accounts, while taking projectibility to be *necessary*, rarely deem it to be *sufficient* for naturalness. A notable exception to this tendency, however, is provided by Sören Häggqvist (2005), who has argued for a projectibility-based approach to natural kinds.⁷ Häggqvist also contrasts his view with HPC. He says:

The second objection to the causal homeostasis view, as elaborated by Boyd and Griffiths, is more important. It is that the demand for underlying mechanisms, even short of demanding internal micro-mechanisms, is still excessive. It is not at all clear why the lack of such mechanisms should impair the soundness of a kind. If there is an underlying mechanism explaining the projectibility of the kind, this may

⁷ Although Häggqvist's view is sometimes included among Bare Property Cluster accounts (see Lemeire 2021; Conix 2017), it is important to highlight an important difference that might set it apart from these views. For, although Häggqvist's view fits in among BPC accounts regarding its neutrality *vis-à-vis* any specific metaphysical grounding for the robust clustering of properties, it departs from these views in incorporating the possibility of certain robust clusters being *brute*. That is, having *no ground*. More precisely, Häggqvist (2005: 82) claims that there is no principled reason to assume that there will always be a causal explanation (be it in terms of essences or more loose causal mechanisms) to account for the clustering of properties. Some of these robust clusters, he claims, might simply be a brute matter of fact. He suggests that this could be the case with certain kinds from fundamental physics, when there does not seem to be any causal explanation for the perfect clustering of properties (2005: 81).

guarantee its unity and reality. But I would suggest that what matters is the reality of the kind, or, to put it another way, its genuine, as opposed to apparent, projectibility. Micro-essences, underlying mechanisms and other essences may be sufficient for the reality of a kind. Are they also necessary conditions for it? (Häggqvist 2005: 80)

My proposal, then, follows the path opened by Häggqvist, but takes two significant steps further, as follows. First, I elaborate the account in response to a serious challenge that is overlooked by Häggqvist. Then I develop some implications that follow from identifying naturalness with projectibility, and which make the proposal depart radically (hence the name) from most traditional approaches to natural kinds. Let us consider each of these ideas in turn.

Häggqvist rightly assumes that a significant benefit of a projectibly-based account of natural kinds is its inclusiveness. As mentioned above, the abundance and neutrality of projectibility makes it difficult for it to exclude any potential natural kind candidates. Häggqvist, however, does not seem to notice that this abundance is a double-edged sword, as it might make the account *overly* inclusive. More precisely, identifying naturalness with projectibility threatens to violate the contrast desideratum, insofar as categories on both sides of the contrast seem to be, at least, minimally projectible. The challenge for a projectibility-based approach to natural kinds, then, is not to preserve the science constraint—which seems easily satisfied—but to meet the contrast desideratum.⁸

The other aspect that differentiates Radical Projectibilism from Häggqvist's proposal is its emphasis on an aspect of projectibility that Häggqvist does not consider: *graduality*. For, crucially, projectibility is not an on-off feature of kinds, but rather, a *gradual* property that can be instantiated in varying degrees. Although not often fully exploited, the idea that projectibility is gradual is not an original one (see Dorr 2019: 42; Griffiths 1999; Khalidi 2018; Magnus 2012: 12; Millikan 2000). Furthermore, it is often acknowledged that kinds can be projectible along two different gradual dimensions (Griffiths 1999: 217; Khalidi 2018: 1383; Millikan 2000: 26): on the one hand, the projections or generalizations in which a kind enters can be more or less *robust*. On the other hand, kinds can be more or less projectible depending on the number or *variety* of generalizations they allow for. Let us flesh this out.

Following Khalidi (2018), we can roughly characterize the *robustness* of a generalization by the number of exceptions it has. While some generalizations are

⁸ I address objections challenging the *necessity* of projectibility in section 7.

universally true and hold under all circumstances, others, although not universal, hold under an exceptionally large range of circumstances, while yet others hold only under rather specific circumstances and require significant *ceteris paribus* clauses (Khalidi 2018: 1382).⁹

The *variety* dimension, instead, corresponds to the number of generalizations in which kinds can enter. Although it is generally expected that paradigmatic natural kinds can figure in numerous generalizations—Mill went as far as to hold that “Real Kinds” could enter into *infinitely many* generalizations (1843/1974, I vii §4)—Khalidi suggests that some paradigmatic natural kinds might actually figure in very few (e.g. *electron*). Khalidi quickly adds, though, that the latter’s poor performance in the variety dimension is compensated by the great (or even universal) *robustness* of the generalizations into which they enter. As we shall see in section 7.1, distinguishing these two dimensions of projectibility will be useful for defending the strong projectibility of certain scientific categories against accusations to the contrary (Spencer 2015; Magnus 2012).

Now, identifying naturalness with a seemingly abundant and gradual property constitutes a significant departure from traditional natural kind theories that have generally focused on drawing a demarcatory line between natural and non-natural kinds. Radical Projectibilism, instead, takes naturalness to be a gradual property and, as such, presents a novel framework where the relevant question is not whether a kind is natural or not (given that most kinds, as we have seen, are at least minimally natural), but instead, its *degree of naturalness*.

As I will argue next, it is precisely the emphasis on the graduality of projectibility, and hence the graduality of naturalness, that will help Radical Projectibilism to address the challenge introduced above and meet the contrast desideratum.

5 The Challenge of Projectibility-Based Accounts of Naturalness: Meeting the Contrast Desideratum

According to the reconstruction provided above, two main theoretical constraints have driven natural kind research. On the one hand, I have emphasized that the main goal or desideratum of natural kind theories has been to articulate an intuitive contrast between

⁹ See Woodward (2000) for a more detailed discussion and characterization of non-accidental generalizations.

arbitrary categories and those that, following the traditional metaphor, carve nature at its joints (i.e. the contrast desideratum). On the other hand, I have shown how the main attempts to account for this contrast have successively violated the science constraint by excluding scientifically legitimate categories. We have also seen that even arguably very inclusive accounts such as Slater's Stable Property Cluster theory might exhibit this problem.

On the face of it, one sensible alternative, I suggested, is to follow Häggqvist and focus exclusively on the presumably abundant and neutral notion of projectibility. As discussed above, however, the main benefit of projectibility can also constitute a potential weakness, since it is not immediately obvious, given this abundance, how a projectibility-based account would meet the contrast desideratum.

To address this challenge, I suggest that we focus on the *graduality* of projectibility. For, although the *abundance* of projectibility threatens to blur the contrast, its graduality allows us to highlight that *not every kind is projectible to the same degree* and hence, not natural to the same degree. According to Radical Projectibilism, then, the intuitive contrast between categories such as *discovered-on-a-Tuesday* and *water*, is just the contrast between the two extremes of a spectrum. Notice that by identifying naturalness with a gradual property, Radical Projectibilism departs from the tradition of drawing a sharp demarcatory line between natural and non-natural kinds. In what follows, I argue that, far from being a shortcoming of this view, understanding naturalness as a gradual property is the appropriate way to counter the relevant notion of arbitrariness.

To elaborate, notice that if we want our notion of naturalness to stand in appropriate contrast with the notion of arbitrariness, we need a characterization that captures nuanced differences and not only extreme ones. Consider for instance, the kind *pet*.¹⁰ Although this kind seems *more arbitrary* than the kind *tiger*, it does seem *less arbitrary* than the kind *animals-belonging-to-the-emperor*.¹¹ Similarly, although everything seems to suggest that the kind *tiger* is not arbitrary, we also have reasons to think that it is *more arbitrary* than the kind *gold*. The more examples we consider, the clearer it will be that it does not seem possible to separate all kinds into two perfectly discrete boxes, as the contrast desideratum would have us believe. What this suggests,

¹⁰ Interestingly, Khalidi (2018) uses this kind as an example of a paradigmatic *non-natural* kind.

¹¹ See Borges's (1942) essay "The Analytical Language of John Wilkins" and the curious taxonomy of animals suggested there.

instead, is that the difference considered in the contrast desideratum is but one particular *extreme* instance of a more general and ubiquitous relation: *more (or less) natural than*.

As soon as we appreciate this, we can see that Radical Projectibilism is in a better position than alternative dichotomic accounts to articulate this more general relation. For dichotomic accounts, insofar as they posit a single sharp demarcatory line, are only able to capture the particular extreme case, and not the more intermediate ones. In this sense, they cannot account for the more general relation *more natural than* of which the contrast desideratum is but one (extreme) instance.

That being so, a projectibility-based account which characterizes naturalness in terms of a gradual property seems particularly well suited to counter the relevant notion of arbitrariness satisfactorily, and to accommodate both extreme and nuanced contrasts. This is the sense in which I contend that Radical Projectibilism not only meets the contrast desideratum, but does so in a more appropriate way, as it also accommodates the more general cases that dichotomic accounts do not accommodate.

As an additional illustration of the potential limitations of dichotomic accounts of naturalness, consider the notorious revision of the concept FISH which, roughly, went from tracking the kind *aquatic animal*—which included whales—to tracking the kind *cold-blooded vertebrate with gills*. Let us call the former *fish_{AQUATIC}* and the latter *fish_{GILLS}*. Although it is uncontroversial that the current English term ‘fish’ refers to *fish_{GILLS}*, philosophers disagree on the “natural kind” status of these two kinds. According to John Dupré’s (1993) Promiscuous Realism, insofar as both kinds stress important sameness relations and serve legitimate purposes, they should both count as natural (1981: 92). On Khalidi’s more restrictive view, in contrast, only the alleged scientific category *fish_{GILLS}* counts as a natural kind (2013: 62).¹²

I want to use this case to illustrate that, independently of whether or not one counts *fish_{AQUATIC}* as a natural kind, a dichotomic approach will face some significant limitations and will lead to some counterintuitive results. As such, I contend that the problem of these views does not stem from *where* they draw the natural/non-natural demarcatory line but, rather, from drawing such a line at all. Let us consider this case in more detail.

Dupré’s Promiscuous Realism tells us that, provided we do not associate the notion of a natural kind with essentialist views, we have good reasons to think of *fish_{AQUATIC}* and *fish_{GILLS}* as equally natural. More particularly, Dupré believes that

¹² Dupré (1999) casts serious doubt on the status of *fish_{GILLS}* as a scientific category.

scientific categories are not fundamentally different from folk ones (1999: 462) and, as such, does not see any reason to dismiss the folk category *fish*_{AQUATIC} as non-natural. Although Dupré is certainly right to emphasize the utility of this kind and the fact that there does not seem to be any fundamental difference between *fish*_{AQUATIC} and *fish*_{GILLS}, his account does not tell us anything about the intuitively plausible *improvement* that has occurred in the conceptual transition from *fish*_{AQUATIC} to *fish*_{GILLS} (more on the notion of conceptual progress in the following section). While I agree with Dupré in not thinking that there is any fundamental or metaphysical difference between these two kinds, I believe, however, that it makes sense to think of *fish*_{GILLS} as being *more natural* than *fish*_{AQUATIC}. For one thing, the kind *fish*_{GILLS} groups particulars in a way that seems to allow for more interesting generalizations than the kind *fish*_{AQUATIC} and, additionally, seems to provide a deeper understanding of the aspect of reality it represents. Notice that Khalidi also emphasizes this apparent contrast, and after insisting that the category *fish*_{AQUATIC} has no inductive value (2013: 62), he suggests that the category *fish*_{GILLS}, in turn, is scientifically useful. He says:

It is instructive to contrast this inclusive use of the term ‘fish’ [*fish*_{AQUATIC}] with the ‘scientific one’ [*fish*_{GILLS}]. [...] Despite the fact that it is not a unitary taxon from the evolutionary or phylogenetic point of view, the category *fish* [*fish*_{GILLS}] has undisputed value as an epistemic kind. There are a number of branches of science, such as ichthyology and marine biology, which use this category to explain and predict natural phenomena. (Khalidi 2013: 62-63)

Although I will ultimately suggest that Khalidi goes too far in positing a fundamental difference between these two categories, a permissive dichotomic account such as Dupré’s, which locates both *fish*_{AQUATIC} and *fish*_{GILLS} on the “natural side” of the divide, is not satisfactory either, as it is unable to articulate this intuitive difference in terms of naturalness. To be clear, my contention against Dupré’s account does not target its promiscuity or permissiveness. I am actually very sympathetic to this attitude. My complaint is, rather, that we need to complement this permissive picture with a gradual account in order to emphasize significant differences that will otherwise remain overlooked. Let us turn to consider the other side of the picture: Khalidi’s more restrictive approach to the case.

Khalidi believes that “not all purposes are created equal” (2013: 62) and that kinds introduced for epistemic purposes have to be prioritized over those that serve other non-

epistemic or pragmatic purposes. As such, he argues that scientific categories will tend to correspond to natural kinds, whereas folk ones will not. Unsurprisingly, then, Khalidi dismisses *fish_{AQUATIC}* as non-natural, while insisting that *fish_{GILLS}* is a natural kind. The problem, again, is that a single sharp demarcatory line is not enough to capture the nuanced differences that kinds may exhibit. For, while it is plausible to think that there is a contrast between *fish_{AQUATIC}* and *fish_{GILLS}* in terms of naturalness (as Khalidi duly emphasizes), it is equally plausible to think that a similar contrast arises when we compare the allegedly non-natural *fish_{AQUATIC}* with a random category such as *wet creature*; a contrast which, I contend, an account of naturalness ought to capture. Khalidi, however, is unable to account for such differences in terms of naturalness given that, on his view, both *fish_{AQUATIC}* and *wet creature* are equally non-natural.

The limitation of having only two discrete options (either natural or non-natural) also explains why Khalidi seems forced to overstress the difference between *fish_{AQUATIC}* and *fish_{GILLS}*, and refer to the former as if it were inductively worse than it actually is. He says: “When the category fish includes aquatic animals, such as crayfish, jellyfish, starfish, and some mollusks, as well as whales and dolphins, it ceases to have value as an inductive category” (Khalidi 2013: 62). I believe, however, that Khalidi is too quick in making this assessment. Not only, as we noted above, is it the case that the abundance of projectibility guarantees that most categories, including *fish_{AQUATIC}*, will exhibit some degree of projectibility; but, additionally, in this particular case, many of the inductive projections that *fish_{AQUATIC}* allows for seem particularly useful and relevant for certain purposes.¹³

Hopefully, this discussion has served to illustrate that approaching these cases equipped with only two discrete boxes constitutes a serious limitation for dichotomic accounts of naturalness. A gradual account such as Radical Projectibilism, in contrast, seems better able to accommodate both the extreme contrasts and the more nuanced ones. In the following section, I turn to presenting a related advantage of a gradual account: this time in relation to conceptual progress.

¹³ See Souza and Begossi (2007).

6 Putting the Gradual View to Work: Scientific Conceptual Progress

While philosophical discussion of scientific progress generally focuses on theory rather than on taxonomy (see Ikka 2019), we have good reasons to believe that scientific progress also affects scientific categories themselves. Although this taxonomic progress can take different forms, I take it that the notion of naturalness can be useful for accounting for some of its instances. On this view, then, conceptual change in science constitutes conceptual progress when categories become (more) joint-carving.¹⁴

Assuming that this is a plausible approach, the goal of this section is to argue that a gradual notion of naturalness is better suited to account for this progress than any dichotomic alternative. More precisely, I contend that a gradual account of naturalness is better able to encompass a broader range of putative instances of taxonomic progress and, as such, in a better position to accommodate actual conceptual practices within science. As we shall see, an important problem of dichotomic accounts is that their usefulness is restricted to cases where the conceptual improvement involves a transition from a non-natural to a natural kind. But scientific conceptual progress, I will suggest, rarely fits this demanding dichotomic picture, as it often involves gradual or incremental revisions and modifications, rather than drastic transitions.

As an illustration, consider again the revision of the concept FISH. In the previous section, I discussed this well-known revision which, we may recall, involved a transition from the category *fish_{AQUATIC}* to the category *fish_{GILLS}*. I argued against both Khalidi's and Dupré's views of this revision by showing that drawing a sharp natural/non-natural divide, no matter where, inevitably leads to counterintuitive results (either by committing to the absolute arbitrariness of *fish_{AQUATIC}*, or by obscuring a significant difference in naturalness between these two categories). The problem, I insisted, is that being equipped with only two boxes does not enable us to highlight the more nuanced distinctions that kinds may exhibit in this and many other cases.

Interestingly for our current purposes, the revision of 'fish' is often taken to be a paradigmatic example of conceptual progress. Now, assuming that this is indeed the case, if one expects the notion of naturalness to play any role in accounting for this progress, I contend that a gradual approach will prove more appropriate. To see why, notice that the

¹⁴ Notice that this does not exclude the possibility of conceptual progress occurring by other means that do not involve naturalness in any way.

defender of a dichotomic approach will be burdened with strong and problematic commitments if she insists on reading this case through a “dichotomic lens”. For, since she is limited to assessing categories as either natural or non-natural, the only form of progress (in terms of naturalness) that she can capture will necessarily involve a transition from a non-natural to a natural kind. As discussed in the previous section, however, this sort of approach does not seem very promising for the case under consideration. For, although there does seem to be a significant difference in non-arbitrariness between *fish_{AQUATIC}* and *fish_{GILLS}* (one that we are trying to account for), this difference does not seem to amount to a deep metaphysical contrast between a non-natural and a natural kind. Relatedly, an apparently undesirable consequence of the dichotomic reading of this revision is that it commits us to the view that *fish_{AQUATIC}* is just as non-natural as any other blatantly arbitrary category we might think of (e.g. *wet animal*). As we already saw in the previous section, though, there are strong reasons to doubt that this is the case.

A gradual account such as Radical Projectibilism, in contrast, does not exhibit these problems. More specifically, it can successfully account for this instance of progress in terms of naturalness, without positing any metaphysically deep difference between the two categories, and thus without committing to the absolute arbitrariness of *fish_{AQUATIC}*. That is, it is permissive enough to count both *fish_{AQUATIC}* and *fish_{GILLS}* as natural, while further allowing us to explain the type of progress that seemingly occurs in this transition.

Last but not least, notice that articulating this graduality in terms of projectibility, as I have recommended, seems particularly promising to deal with this case. For, crucially, the increase in naturalness that very plausibly occurs in the transition from *fish_{AQUATIC}* to *fish_{GILLS}* can be captured by a correspondingly plausible increase in projectibility. Indeed, as suggested above, the category *fish_{GILLS}* is significantly more projectible (Khalidi 2013:62) than the less scientifically-informed, *fish_{AQUATIC}*, as it seems to support a higher number of inductive generalizations, relevant for various scientific domains.

Before concluding this section, a clarification may be in order. My aim so far has been to argue that a gradual account of naturalness stands in a better position than a dichotomic one to contribute towards an account of scientific conceptual progress. Now, this claim has to be distinguished from the more radical view according to which *all* instances of scientific conceptual progress are constituted by an improvement in naturalness—that is, the view according to which a gain in naturalness is necessary and sufficient for scientific conceptual progress. Far from endorsing this position, I have only

committed to the idea that an improvement in naturalness is an important way for conceptual progress to come about, while allowing that it is likely not the only one. Very plausibly, other considerations (be they epistemic or not) apart from naturalness will often also drive scientific taxonomy and, as such, will be relevant for assessing scientific conceptual progress. What I have argued, moreover, is compatible with positions according to which other important considerations—say, moral or political—sometimes *override* concerns about the naturalness of categories. The modest idea I have intended to convey here is simply that if we want a notion of naturalness to play at least some role in explaining certain intuitive instances of scientific conceptual progress, then a gradual notion of naturalness will be better suited to this task than a dichotomic one.

7 Necessity Worries

At the beginning of this paper, I introduced projectibility as a feature that has often been taken to be distinctive of natural kinds. The originality of Radical Projectibilism, I argued, lies in identifying naturalness with projectibility and meeting the contrast desideratum by emphasizing its graduality.

I also suggested that one of the reasons why many authors have resisted this identification is a fear of being overly inclusive (recall Magnus's (2012: 12) resistance to counting *jade* as a natural kind). As such, many natural kind theorists have taken projectibility to be necessary but insufficient for naturalness and have thus come up with further conditions for demarcating natural from non-natural kinds.

Interestingly, some authors have voiced concerns with projectibility-based accounts that take the opposite direction, as it has also been argued that projectibility may not, after all, be *necessary* for naturalness. More precisely, some authors have argued that some scientifically legitimate categories are not very projectible and, as such, that a projectibility-based account will fail to be appropriately inclusive. This worry is to be taken seriously; for, if these considerations were right, Radical Projectibilism would, in its own way, also violate the science constraint. I will consider two such arguments. First, I will present Quayshawn Spencer's argument regarding the poor inductive power of superheavy elements. Then, I will turn to considering a similar contention from Magnus involving polymorphic species. My strategy for resisting these potential counterexamples will consist in arguing that neither Spencer nor Magnus succeed in making the case for

the poor projectibility of their respective examples. I will thus argue that both superheavy elements and polymorphic species are significantly projectible categories.

Before getting into the details, however, we must clarify what it takes for a projectibility-based account to violate the science constraint. For one could worry that counterexamples of the kind that Magnus and Spencer have in mind are, given the abundance of projectibility, non-starters. Indeed, once we have conceded that even kinds such as *discovered-on-a-Tuesday* may be minimally projectible, scientific categories such as superheavy elements and polymorphic species will surely be, at least, minimally projectible too. I find this way of addressing these counterexamples unconvincing, however. Avoiding violating the science constraint by means of responding that “every kind is at least minimally projectible”, seems to trivialize the science constraint. Given the well-attested important role that this theoretical constraint has played throughout the development of natural kind theories, this trivialization certainly looks undesirable. Indeed, the science constraint has received so much attention because we have strong intuitions to the effect that the way systematic scientific inquiry groups things together is *particularly* non-arbitrary. As such, I defend that for a projectibility-based account to meaningfully respect the science constraint it is not enough to insist that every (scientific) category is at least minimally projectible, and hence, “a little bit” natural. In order for a projectibility-based account of naturalness to non-trivially preserve the science constraint, scientifically legitimate categories must support a significant number of relevant inductive generalizations. Although this restriction is certainly not very specific, it is enough to make the satisfaction of the science constraint non-trivial, as it establishes a higher standard for scientific categories to reach, and, in doing so, sets them apart from minimally projectible categories such as *discovered-on-a-Tuesday*, *pet*, etc.

More specifically, when it comes to the two counterexamples under consideration, I will insist, *contra* Spencer and Magnus, that they are unproblematic for a projectibility-based account, as they both support a significant number of inductive generalizations, relevant for various scientific domains. As such, I will argue that they (non-trivially) satisfy the science constraint, and thus do not constitute genuine counterexamples to a projectibility-based account of natural kindness.

Having anticipated this, let us consider Spencer and Magnus’s potential counterexamples in more detail.

7.1 Superheavy Element 117

Spencer argues that a natural kind theory that focuses exclusively on the inductive power of kinds (i.e. projectibility) will fail to include certain paradigmatic natural kinds which, he claims, are “notoriously inductively weak” (2016: 162). To substantiate this idea Spencer presents the case of superheavy elements, and focuses in particular on element 117, also known as “tennessine”. Indeed, given the seemingly indisputable status of chemical elements as paradigmatic natural kinds, it would be problematic for any theory of naturalness to exclude such paradigmatic exemplars or, in the case of a gradual account, to ascribe them the same degree of arbitrariness as categories such as *discovered-on-a-Tuesday* and the like.

I will suggest, however, that Spencer does not succeed in making the case for the weak projectibility of element 117. More precisely, I will argue that this element supports relevant inductive generalizations and that Spencer’s incorrect assessment derives from conflating projectibility with other notions in the vicinity, such as our *capacity* to draw inductive inferences, or the inductive *method*. Let’s consider Spencer’s view in more detail. Concerning chemical element 117, he says:

Since only six atoms of element 117 have ever been synthesized, and since the atoms that have been synthesized have existed for less than a second, nuclear chemists have not been able to get “many inductive generalizations” out of 117. Furthermore, the latter is not a temporary setback. Due to the nuclear instability of 117, it is not the sort of kind that we *can* generate many inductive generalizations with it. Thus, unlike other elements, we know nothing about 117’s properties at standard temperature and pressure – such as its phase, its density, its melting point, its boiling point, its ionization energies, or its atomic radius. [...] So, natural kind theories that require natural kinds to be inductively powerful fail to predict the existence of inductively weak paradigm natural kinds, such as superheavy elements. (Spencer 2016: 162)

The first thing to notice is that Spencer’s claim regarding the weak projectibility of *tennessine* should not be understood merely as stating that, given its nuclear instability, we lack the *capacity* to learn as many things about it as we can about other, more stable elements. For this limitation would simply amount to us *knowing* comparatively fewer projections supported by this category, but would not be indicative of the category being

projectibly poor.¹⁵ Rather, Spencer's claim must not only be that we cannot *learn* tennesine's properties but, more radically, that tennesine *lacks* the relevant properties typical of other non-superheavy elements (e.g. melting point, density, etc.) and, as such, that there are few projections we can make about it.

With this clarification in mind, in what follows I put forward various considerations that cast serious doubt on this view. I argue that we have no compelling reasons to believe that *tennessine* (along with other superheavy elements) is significantly less projectible than other chemical elements.

First of all, notice that the intrinsic instability of tennesine already constitutes a very *robust* general fact about this element; one on which the experiments to synthesize it heavily relied. Interestingly, in presenting his Stable Property Cluster approach to natural kinds, Matthew Slater makes a similar point about another unstable element: *uranium-235*. He says:

Rather than focusing on particular instances of kinds—*this* bar of gold, for example—we need to focus instead on the *pattern* of property instantiations. Once we see this point, we can readily draw a contrast between *objects* that are unstable, even though the pattern of properties associated with the *kind* is stable. Trade your gold bar for one made of uranium-235. It's unstable and hence radioactive (careful!). But the properties of the *kind* uranium-235 are stable; the *pattern* of having a certain atomic number, mass, density, half-life, and so on stably come as a package deal. That's why we knew to advise caution. It's a stable, reliable fact about the kind uranium that any particular *chunk* of uranium is unstable. (Slater 2013: 147)

Although I will also contend that tennesine has many other projectible properties, notice that having a very robust one (i.e. instability) is already a good indicator that this category is inductively powerful. For, we may recall, the projectibility of a category depends not only on the *variety* of projections that it supports, but also on the *robustness* of those projections. As Khalidi suggests, the fact that projectibility ranges over two dimensions allows some very projectible categories to be so, not in virtue of supporting *many* inductive generalizations, but instead in virtue of the (few) generalizations they support

¹⁵ To see this through an example, consider the case of *Phobaeticus chani*, a stick insect with outstanding camouflage skills. We know very little about this insect, partly because only a few specimens have been observed to date. It seems clear, though, that it cannot be deduced from this epistemic limitation and our corresponding lack of knowledge about this insect that this category is inductively weaker than any other species category that is more easily observed and studied.

being very robust (or the other way around). This could be the case, Khalidi suggests, with some kinds from fundamental physics, such as *electron*, which although generally characterized only in terms of three properties (spin, charge, and weight), is a very projectible category due to the fact that these properties are perfectly clustered. More generally, Khalidi suggests that when it comes to the utility of kinds for scientific inquiry, low performance in one of the two dimensions can be compensated by a high score in the other. He says:

The two dimensions would seem to be orthogonal, though strength along one of these dimensions may help compensate for weakness on the other, when it comes to their utility for scientific inquiry. A natural kind predicate that has a multitude of other predicates associated with it might be useful in scientific inquiry even if these associations are not perfectly strict or exceptionless, and a natural kind predicate associated invariably or universally with certain other predicates may get away with not being associated with a great variety of such predicates. (Khalidi 2018: 1383)

Now, apart from the robust instability of tennessine, notice that, even if we cannot observe and measure the behavior and properties of this element by conventional means, we can nonetheless build models to *predict* many of its properties. This is crucial, as it suggests that, even when it comes to the *variety* dimension of projectibility, *tennessine* performs significantly better than what Spencer would have us believe. More specifically, some of these models have predicted that tennessine's melting point will be somewhere in the range of 350–550 °C (Hoffman, Lee & Pershina 2010: 1728), that its boiling point is 610 °C (Takahashi 2002), that it has a density between 7.1 and 7.3 g/cm³ (Bonchev & Kamenska 1981),¹⁶ and that it is *solid* at standard temperature and pressure (Bonchev & Kamenska 1981). Additionally, values for its ionization energies (Chang, Li, Dong 2010) and atomic radius (Bonchev & Kamenska 1981) have also been predicted.

Unfortunately for Spencer, these predictions are clearly at odds with the view that tennessine *lacks* the relevant properties and, in this sense, in stark tension with his assessment regarding its poor projectibility. They suggest not that tennessine is “notoriously inductively weak” but, quite to the contrary, that it can support a significant number of inductive generalizations.

¹⁶ Notice that these results' being presented in terms of intervals is again, due to an epistemic limitation. The idea is not that tennessine has no precise melting point or density, but rather that our current means of prediction do not allow us to go beyond predicting ranges.

Now, while I take these considerations to be sufficient to make the case that *element 117* is significantly projectible, there is another idea that might serve to strengthen the case, and which is thus worth presenting. According to the standard view of quantum physics, radioactive decay—the phenomenon in virtue of which unstable elements are short-lived—is probabilistic. This is important, as it entails that there is always an infinitesimal chance of a sample of tennessine lasting long enough to be tested, manipulated, measured, etc. Although extremely improbable, the fact that this constitutes a possibility gives us further reason to believe that the nuclear instability of tennessine, although an important epistemic limitation, does not affect its metaphysical status and, as such, does not constitute a reason to doubt that this element is *as projectible as* any other chemical element.

Finally, one could object that the fact that tennessine's properties have been discovered through an abductive rather than an inductive method suggests that this category is not very projectible. Indeed, Spencer seems to have something like this in mind when he suggests, later on that "117's lack of inductive power does not undercut its epistemic utility in nuclear chemistry. It's just that its epistemic utility is different. It's abductive, not inductive" (2016: 162). I contend, however, that deeming element 117 weakly projectible for such reasons would amount to confusing projectibility with the inductive *method*. In conflating these two notions one would fail to notice that the reason that projectibility is often considered distinctive of natural categories is not connected with the *method* through which we learn generalizations about them, but rather with the very fact that they support such generalizations. To insist, the distinctive feature of natural kinds—and the reason for which projectibility has generally been taken to be characteristic of them—is not that we learn things about them through any particular method (e.g. observation of particular members, followed by inductive generalization to the whole kind), but rather, that what we know and learn about them is projectible to all the members of the kind. In this sense, I conclude, *contra* Spencer, that *element 117* is significantly projectible and, accordingly, does not constitute a successful counterexample to a projectibility-based account of naturalness.

7.2 *Magnus's Polymorphism*

A similar case against the *necessity* of projectibility for naturalness is put forward by P. D. Magnus. He contends that focusing only on the inductive power of categories to

determine natural kindhood risks overlooking certain legitimate scientific categories which do not appear to be very projectible.

More to the point, Magnus suggests that focusing on projectibility ultimately leads to focusing on—and eventually overemphasizing—*similarity*. This is so, he insists, because the projectibility of a category is grounded in its members' sharing many relevant properties. He says:

Coming at natural kinds in this way [by focusing on projectibility] leads us to suppose that members of a natural kind are connected by similarity. The reason that this *A* can be used as a proxy for other *As* is that they all resemble one another in many respects. (Magnus 2012: 11)

With this connection in mind, Magnus goes on to complain that certain natural kind theories such as HPC have focused too much on similarity, and have therefore failed to see that scientific taxonomy does not always seek to individuate categories by stressing similarities. He refers to this alleged tendency of overemphasizing similarity and projectibility as *similarity fetishism*. He says:

Quine is part of a tradition, going back to Mill, which assumes that membership in the same kind is a matter of having a large number of properties in common. Call this similarity fetishism. The yoke of similarity fetishism makes the induction assumption unable to accommodate kinds which are not joined by similarity and thus makes it insufficient to serve as a definition of 'natural kind'. (Magnus 2012: 12)

To illustrate his point, Magnus focusses on the case of polymorphic species. That is, species whose members can be grouped in different subcategories according to significant recurring differences. Although many species are polymorphic (one of the most common examples being sexual dimorphism in mammals), some species stand out from the rest by exhibiting remarkably extreme differences (in morphology, behavior, etc.). According to Magnus, a projectibility-based approach to naturalness that “fetishizes” similarity among the instances of a kind would thus have no reason to group these extremely divergent morphs under the same category. To make his case more vivid, Magnus offers the example of the highly sexually dimorphic seadevil.¹⁷ He says:

¹⁷ To get an idea of how significant the differences between the female and male morphs of this species are, notice that the males, which are five times smaller than the females, were for a long time thought to be parasites attached to the females' bodies, until it was later discovered that they were essential for reproduction (these cases are known as “sexual parasites”).

Take a specific seadevil species, such as *Linophryne arborifera* [...]. Females and males are so dissimilar that there are few inductions one can make about the species in general from a single sample. If one were simply looking for projectible predicates, then the species would not be a relevant kind at all. (Magnus 2012: 160)

In what follows I will try to make the case, *contra* Magnus, that highly polymorphic species such as *Linophryne arborifera* are significantly projectible, or at least substantially more so than what he suggests. More precisely, I will argue that polymorphic species, despite diverging significantly in morphological and behavioral features, still share some very important diachronic properties (e.g. shared ancestry), in a way that supports many relevant inductive generalizations. Additionally, I will also suggest that polymorphic species share many relevant synchronic properties related to their impact on ecosystems, their habitat, and even their morphology.

Before getting into the details of Magnus's case, though, notice that Ereshefsky and Reydon (2015, forthcoming) raise a similar worry against projectibility-based views. Their contention is that biological taxonomy often focuses on highlighting *history* or *ancestry*, which, they suggest, does not always overlap with similarity. They say:

The challenge for those that assert that natural kinds are groups of entities with numerous similarities is that classifying by similarity and classifying by history can conflict. And when they do conflict, the view that natural kinds are inductive kinds fails to capture the classificatory practices of those biologists that classify by history. (Ereshefsky and Reydon forthcoming)

Magnus also seems to draw this contrast between *history* and *similarity* when he suggests that what unifies the members of a (dimorphic) species is not similarity but, rather, “a common causal history over evolutionary time” (2012: 162).

I will argue, however, that Ereshefsky and Reydon (2015, forthcoming), and Magnus (2012), are too quick to assume that similarity amounts to “superficial similarity” or, more precisely, to *synchronic similarity*. For there does not seem to be any principled reason not to count “shared history” or “shared ancestry” among the relevant properties that members of a (dimorphic) species *share*, and thus among the properties in virtue of which they can be considered to be significantly *similar*. Moreover, not only are these *diachronic similarities* relevant from the point of view of evolutionary biology, but crucially for our purposes, they ground many important inductive generalizations. This point is vividly made by Chakravartty who, against Ereshefsky and Reydon, insists that

the focus of biological taxonomy on ancestry is not in tension with highlighting inductively powerful categories. Quite to the contrary, the aim is still to make inductive inferences. He says:

This skepticism about induction, however, is unfounded on both counts. Regarding the first charge, that inductive success is immaterial to theorizing about kinds in at least some cases, crucial weight here is borne by the suggestion that ‘different aims’ may be served instead. In the case of Cladism and Evolutionary Taxonomy, ‘the aim [...] is to classify distinct branches on the Tree of Life’ [Ereshefsky and Reydon 2015] (p. 977). But to what ends? Biological taxonomy is not stamp collecting. *It serves inductive success* [my emphasis]. The point of demarcating on the basis of lineages is not to identify branches on a tree for its own sake and stop there, but rather to identify groups of organisms that are evolutionarily significant: ones that are subject, as collectives over time, to biological processes described in terms of parameters including selection, adaptation, and genetic drift. Investigating and illuminating all of this involves inductive reasoning in the service of retrodictions and evolutionary explanations regarding those kinds. These categories of organisms thus facilitate biologically important, inductive inferences. (Chakravartty forthcoming: 6)

Khalidi (2021), in investigating the aptness of etiological kinds as natural kind candidates, also suggests that these kinds, characterized by sharing diachronic rather than synchronic properties, support *retrodictions* (i.e. predictions of the past), which are a particular form of projection. He says: “For instance, if we identify a rock as a meteorite based on its fusion crust, we can infer that it had an extra-terrestrial origin and a certain causal trajectory through the earth’s atmosphere” (2021: 14). Similarly, if we identify an organism as a *Linophryne arborifera*, we can infer, for instance, how closely related it is to another given organism. Faced with these considerations, I argue that we have reasons to think that *Linophryne arborifera*, as well as other polymorphic species, will support many important retrodictions involving their evolutionary history (e.g. evolutionary closeness to other species, developmental pathway, etc.).

Additionally, as if acknowledging these shared diachronic similarities were not enough to defend *Linophryne arborifera*’s status as a significantly projectible category, notice that members of this species also share relevant synchronic properties related to their impact on ecosystems, their habitat, and even their morphology. Interestingly, even Magnus acknowledges that members of this species category share important

morphological traits. (He references Pietsch (2009: 24-30) as providing an extended account of morphological traits shared by both morphs.) Somewhat surprisingly, though, Magnus does not seem to take these morphological similarities into consideration when it comes to assessing the projectibility of the category. The reason for this, he suggests, is that the “properties of males are insufficient to diagnose species” (Magnus 2012: 161). This consideration, however, even if true, does not jeopardize the projectibility of the category as a whole. For, independently of whether the morphological traits of males are enough to individuate the species or not, inasmuch as both morphs share morphological properties that are relevant from a biological standpoint, these shared similarities contribute towards making the species category more projectible in the relevant sense.¹⁸

Overall, these considerations suggest that Magnus overestimates the impact of the divergent female and male morphologies on the projectibility of polymorphic species categories such as *L. arborifera*. As such, I conclude, *contra* Magnus, that it is not true that “if one were simply looking for projectible predicates, then the species would not be relevant at all” (2012: 160), and that *L. arborifera* is significantly projectible (in the specific non-trivial sense specified above).

Finally, it could be objected that Magnus’s case against projectibility-based accounts is not only based on the idea that this category is weakly projectible—which, as we just saw, seems doubtful—but, additionally, on the claim that projectibility is not the relevant feature in virtue of which different domains of biology favor this category. More precisely, Magnus suggests that the rationale for grouping together the members of a (dimorphic) species is not similarity or projectibility but, instead, *explanatory* considerations. He says: “Explanatory considerations identify *L. arborifera* as a legitimate taxon, even if it is not an inductively robust category” (2012: 162). The picture Magnus presents, then, is one where explanatory considerations are in tension with, and (sometimes) prioritized over, similarity and projectibility.

This view, however, is not without controversy. As Miles MacLeod (2014) suggests in his review of Magnus’s monograph, not only does Magnus provide no account of what makes a kind explanatory *qua* kind but, moreover, “it is also arguable that what grounds a kind as explanatory is similarity among its members in the first place”

¹⁸ To draw a simple comparison, consider two important shared morphological similarities of tigers: *having stripes*, and *having four legs*. It is clear that these two morphological properties by themselves are not enough to individuate the species (i.e. *Panthera tigris*). Still, these similarities contribute towards making the tiger category *more projectible*. The same goes, I suggest, for more extreme cases of dimorphic species such as *Linophryne arborifera*.

(MacLeod 2014: 337). Importantly for our purposes, if something along MacLeod's view is correct, then, by emphasizing the explanatory value of *L. arborifera*, Magnus would not thereby discount its projectibility but, quite to the contrary, provide further reasons in favor of this category being projectible in the relevant sense.

Notice that Magnus's own example seems to point in this direction. Indeed, we have seen that he identifies as explanatorily relevant the fact that members of *Linophryne arborifera* have a *common causal history*. But, if the above considerations are on the right track, focusing on a common causal history amounts to focusing on *similar* diachronic properties. In this sense, his example does not involve any tension between explanatoriness and projectibility but, instead, a case where both dimensions are grounded in the diachronic similarities of the category.

Accordingly, I conclude that Magnus's case does not succeed as a counterexample to projectibility-based accounts of naturalness. I have argued that, in virtue of their synchronic and diachronic similarities, the categories that correspond to polymorphic species support many relevant inductive generalizations. Moreover, I have shown that even if Magnus is right in his claim that the rationale for grouping polymorphic organisms under a single species category is the explanatory potential of the resulting category, this does not pose a challenge to its projectibility but, quite to the contrary, provides further reason not to doubt it.

8 Conclusion

In this paper I have put forward an original view according to which the naturalness of a kind is to be identified with its degree of projectibility. Although projectibility has traditionally been given a prominent role in natural kind theories, the current proposal departs from other theories in singling out no other additional condition for naturalness. As such, a distinctive characteristic of Radical Projectibilism is that, by identifying naturalness with a gradual property such as projectibility, the notion of naturalness itself becomes one of degree. Rather than constituting a shortcoming of the view, I have argued that understanding naturalness in a gradual way not only appropriately counters the relevant notion of arbitrariness but, moreover, brings important advantages over dichotomic alternatives. Finally, I have addressed objections involving potential

counterexamples to a projectibility-based account. Overall, I take myself to have made the case for Radical Projectibilism to be considered a plausible account of naturalness.

Ameliorating at the Joints: A Permissive Normative Framework for Conceptual Engineering

1 Introduction

Metaphysical naturalness, it can be argued, can play two opposite roles in ameliorative projects: a positive one and a negative one. Metaphysical naturalness can play a positive role if it *motivates* an amelioration. Clark and Chalmers exemplify this attitude when, in arguing for their revision of the notion of ‘belief’, they claim: ‘By using the ‘belief’ notion in a wider way, it picks out something more akin to a natural kind. The notion becomes deeper and more unified, and is more useful in explanation’ (Clark and Chalmers 1998: 14). On the other hand, metaphysical naturalness can play a negative role if it *prevents* or gives reasons against amelioration.

In a recent paper, Mona Simion (2017) presents what I take to be an articulation of this negative role. Simion argues that when a concept carves nature at its joints, we should abstain from ameliorating it—no matter the moral, practical, or other benefits—if the amelioration results in the concept no longer carving nature at its joints. That is, if the amelioration translates into an *epistemic loss* (Simion 2017: 8).

In this paper I argue that Simion’s procedure is too restrictive and yields the wrong predictions about intuitively legitimate ameliorations. I will argue, thus, for a less restrictive view according to which conceptual epistemic losses do not systematically *block* amelioration, but merely *provide reasons* against it. Crucially, this will open the possibility for epistemic losses to be, *contra* Simion, permissible, provided that they are duly compensated.

The paper is structured as follows. In *Section 2*, I present Simion’s proposal. In *Section 3* I critically discuss the notion of ‘epistemic loss’, which I contend, is not sufficiently developed in Simion’s work. I take this to be a crucial endeavour given that, to assess the permissibility—or impermissibility—of epistemic losses, we first need a clear understanding of what they amount to. I consider, thus, three possible interpretations of ‘epistemic loss’ that could be plausibly attributed to Simion and argue that only one of them survives scrutiny. Then, with the favoured notion of ‘epistemic loss’ in mind, I argue in *Section 4* for a more permissive normative procedure. According to this alternative,

epistemic losses can be permissible if the reasons against amelioration that result from the epistemic loss are compensated by the non-epistemic reasons for amelioration. I will then flesh out some of the relevant factors that will need to be taken into account in order to assess this permissibility. Finally, in *Section 5*, I put pressure on Simion's proposal by presenting two case studies that do not seem to fit with her restrictive procedure. My alternative, I argue, manages to accommodate these cases significantly better. In *Section 6* I conclude.

2 Simion's Proposal

In her paper 'The "Should" in Conceptual Engineering', Simion is concerned with the different normative sources for amelioration and the tensions that might result from their interaction. She draws from Cappelen's (2017) taxonomy of conceptual deficiencies and distinguishes the semantic, epistemic, and practical deficiencies that concepts can exhibit. She departs from Cappelen and other authors, however, in arguing that amelioration is not only about *solving deficiencies*, but also consists in *conceptual improvement*.

Simion warns us, though, that extending the scope of conceptual engineering in this way can backfire and generate tensions between different normative sources for amelioration. Indeed, we might have pragmatic or moral reasons to revise a concept in a certain way, while epistemic or semantic reasons push us in a different direction. Which reasons should take primacy?

Simion offers her proposal as a partial solution to this problem. Not every consideration, she argues, has the same normative strength. She distinguishes between reasons that contribute towards the fulfilment of the *function* of the attitude or action under consideration and reasons that do not. When it comes to concepts, Simion suggests that they have a *central epistemic function*. She formulates this in the following terms:

Concepts, just like beliefs, are representational devices, their function is an epistemic one: to represent the world. In virtue of this function, concepts will be properly functioning when responsive to epistemic reasons, and malfunctioning when responsive to practical reasons. Concepts will be good concepts qua concepts when they are epistemically good. (Simion 2017: 10)

This central function, in turn, motivates the Epistemic Limiting Procedure (*ELP*):

A representational device should be ameliorated iff (1) There is all-things-considered reason to do so and (2) The amelioration does not translate into epistemic loss. (Simion 2017: 10)

As an illustration of this procedure, Simion discusses a hypothetical revision of the biological category *deer*. Simion suggests that, one could, for practical and moral reasons, revise the biological concept DEER by excluding the subgroup *roe deer* from the concept. This exclusion, she argues, could bring some practical benefits by providing *roe deer* with a higher degree of environmental protection. Simion claims, however, that revising the concept DEER in such a way would—no matter the practical or moral benefits—amount to an *epistemic loss* that we should not be willing to accept. She says:

Think of the example of ‘deer’ again: insofar as, in the process of revising the concept, we do not end up with less knowledge about the world, the engineer is free to proceed with revision for any good practical, moral, political, etc. reason. However, no practical, moral, political, etc. reason will support conceptual revision that obscures our view of the world. (Simion 2017: 11)

In this sense, I take Simion to be on the right track when considering this hypothetical revision—let us call it DEER*¹⁹—to involve some sort of epistemic loss. I also agree with Simion that this specific revision does not seem legitimate. I contend, however, that Simion is not clear about what exactly happens in this revision and thus, what exactly an epistemic loss amounts to. One might wonder, for instance, what it takes for a revision to, as she puts it, ‘obscure our view of the world’, or result in us ‘having less knowledge about the world’.

I take this point to be in crucial need of clarification. For if we are to apply Simion’s procedure, we must be able to identify epistemic losses whenever they occur. Moreover, different accounts of ‘epistemic loss’ will likely result in different normative stances regarding its permissibility. What follows, thus, is an examination of various alternative ways of understanding ‘epistemic loss’. I consider three different readings and argue that only one of them survives scrutiny. My assessment of these alternatives is

¹⁹ I shall remain uncommitted as to whether revision preserves sameness of concept or not. I believe that the debate remains intelligible independently of whether one thinks of DEER* as *replacing* the old concept or as, instead, being the same concept after undergoing certain modifications.

based on how well they accommodate Simion's *deer* example, as well as their overall plausibility.

3 Epistemic Loss: Three Readings

Simion claims that concepts have a central epistemic function and that therefore, we should never ameliorate—no matter the practical and moral benefits—if the amelioration results in an epistemic loss. But what exactly is it for an amelioration to result in an epistemic loss? Let us consider some alternatives.

3.1 *Epistemic Loss as Ceasing to Represent a Real Kind*

Epistemic Loss₁: A conceptual epistemic loss occurs when, as a result of amelioration, a concept ceases to represent a real or existing kind.

One plausible way of understanding 'epistemic loss' is to think of it as a concept *ceasing to represent a real or existing kind*. This would be the case, for instance, if the amelioration results in concepts such as PHLOGISTON or UNICORN which, clearly, fail to represent real kinds in the world.

This interpretation could well serve Simion's purposes insofar as it does not seem to be generally permissible—fiction being perhaps an interesting exception—to allow for *epistemic losses₁*. Most conceptual engineers, I take it, would dismiss ameliorations resulting in concepts that utterly fail representationally, that is, that do not represent a real kind. We could concede to Simion, therefore, that *epistemic losses₁* should never be permissible in ameliorative projects, no matter their potential moral or practical benefits.²⁰

The main problem with *Epistemic Loss₁*, however, is that it does not seem to be capable of accommodating Simion's own example. Recall that Simion rejects revising the concept DEER by excluding *roe deer* because this revision, she argues, results in an epistemic loss. But does the revised concept DEER* fail to represent a real kind?

²⁰ Notice that if a given concept does not represent a real kind in the first place, it would not be epistemically problematic for it not to represent one after revision. Indeed, for an amelioration to result in an epistemic loss in this sense it must involve a transition from a concept that represents a real kind to one that does not. The revision of the concept UNICORN, for instance, is not susceptible of resulting in an epistemic loss₁. I thank an anonymous referee for raising this point.

Although it seems uncontroversial that DEER* is quite arbitrary according to biological taxonomical standards²¹—remember that DEER* excludes *roe deer* by relying on moral considerations, which are generally not taken to be relevant from a biological standpoint—it is something else altogether to claim that it is not a *real* kind. Indeed, if one adopts a permissivist metaphysical framework (see Schaffer 2009) there is no reason not to count *deer** as a real kind, despite its arbitrariness. In this sense, DEER* is fundamentally different from PHLOGISTON or UNICORN, in that these latter concepts fail to track any real kind in the world, while the former, arbitrary as it is, does track a non-empty grouping of particulars.

Of course, this way of approaching the issue assumes a permissivist metaphysical framework, which does not equate the *reality* of a kind with its *naturalness*. Not everyone, however, might be willing to accept this. Some people, indeed, might disagree and contend that *only natural kinds exist*. Natural kinds are characterized by being *discovered*, whereas non-natural or arbitrary collections, they would argue, are ‘human constructs’ with no independent existence outside of human minds and perspectives. On such a view, one could maintain that DEER tracks a natural and thus, real kind, while DEER* does not. In doing so, one could reconcile Simion’s example with *Epistemic Loss₁*, for transitioning from one concept to the other would amount to *ceasing to represent a real kind*.

As I will show next, however, accommodating the *deer* case by appealing to the natural/non-natural distinction brings its own specific problems.

3.2 *Epistemic Loss as Ceasing to Represent a Natural Kind*

Epistemic Loss₂: A conceptual epistemic loss occurs when, as a result of amelioration, a concept ceases to represent a natural kind.

The second way of understanding ‘epistemic loss’ is to think of it as a concept *ceasing to represent a natural kind*.²² An *Epistemic Loss₂* occurs when amelioration turns a concept that hitherto tracked a natural kind into one which does not.

²¹ See Ereshefsky (2001: 158-162) for a more detailed discussion of what makes a biological species concept legitimate.

²² I follow most authors in the natural kind literature in using the label ‘natural’ to refer to ways in which particulars can be grouped together. ‘Natural’, thus, contrasts with ‘arbitrary’ but not necessarily with ‘social’. Objects from the natural world can be grouped in very arbitrary or non-natural ways, just as artifacts or objects from the social world can be grouped in very natural ways.

Simion seems to have a picture of this sort in mind when she uses the ‘carving at the joints’ metaphor that is so prevalent in the natural kind literature. She says:

Does this justify revising our concept of deer? Should we be responsive to the moral and political factors and abandon a perfectly functional concept, which, by stipulation, carves nature at its biological joints? Think, for instance, of all the epistemic work that such a concept can do for us, and which will be lost after moral amelioration. (Simion 2017: 8)

It would seem, indeed, that *Epistemic Loss*₂ is better suited to accommodate the *deer* example than *Epistemic Loss*₁. More precisely, one might think that DEER represents a natural kind (i.e., *deer*), whereas DEER* does not. A revision involving a transition from the former to the latter, therefore, would constitute an *Epistemic Loss*₂, one which Simion prohibits.

There are, however, several problems with this view. The first, and more specific to the case, is that this approach relies on an account of natural kinds that draws the border between natural and non-natural kinds somewhere in between *deer* and *deer**. Although such a view might be defensible, it will nonetheless be highly controversial.

To see why, notice the vast array of alternative accounts of natural kinds available on the market, many of which draw the natural/non-natural demarcatory line in different places and by different criteria. It should not come as a surprise, then, that some of these accounts would similarly dismiss both *deer* and *deer** as non-natural (e.g., Ellis 2001, 2008, Churchland 1985) or, on the contrary, accept both as natural kinds (Dupré 1993, Ludwig 2018, Hacking 2007).

The second—more serious—problem is that given such an account of epistemic loss, the *range* of Simion’s procedure risks becoming seriously limited. If epistemic losses are understood as involving a transition from natural kinds to a lack thereof, *ELP* will fail to constrain any amelioration that does not involve this sort of transition. But many ameliorative projects, I suggest, fail to feature such transitions and, nevertheless, seem susceptible to involving epistemic losses of some sort.²³

²³ This will of course vary depending on the theory of natural kinds one favors, given that different theories identify as “natural” different sets of kinds. Still, most of them draw the natural/non-natural border somewhere and, as such, exclude some potential ameliorations once incorporated into the reading of epistemic loss under consideration.

The sorts of projects I have in mind can be found on two sides of a spectrum: ameliorations involving transitions from arguably non-natural kinds to other non-natural kinds and, conversely, transitions from arguably natural kinds to other natural kinds. I suggest that, given that transitions of both sorts can be accurately described—in Simion’s own metaphorical terms—as ‘obscuring our view of the world’, it is desirable for our notion of epistemic loss to be successfully applicable to these as well.

Consider, to begin with, transitions of the first sort. That is, ameliorations which start with an arguably non-natural kind. On the reading of epistemic loss under consideration, such ameliorations could never result in an epistemic loss, for the initial concept does not track a natural kind in the first place and, hence, does not have any “natural” status to lose. This, however, seems unduly restrictive given that many ameliorative projects involve kinds that, very plausibly, most of the accounts under consideration would not consider natural. Think—to name just a few—of conceptual debates around ‘sandwich’, ‘journalist’, ‘vehicle’, ‘torture’, ‘marriage’, etc.^{24 25} The fact that the revision of these concepts seems susceptible of, in Simion’s words, ‘obscuring our view of the world’ suggests that epistemic loss₂ is too narrow. To insist, if *ELP* is to constraint ameliorative projects in a significant way it needs to incorporate an account of epistemic loss applicable to this sort of cases too.

The second sort of ameliorative project to which this reading of epistemic loss fails to apply involves a transition from an (arguably) natural kind to another (arguably) natural kind. Although such cases might be rarer, ameliorative projects might feature this sort of transition too. Recall, for instance, that according to some of the accounts of natural kinds we mentioned earlier, both *deer* and *deer** would count as natural. Still, it seems that some sort of epistemic loss occurs in this revision: one that our reading of ‘epistemic loss’ ought to capture. In section 4.2 below, while discussing my second case study (*SPECIES*), I will illustrate this point further.

These considerations suggest that if ‘epistemic loss’ is understood as involving a transition from a natural to non-natural kind, then Simion’s epistemic limiting procedure will fail to constrain many ameliorative projects that seemingly ought to be epistemically constrained.

²⁴ It may happen that some of these kinds turn out to be natural according to these accounts. But to anticipate: the legitimacy of the revision seems independent of this.

²⁵ See Peter Ludlow’s *Living Words* (2014) for a more detailed survey of many of the debates around these concepts.

In addition, there is a related concern with this reading of ‘epistemic loss’. For it assumes naturalness to be a *dichotomic* notion; so that whether naturalness applies to kinds is not a matter of degree. In the following section, I consider this possibility and argue that the notion of *more natural than* provides a better reading of the kind of ‘epistemic loss’ that plausibly occurs in the transition from DEER to DEER*. As we shall see, this seems to be the best reading of ‘epistemic loss’, for it does not exhibit any of the shortcomings of the previous two.

3.3 *Epistemic Loss as Less Natural*

Epistemic Loss₃: A conceptual epistemic loss occurs when, as a result of amelioration, a concept represents a *less natural* kind than the pre-ameliorated concept did.

This third way of understanding ‘epistemic loss’—the one I shall favor—shares with the second the idea that *naturalness* is the main dimension (although plausibly not the only one) along which the epistemic performance of our concepts is to be evaluated. What distinguishes this reading from the previous one, though, is that here the epistemic loss is understood as being a matter of degree. *Epistemic loss₃* does not occur when a concept ceases to track a natural kind, but rather, when the revision results in a concept tracking a *less natural* kind.

Although a full-fledged account of a gradual notion of naturalness cannot be presented here, we can at least give a rough picture of how it would go, making reference to David Lewis’s work. In ‘New Work for a Theory of Universals’ (1983), Lewis argues that distinguishing natural from non-natural properties—where the former are taken to be sparse and the latter abundant and thus, indiscriminating—can help to account for several phenomena, including supervenience, lawhood, causation, etc. More importantly for our purposes, Lewis also suggests that the relevant notion of naturalness might be one of *degree*, and so fall along a spectrum. The perfectly natural end of the spectrum would be occupied by a privileged group of properties that appear in the laws of nature, account for resemblance and causal powers, and ground other properties, whereas the highly non-natural properties at the spectrum’s other end would be indiscriminating and generally useless. This view nicely captures the common intuition that the property *green*—although not perfectly joint-carving—is more joint-carving than *grue*, or that the property

of *being a planet* is more joint-carving or natural than the property *discovered-on-a-Tuesday*. Moreover, understanding naturalness in this way frees us from the daunting task of finding a dichotomic (non-stipulative) distinction between natural kinds and mere groupings of particulars.

Notice, in any case, that there are different ways of articulating in detail a gradual notion of naturalness. Given that this is not the central purpose of this paper, I shall rely on the often-held view that inductive and explanatory power are the main indicators of a kind's naturalness. Although this picture is not without controversy, I take it to be common and extensive enough to serve the purposes of this work.

Coming back to the case at hand, this gradual conception of naturalness—and the complementary notion of *Epistemic Loss*₃—is better suited, I argue, to accommodate Simion's *deer* example. For, very plausibly, *deer** is *less natural* than *deer*. What is going on in this revision, then, is not a drastic transition from a real to an unreal kind, nor from a natural to a non-natural kind. Rather, there is a loss of naturalness which, depending on how significant it is taken to be, can be responded to in different ways (I shall develop this point further in the next section).

Finally, this account of epistemic loss does not suffer the problems associated with the previous reading discussed above. More precisely, this way of understanding 'epistemic loss' can accommodate a broader range of cases since it does not rely on a dichotomic notion of naturalness and thus, does not require a natural kind to be involved in the first place. In other words, the amelioration of any concept is, under this reading, susceptible of translating into an epistemic loss (i.e., ameliorating concepts such as MARRIAGE, TORTURE, or SANDWICH—or any paradigmatic non-natural kind—could translate into an epistemic loss).

4 An Alternative Normative Framework

Now, does understanding 'epistemic loss' in a gradual way jeopardize Simion's procedure? Not much, it could be argued. For Simion could concede that epistemic losses are a matter of degree while maintaining that *any* epistemic loss—no matter how big or small—blocks amelioration.

I suggest, however, that, though defensible, this amendment would make Simion's procedure more controversial. Indeed, while many would grant that transitions from real to unreal kinds, or from natural to non-natural kinds, have enough normative force to

systematically block amelioration, fewer would hold that the same goes for *any* epistemic loss, regardless of how big or small it is. After all, according to the gradual reading I have argued for, a concept suffering an *Epistemic Loss*₃ may still carve nature reasonably well.²⁶ According to any of the dichotomic readings we have considered and dismissed, on the contrary, a concept that has suffered an epistemic loss is, by definition, either utterly arbitrary or does not track a real kind. In this sense, it seems reasonable to think that, when ameliorating, a transition to a concept that, although less natural, still carves nature reasonably well, does not generate as strong reasons against amelioration as would a transition to an utterly arbitrary or unreal kind.

My alternative normative framework thus departs from Simion's and drops the idea that epistemic losses *systematically block amelioration*. I contend, instead, that epistemic losses merely *provide reasons against amelioration*. This entails that an epistemic loss will have to be *weighed* against the non-epistemic reasons motivating the amelioration in order to determine whether it is permissible or not.

Although this is a substantial difference with regard to Simion's procedure, my proposal preserves the idea that there are epistemic constraints on ameliorating our representational devices. These constraints, though, as we shall see, are less restrictive than Simion's.

More precisely, on my view, for an epistemic loss to be permissible it will have to be *compensated* by the non-epistemic gains of the amelioration. Although assessing when this obtains is not an easy task, we can at least identify various relevant factors that will need to be taken into account. One such factor concerns the magnitude of the epistemic loss. Other things being equal, the bigger an epistemic loss, the stronger the reason it generates against amelioration (i.e., minor epistemic losses will tend to be easier to outweigh than "gross" ones). We should also factor the strength, so to say, of the non-epistemic reasons for amelioration, for these can vary from the completely illegitimate to the most pressing and imperative. Moreover, how these two are balanced will be relative to the concept under consideration.²⁷ Indeed, concepts might serve a variety of functions

²⁶ For stylistic reasons I speak here and elsewhere as if concepts themselves suffered the epistemic loss. Keep in mind, though, that concepts themselves are not the locus of the epistemic loss. Rather, the epistemic loss is the result of the ameliorative process, and it is the cognizer who uses or has the concept that "suffers" such loss.

²⁷ Take the case of the attempted revision of the concept PLANET by the state of New Mexico, which, in conflict with the consensus of the astronomical community, counts Pluto as a planet. Among the reasons for this decision, the corresponding authorities are proud to include the fact that the discoverer of Pluto, Clyde Tombaugh, was based at the university of New Mexico. I leave it to the reader to consider the strength

and, as such, it would be misguided to expect that they all should be assessed with identical criteria and by the same standards. For instance, it does not seem appropriate to expect the same level of epistemic performance from scientific and folk concepts alike. One way of fleshing out this idea thus consists in signalling that the epistemic losses affecting ameliorations of epistemically driven concepts (e.g., scientific concepts) will generally generate stronger reasons against amelioration and, as such, will tend to be harder to compensate. Conversely, epistemic losses affecting the amelioration of folk concepts will, other things being equal, be easier to compensate.

Now, these considerations and the incorporation of the gradual reading of epistemic loss, by themselves, are not definitive reasons to reject Simion's procedure for, as I said before, one could still hold that *any* epistemic loss, no matter how big or small, should systematically block amelioration. The problem with this position, though, is that it seems very distant from many of our conceptual practices. What follows, then, is an attempt to put more pressure on this restrictive position by presenting two case studies that do not seem to fit with this framework. The more permissive alternative that I just presented, I argue, fares significantly better.

5 Two Case Studies: MARRIAGE and SPECIES

In this section I shall put pressure on Simion by presenting specific examples that seem difficult to accommodate within her framework. I discuss two case studies: the revision of MARRIAGE and the conceptual choice between alternative SPECIES concepts.

These examples have in common two elements which, I contend, expose a serious challenge for Simion's epistemic limiting procedure. The first is that they seem to be paradigmatic cases of successful amelioration. The second is that they both have likely involved an epistemic loss.²⁸ This is problematic for Simion, for these two elements are, according to her procedure, incompatible; recall that on her account, amelioration cannot succeed if it involves an epistemic loss.

of this reason for revision in contrast, for instance, with the non-epistemic reasons for the amelioration of MARRIAGE.

²⁸ Notice here the use of 'likely'. I cannot commit beyond this probabilistic claim, for epistemic losses are not always easy to identify or measure. Although extreme cases are easy to compare (e.g. GREEN versus GRUE), the same does not hold for kinds that oscillate along similar ranges of naturalness, such as those I will be dealing with.

Simion's procedure thus does not seem to yield the right predictions about these cases. On the contrary, were she to maintain that these revisions are, as everything seems to suggest, legitimate, she would be forced to conditionalize such legitimacy on the total absence of any epistemic loss. On the alternative view defended here, in contrast, we only need to conditionalize their legitimacy on any epistemic loss being compensated by the moral and political considerations that motivate the amelioration.

To see this problem more clearly, let us begin by considering the revision of MARRIAGE.

5.1 *The Revision of MARRIAGE*

I take the revision of 'marriage' to be a paradigmatic case of successful amelioration. Its success, moreover, does not stem from an improved epistemic performance. On the contrary, the revision of MARRIAGE seems successful because of the political and moral benefits it has brought and, crucially, quite independently from any epistemic function that we could ascribe to it.

Without going into great detail, it can be said that the concept MARRIAGE has transitioned from including only heterosexual couples, to also including gay couples. Let us use MARRIAGE_{EX} for the old exclusionary concept and MARRIAGE_{IN} for the revised inclusive one. Keep in mind, also, that the lexical item 'marriage' has remained the same throughout this revision.

Now, although it is certainly difficult to assess whether the kind *marriage*_{IN} is more or less natural than *marriage*_{EX} and thus, difficult to determine whether an epistemic loss has occurred through this revision, I want to suggest that this possibility is not unlikely. To see how this might be so, notice that the exclusionary concept MARRIAGE_{EX} not only tracked the institutional agreement that a man and a woman might endorse, but also overlapped with a cluster of social heteronormative practices that are deeply rooted in most Western societies. Given this, it is likely that the kind *marriage*_{EX} could have supported some inductive generalizations involving classic heteronormative behaviors (e.g., generalizations involving household task distribution, classic different-sex couple dynamics, care-work distribution, etc.), some of which are susceptible to being lost or weakened after revision and the consequent inclusion of non-heteronormative couples under the concept.

Before discussing the implications of this possible epistemic loss, an important remark is due. According to Simion, when it comes to social concepts such as MARRIAGE, there is no risk of epistemic loss, and thus no epistemic constraint preventing their amelioration. This is so, she argues, because social concepts *constitute* the reality they represent and, thus, any change in the concept entails a change of the reality thereby represented. She says: ‘[S]ince as soon as one engineers a social kind concept, reality follows suit, it is unlikely to ever get epistemic deficiencies’ (Simion 2017: 13).

Although it is certainly the case that social concepts can (partly) constitute the reality they represent, it does not follow from this that their amelioration cannot translate into an epistemic loss—provided, of course, that we have a good articulation of the notion of epistemic loss at hand. In fact, I believe that Simion’s position on the epistemic “immunity” of social concepts rests on the somehow underdeveloped notion of ‘epistemic loss’ that I tried to disentangle in the first half of this paper. With the more precise reading of ‘epistemic loss’ (i.e., as loss of naturalness) at hand, however, there is no reason to think that the amelioration of social concepts cannot translate into epistemic losses too. Indeed, as shown above, the amelioration of MARRIAGE could have involved an epistemic loss.²⁹

Now, when confronted with this possible epistemic loss, Simion faces an important difficulty. For, according to her procedure, the revision of MARRIAGE succeeds *only* under the condition that no epistemic loss has occurred. That is, Simion would be forced to reject the inclusive concept MARRIAGE_{IN}, were we to confirm that such loss, no matter how big or small, has indeed occurred. Given our strong intuitions regarding the success of the revision of marriage—which seem unaffected by any hypothetical epistemic loss—this prediction seems highly counterintuitive.

In contrast, the alternative framework I have put forward is more flexible and seems better equipped to deal with the amelioration of MARRIAGE. For, according to this view, the legitimacy of this amelioration does not depend on the *absence* of an epistemic loss but can, on the contrary, tolerate it. The legitimacy of the revision of MARRIAGE, on my view, depends on the far more likely scenario of this epistemic loss having been reasonably *compensated*. That is, for the epistemic loss to be permissible, the reasons against amelioration that the epistemic loss generates need to have been outweighed by

²⁹ See Podosky’s (2018) reply to Simion for an alternative discussion of the relation between epistemic losses and the amelioration of social kind concepts.

the non-epistemic reasons for amelioration. Let us consider how this might have gone in the marriage case.

Setting aside for now the difficult task of assessing the *magnitude* of the epistemic loss (if any)—recall that the bigger the loss, the stronger the reason against amelioration—the most compelling reason to think that this compensation occurred is to be found, I believe, in the strength of the non-epistemic gains of the amelioration of MARRIAGE. Extending the concept of MARRIAGE to gay couples is a matter of social justice: important and pressing enough to compensate, it seems, any potential epistemic loss. Moreover, recall that the nature of the concept being ameliorated plays an important role in how the reasons for and against amelioration are balanced against one another. More precisely, an epistemic loss affecting a folk concept such as MARRIAGE is less significant than one that would affect a scientific concept, which, arguably, is more epistemically driven.

That being the case, the reason against amelioration that the epistemic loss generated in the marriage case is weaker than any that would affect the amelioration of, say, ELECTRON or VACCINE. As such, the epistemic loss involved in the amelioration of MARRIAGE is more easily compensated by the non-epistemic reasons that motivate the amelioration. Although these considerations are not definitive, they show how the epistemic loss of this specific amelioration might have been compensated.

At this point, however, an objection could be raised. It could be argued that my account is still too restrictive, for it leaves open the possibility that the amelioration of MARRIAGE is illegitimate. Indeed, if part of the strength of the reasons against amelioration derives from the magnitude of the epistemic loss, then there is, in principle, room for a huge epistemic loss to outweigh the non-epistemic reasons for amelioration and thereby block it or render it illegitimate. Some would take this result to be undesirable for, they would argue, the political and moral reasons motivating the amelioration of MARRIAGE should compensate for *any* epistemic loss, no matter its magnitude.

One way to further develop this objection would be to argue that some concepts serve non-epistemic functions and, as such, the epistemic losses in which they are involved are irrelevant and do not generate reasons against amelioration. To insist on this point, one might say that concepts such as MARRIAGE are “indifferent” to naturalness and, in this sense, epistemic losses should not play any part in assessing their amelioration.

Although I am sympathetic to this position—at least when it comes to the intuition that not all concepts ought to be assessed with identical criteria—I think that it is

ultimately misguided. I want to insist, then, on the claim that epistemic losses *always* generate reasons against amelioration—weak as they might sometimes be. Moreover, I believe that my proposal can respect the spirit of this objection, if not the letter.

It seems, indeed, that it is not the main function of concepts such as MARRIAGE to carve nature in the most natural way possible. Some concepts, it would seem, have other main functions, pragmatic or otherwise, in relation to which metaphysical naturalness does not appear to be especially important. One could reasonably wonder, then, why naturalness should play any role in assessing the amelioration of such concepts.

To understand why, we need to draw a distinction between two alternative ways in which naturalness can be valuable and, in this sense, generate reasons for and against amelioration. Naturalness, I submit, can be valued both *intrinsically* and *instrumentally*. I wish to suggest, moreover, that even when concepts do not pursue naturalness as an intrinsic good, they still value it instrumentally as a *means* for other non-epistemic purposes.

Metaphysical naturalness is valued *intrinsically* as an epistemic good. We take ourselves to be in a better epistemic position as cognizers when we have green-like concepts rather than grue-like ones. We value naturalness, in these cases, because having a more natural-kind-tracking concept gives us a better access to the objective structure of the world. Independently of any possible collateral non-epistemic benefit, it is better, for instance, to have a concept of FISH that excludes whales, and a concept of PLANETS that excludes Pluto. We are in a better epistemic position when we have these arguably more natural concepts than we were before amelioration. Our grasp of the structure of the world is, in a basic sense, better with these arguably more natural concepts.

This attitude is best exemplified in scientific conceptual and classificatory practices, which have as one of their main goals to carve reality in the most natural way possible. Naturalness, in such cases, is valued intrinsically as an epistemic good. Of course, as we will see when discussing the case of species, it would be mistaken to believe that this quest for naturalness is the *only and always primary* goal of scientific conceptual practices. It would be just as wrong, however, not to acknowledge it among one of the most important goals.

But naturalness can also be valued *instrumentally*, as a *means* for other non-epistemic purposes. That is, even when the function of a given concept is not epistemic, it will often be the case that the fulfilment of this non-epistemic function depends on the concept tracking a kind with a certain degree of naturalness.

To see why, notice that having concepts that track natural kinds not only gives us access to the objective structure of the world and puts us in a better epistemic position. Apart from this, having such concepts is essential for innumerable non-epistemic activities which would be completely impossible were we to have concepts that didn't allow for any generalizations or inductive projections. Perhaps, the spirit of the instrumental value of naturalness is best captured by Richard Boyd's words when he suggests that 'Gruified gardening would be as unsuccessful as gruified mineralogy' (Boyd 1999: 161), or by William James's metaphor of the seven-league boots:

Kinds, and sameness of kind—what colossally useful *denkmittel* for finding our way among the many! The manyness might conceivably have been absolute. Experiences might have all been singulars, no one of them occurring twice. In such a world logic would have had no application; for kind and sameness of kind are logic's only instruments. Once we know that whatever is of a kind is also of that kind's kind, we can travel through the universe as if with seven-league boots. (James 1907: 179)

With these considerations in mind, and coming back to the objection to the marriage case that motivated this detour, it can hopefully be seen now how it is problematic to argue that *any* epistemic loss, no matter how big, can be compensated by the moral and political reasons for amelioration. For, very plausibly, a dramatic epistemic loss in the amelioration of MARRIAGE could result in a concept that does not even serve the political and moral purposes that motivated the revision. Consider, for instance, the hypothetical amelioration of MARRIAGE resulting in a completely gerrymandered kind such as, for instance, *marriage_{IN-or-bigger-than-my-wallet}*. Apart from the seemingly impossible task of successfully implementing such a concept in any linguistic community, I take it that this hypothetical epistemic loss would be harder—if possible at all—for any moral or political gain to compensate.

In the light of this, and against the objection just considered, I believe that it is a virtue of my normative framework that it always counts epistemic losses as generating reasons against revision. This captures well enough our unwillingness to use very gerrymandered kind-tracking concepts and prevents extreme cases like the one just considered.

Let us now turn to discussing the case of the concept of SPECIES.

5.2 *Conceptual Choice: The Concept of SPECIES*

Before going into the details of this case, an important clarification is due. Whereas the case of MARRIAGE involved a paradigmatic case of revision—where a concept is changed or replaced by another—this is not so in the case of the SPECIES concept. What I shall discuss in relation to this case is better described as an instance of *conceptual choice*. As I will be using the phrase, ‘conceptual choice’ involves choosing between alternative concepts that *coexist* in the vicinity of one another. In this sense, conceptual choice does not consist in *replacing* a concept by another in *all* relevant contexts but, instead, assumes a form of *pluralism* about the relevant concept—in this case, SPECIES. Although there are undoubtedly significant differences to be acknowledged between revision and conceptual choice, I take both of these to be alternative instances of amelioration, and as such, I take the normative concerns driving this discussion to affect them in similar enough ways. More precisely, what is important for our purposes is to identify the normative considerations relevant to assessing, roughly, which concept we ought to favor in a particular context. Now, and interesting as this further question may be, whether or not this involves replacement, change, elimination, or even coexistence, need not greatly divert our discussion.

The species case, moreover, is interesting for yet another reason: we are dealing here with a *scientific* concept. It could be argued that whereas epistemic losses are permissible when it comes to revisions and conceptual choices involving folk concepts (whose main function is arguably not epistemic), the same does not hold for scientific concepts. When it comes to the latter, one could argue, epistemic losses are always impermissible, no matter their magnitude. Indeed, if this were right, Simion’s procedure could be preserved if restricted to scientific concepts only.

But, as our subsequent discussion will illustrate, Simion’s procedure is still overly restrictive when it comes to scientific concepts. As in the marriage case, this second example features an (apparently) legitimate conceptual choice between alternative SPECIES concepts which, crucially, could have involved an epistemic loss. This possibility constitutes a serious problem for Simion since, again, she is forced to either commit to the controversial claim that no epistemic loss has occurred, or else to contend, in opposition to biological expertise, that this conceptual choice is illegitimate. As we will see, neither of these alternatives seems very attractive. Let us consider this point in some detail.

The concept SPECIES refers to the basic unit of classification of living organisms and it is the focus of important debates among biologists. Taxonomists are far from reaching any consensus on how to define ‘species’ and, in fact, *species pluralism*—the view that multiple concepts of species can coexist and serve different fields and purposes—is today the most popular stance (Dupré 1993; Ereshefsky 2001; Kitcher 1984, 1987). A significant point in favor of pluralism is that what appears to be a particularly useful concept in one context can be especially ineffective in another.

Take, for instance, the *biological species concept* (BSC). According to the BSC, ‘a species is a group of interbreeding natural populations that is reproductively isolated from other such groups and can produce fertile offspring’ (Mayr and Ashlock 1991: 26). This species concept is one of the most popular, since testing whether two organisms can interbreed and produce fertile offspring gives us an exceptionally good idea of how closely connected they are. The *biological species concept*, however, fails notoriously when it comes to asexual organisms and, more importantly for present purposes, fossils; given that the former do not interbreed and, the latter, being extinct, cannot be checked as to their reproductive behavior. On the face of it, biologists working on these fields turn to other—more operative—concepts, better suited for their objects of study. Paleontologists, in particular, tend to rely on the *phylogenetic species concept* (PSC), which places special emphasis on tracking phylogenetic relations (ancestry), and draws distinctions primarily on the basis of morphological traits.

Although, in the context of a particular subfield, it is easy to see how some concepts outperform others, it is certainly difficult to compare different species concepts in terms of their *naturalness*. Attempting to determine, for instance, whether BSC is more or less natural than PSC does not seem a promising enterprise, if possible at all.

What is interesting to acknowledge about this case, however, is that *even if we were* to find out that PSC tracks a less natural kind than BSC, this would hardly dissuade paleontologists from using the former—hypothetically less natural—concept. This shouldn’t be too surprising, though. For, while it is certainly the case that paleontologists require a concept of SPECIES that allows for strong predictions and generalizations, they also need a concept that can be satisfactorily applied to fossils. It seems reasonable, therefore, to think that paleontologists would be ready to assume an epistemic loss in exchange for a concept that satisfies this function.

Simion, however, cannot accommodate this possibility. For according to her procedure, were we to find out that the PSC tracks a less natural kind than BSC, this would

immediately turn the conceptual choice of the paleontologist into an illegitimate one. Simion, again, needs to pay a high price and conditionalize the success of this conceptual choice on the total absence of any epistemic loss.

My alternative, I contend, better accommodates this case, for it is permissive enough to allow for epistemic losses provided that they are reasonably compensated. As such, it can make sense of the fact that paleontologists would be ready to assume an epistemic loss in exchange for a concept that fits the particularities of their field.

In addition, this permissive framework is still restrictive enough to yield the right predictions regarding cases which apparently ought to be excluded and which, I venture, paleontologists themselves would exclude. For instance, no paleontologist would be willing to use a SPECIES concept that, although somehow applicable to fossils, would be so gerrymandered as to allow no interesting generalizations or predictions. My account accommodates this sort of scenario too. To see why, it suffices to recall that the bigger the epistemic loss, the stronger the reason against amelioration, and, as a result, the harder it is to compensate the loss. As such, in all probability, an epistemic loss as huge as the one we are imagining will not be compensated by the pragmatic motivations of the paleontologist.

The species case, far from being an isolated example, illustrates a wider phenomenon prevalent in scientific conceptual and classificatory practices. The idealistic picture according to which the scientist *discovers* the kinds that exist “out there” without any sort of intervention or choice-making is untenable. Quite the contrary, pragmatic limitations and concerns often drive scientific conceptual practices and, as a result, epistemic and non-epistemic considerations often stand in tension and have to be weighed against one another. In this sense, I believe that a flexible and permissive normative framework such as the one I have put forward captures this messy situation quite well.

6 Conclusion

In this paper, I have argued against Simion’s Epistemic Limiting Procedure for amelioration and presented a more permissive alternative. More precisely, I have argued that epistemic losses do not systematically prevent amelioration but merely provide reasons against it. This constitutes a substantive difference in relation to Simion’s procedure, for it opens the possibility for epistemic losses to be permissible, provided that they are compensated by the non-epistemic reasons that motivate the amelioration.

To defend this alternative, and after trying to articulate the notion of ‘epistemic loss’, I have put pressure on Simion’s view by presenting two case studies that do not fit her normative picture. These are cases of seemingly successful revisions that could have involved epistemic loss; something that Simion’s restrictive picture cannot accommodate. I have argued that my alternative, in contrast, succeeds in accommodating these cases and more generally, better captures our conceptual practices and intuitions regarding the legitimacy or illegitimacy of ameliorative projects.

Is Sex Binary? Yes and No, and a Terminological Proposal

1 Introduction

Is sex binary? Interestingly, we are likely to get very different answers to this question depending on who we ask.

In an article published in *Nature*, Claire Ainsworth (2015) claims that “the idea of two sexes is simplistic. Biologists now think there is a wider spectrum than that.” In a similar vein, biologist Anne Fausto-Sterling claims, in her 2018 *New York Times* op-ed “Why Sex is Not Binary”, that “[t]wo sexes have never been enough to describe human variety”, and that “[i]t has long been known that there is no single biological measure that unassailably places each and every human into one of two categories—male or female.” Along similar lines, in a piece published in *Scientific American*, Sun (2019) claims that “actual research shows that sex is anything but binary”.

In what might seem to be a sharp contrast, philosopher Alex Byrne (2018) argues in his reply to Fausto-Sterling (2018) that there are only two sexes, which are actually not that difficult to identify. He says: “Specifically, females produce large gametes (reproductive cells), and males produce small ones. (Since there are no species with a third intermediate gamete size, there are only two sexes.)” Similarly, Kathleen Stock (2021), in her recent book *Material Girls*, devotes a chapter to defending the claim that “binary sex exists”, that is, the view that “humans are divided into females and males, and that this binary division is a natural state of affairs rooted in stable biological fact” (2021: 45).

In this paper, I argue that, contrary to appearances, this dispute does not involve any disagreement concerning *facts* about human sexual biology. I will argue, moreover, that at least part of the dispute derives from a tension around the term ‘binary’, which, when used to characterize sex, allows for alternative interpretations. To anticipate, ‘binary’ can convey, on the one hand, (i) that there are just two relevant kinds, and on the other hand, (ii) that there are just two relevant groups of people. In other words, “Sex is binary” can be interpreted as the claim that there are just two sexes or, instead, as the claim that everyone is either one or the other. As we will see below, however, while (i) is consistent with the relevant biological facts, (ii) is false. As Paul Griffiths (2020) puts it: “Yes, there are only two sexes. No, this does not mean that every living being is either

one or the other.” That being so, although sex might be said to be binary in the sense of (i), it clearly isn’t in the sense of (ii).

The first part of this paper, then, is devoted to partly summarizing the views of some of the contenders in this debate and showing how, initial appearances notwithstanding, there is no factual disagreement involved concerning the relevant biological facts—at least, as we will see, once they are appropriately formulated. In this sense, this first part of the paper is somehow conciliatory, in that I will contend that all parties end up agreeing on the relevant biological facts—which is not to say that there are no remaining issues in dispute. What these remaining issues are, however, is not something that will concern me here. (Although I will note that they are arguably partly terminological, in a way that may reflect an underlying disagreement in values.)

The second part of the paper, in contrast, has a more normative character. There I will argue that characterizing sex as “binary” can be both epistemically and morally problematic. In light of this, I propose to avoid ‘binary’ and to use ‘dimorphic’ instead. The main reason for this terminological move, I shall argue, is that ‘dimorphic’ preserves the important function that ‘binary’ presumably fulfils, while avoiding its potentially undesirable implications.

2 The dispute

In this section, I introduce what I take to be the two most important facts about human sexual biology for the dispute that concerns us here. Then, I summarize some rival views of the nature of sex by paying special attention to how their various authors position themselves *vis-à-vis* these two important facts. In doing so, I show how, contrary to appearances, there is no disagreement concerning either of them.

2.1 The relevant biological facts

Before introducing the two relevant facts around which much of our subsequent discussion will revolve, an important disclaimer is due. In order to successfully assess whether or not the targeted dispute involves a factual disagreement, we need to formulate the relevant facts without using highly controversial terms such as ‘sex’, ‘male’, or ‘female’. Indeed, significant issues depend on how these terms are used (see Chalmers 2011; López de Sa MS), and positioning oneself in that regard already constitutes an

important commitment which, at this early stage, it is preferable to avoid. Moreover, it is precisely by stripping the relevant facts of this contested terminology that we will be able to appreciate that the contenders in this debate do not disagree with either of them. That being said, let us introduce these two facts, which I will label *DIMORPHISM* and *NON-EXHAUSTIVITY*:

DIMORPHISM: Some species reproduce sexually by fusing a relatively large gamete with a relatively small gamete. Many species exhibit further morphological differences³⁰ associated with the production of each of these two gamete types. In the case of humans, producing³¹ relatively large gametes tends to cluster with other physiological properties such as having XX chromosomes, having a vagina, having ovaries, etc., whereas producing relatively small gametes tends to cluster with other physiological properties such as having XY chromosomes, having a penis, having testes, etc.

NON-EXHAUSTIVITY:³² The clustering of the property of producing one of the two gamete types and the other physiological features typically associated with it, although statistically disproportionate, is not strict. That is, sometimes, properties which in the vast majority of cases are associated with one of the two gametes co-occur with the other gamete. As such, there are individuals who instantiate properties of both clusters, and thus cannot be located in either of them.³³ In other words, the two groups of people that each have one gamete type and the properties typically associated with it *do not exhaust* the entire population.

³⁰ Monomorphic species do not exhibit additional morphological differences associated with the production of each of the gametes. Also, in some species called “simultaneous hermaphrodites”, the same individual organism produces both types of gametes.

³¹ There are multiple reasons that might prevent an individual organism from producing either type of gamete. To name the most common, humans do not produce gametes throughout their entire life cycle. For this reason, it would perhaps be more precise to characterize this property in dispositional terms, or, as Byrne (2018) does, speak instead of the property of “being in the ‘developmental pathway’ that produces one gamete type or the other”.

³² Later on, I will use *EXHAUSTIVITY* to refer to the negation of *NON-EXHAUSTIVITY*. That is, the view according to which sex characteristics strictly co-vary such that there are no cases that cannot be located in either of the two main clusters.

³³ Additionally, in a minority of cases, certain features that typically occur in one of two statistically dominant forms (e.g. genitalia as either a penis or a vagina; chromosomes as either XX or XY, gonads as either ovaries or testes) exhibit different forms. As such, it is impossible to associate them with one of the two gamete types (e.g. a 1 cm phallus: “too big to be a clitoris, too small to count as a penis”, XXY chromosomes, ovotestes, etc.). For a satirical depiction of the medical handling of “abnormal” phallus sizes in newborns see: <http://alicedreger.com/phallometer>.

In order to illustrate that there is, indeed, no disagreement regarding these facts, I will summarize some of the most important views of this issue. As shall be seen, none of the authors considered here disagrees with either *DIMORPHISM* or *NON-EXHAUSTIVITY*. I shall focus on two authors who argue against the view that sex is binary and two authors who defend this view, or at least a version of it.

2.2 *The contenders*

(a) *Fausto-Sterling*

Anne Fausto-Sterling, who entitles her 2018 *New York Times* op-ed “Why Sex is Not Binary”, has been arguing for several years (Fausto-Sterling 1993, 2000, 2016, 2018, 2020) against what she refers to as “absolute dimorphism” (2016: 194). That is, the view according to which each and every individual falls into either the male or the female category. She says:

On close inspection, absolute dimorphism disintegrates even at the level of basic biology. Chromosomes, hormones, the internal sex structures, the gonads and the external genitalia all vary more than most people realize. Those born outside of the Platonic dimorphic mold are called intersexuals. (Fausto-Sterling 2000: 20)

Fausto-Sterling takes absolute dimorphism to be not only false—as well-documented cases of intersex people show—but also morally problematic. This is so, she says, because assuming absolute dimorphism might lead to conceptualizing any sex ambiguity as abnormal, and consequently might motivate so-called “normalizing” surgeries on intersex infants. Fausto-Sterling says:

Complete maleness and complete femaleness represent the extreme ends of a spectrum of possible body types. That these extreme ends are the most frequent has lent credence to the idea that they are not only natural (that is, produced by nature) but normal (that is, they represent both a statistical and a social ideal). Knowledge of biological variation, however, allows us to conceptualize the less frequent middle spaces as natural, although statistically unusual. (Fausto-Sterling 2020: 81)

Notice, then, that if we contrast Fausto-Sterling’s ideas with the two relevant facts introduced above, we can safely say that her main goal consists in defending *NON-EXHAUSTIVITY* against assumptions to the contrary. That is, Fausto-Sterling stresses the non-strict co-variation of the different sex-relevant properties and, consequently, the

existence of cases that do not fall into either of the two major categories. In this sense, although Fausto-Sterling rejects, as she puts it, “absolute dimorphism”, she does not argue against the more modest *DIMORPHISM*. Quite to the contrary, she often assumes it and sometimes even explicitly endorses it. For instance, in some of her earlier work (1993) she suggests introducing three additional sex categories (‘herms’, ‘ferms’, and ‘merms’) in order to acknowledge and give visibility to those people whose bodies exhibit characteristics of both typical males and typical females.³⁴ Importantly, however, these new categories are introduced to *complement* the extant ‘male’ and ‘female’, not to replace them. Fausto-Sterling, moreover, acknowledges that the male and female categories include the vast majority of people, as she estimates the proportion of intersex people to be 1.7%.³⁵ That is, according to her own view, the remaining 98.3% of the population falls neatly into either the female or the male category. Even more straightforwardly, she acknowledges: “Our data on humans show that anatomically and physiologically, humans are almost dimorphic with regard to genitalia and chromosomes but that when one considers intersex conditions, there are infrequent intermediate states” (2016: 190).

Accordingly, Fausto-Sterling’s goal in characterizing sex as “non-binary” should not be interpreted as aiming to deny *DIMORPHISM*, but rather, as intending to emphasize *NON-EXHAUSTIVITY*. We may conclude, therefore, that Fausto-Sterling does not disagree with either of the relevant biological facts we are focusing on.³⁶

(b) *Ziemińska*

Renata Ziemińska (2020) has also challenged the “binary notion” of sex. Like Fausto-Sterling, her main objection to this notion is that it fails to acknowledge the existence of people with intersex conditions. Ziemińska argues that the binary notion is too simple to

³⁴ Fausto-Sterling (2020: 112-113) has since retreated from this specific proposal, on the basis that in attempting to legitimize other sets of genitals she ended up focusing too much on genitals as an identity-determining feature.

³⁵ The frequency of people with intersex traits is a matter of dispute. See, for instance, Sax (2002), who criticizes Fausto-Sterling’s numbers. Notice, however, that at least part of the problem is conceptual in nature, as it involves defining INTERSEX in a broader or narrower way. See footnote 42 for more on this issue.

³⁶ Fausto-Sterling has been criticized (see Stock 2021: 57) for suggesting that sex should be conceptualized as a “continuum”, which, indeed, does not seem the best terminology to convey *DIMORPHISM*. Although it is true that Fausto-Sterling uses this terminology in some of her very early work, she has since clarified that sex, as a whole, cannot be conceptualized as a continuum (Fausto-Sterling 2016: 189).

capture the diversity of sex characteristics, and that this misrepresentation is a form of epistemic injustice (Fricker 2007) that can turn into practical harm. She says:

The simple, binary female/male divide is a kind of epistemic oversimplification that ignores a great deal of empirical data about people with intersex traits. I claim that “normalizing” surgery is the expression and reinforcement of epistemic injustice, as it is a visible, physical harm made by public institutions and allowed by the law; however, the epistemic injustice exists before the surgery. Before physical violence, there is symbolic violence. It is a physical violence towards a minority group of people in the name of norms created by the dominant group. Children with intersex traits feel “the knife of the norm” on their bodies. (Ziemińska 2020: 60-61)

Again, although Ziemińska’s main goal in arguing against the binary notion is to emphasize *NON-EXHAUSTIVITY*, she does not reject *DIMORPHISM* in order to do so. Quite to the contrary, she concedes it when she says that “the idea of sex dimorphism applies to most people, but it cannot be applied to all people and serve as a criterion to divide people into two groups” (Ziemińska 2018: 179).³⁷ That is, she acknowledges that most people do fall within the male and female categories, but stresses the importance of recognizing that not everyone does. We may conclude, then, that although she challenges the binary notion, Ziemińska does so to emphasize *NON-EXHAUSTIVITY*, and not to deny *DIMORPHISM*.

Let us turn now to considering those who sympathize with the binary notion.

(c) *Stock*

In her recent book, *Material Girls*, Kathleen Stock (2021) devotes a whole chapter to arguing that “binary sex exists” (2021: 45); that is, in her words, the view that “humans are divided into females and males and that this binary division is a natural state of affairs rooted in stable biological facts” (Stock 2021: 45). In order to do so, Stock distinguishes three alternative accounts of sex—“the gamete account”, “the chromosome account”, and “the cluster account”—and confronts them with potential objections.³⁸ According to Stock, none of these objections threatens the idea of there being “two sexes”.

³⁷ In this paper, Ziemińska presents her proposal as an argument against “sex dimorphism”. Within our current framework, though, it would be more accurate to characterize her argument as challenging “*absolute* dimorphism”, in order to make clear that her main contention is directed against *NON-EXHAUSTIVITY*, and that she does not deny *DIMORPHISM* as formulated above in section 2.1.

³⁸ Most of these objections go beyond the scope of this paper and do not interfere with its main argument. As such, I won’t address them here.

It seems clear, then, that at least part of Stock’s goal when she insists on the reality of “binary sex” is to defend *DIMORPHISM*. Indeed, she further clarifies that—no matter which of the three accounts of sex one favors—“for the majority of humans, there will be a clear answer as to whether someone is male or female” (Stock 2021: 48). Notice, however, that Stock *also* concedes *NON-EXHAUSTIVITY*, as she acknowledges that “there are occasional cases of DSDs not easily characterised as either male or female” (Stock 2021: 43).³⁹ Although the idea is formulated here in epistemic terms (i.e. she seems to suggest that the difficulty involves a limitation on *knowing* whether an individual is female or male), she actually commits to the more radical view that there are cases where there is simply no unequivocal answer as to whether an individual *is* male or female. Instead, Stock acknowledges that different accounts of the sexes might result in these categories having different membership conditions. She says:

On both the chromosome and gamete accounts, ‘male pseudohermaphrodites’ are still male because they have a Y chromosome and are on a small-gamete-producing pathway, albeit disrupted and with a non-standard sexed body shape, relative to the norm. Equally, ‘female pseudohermaphrodites’ are female because they lack a Y chromosome and are on a large-gamete-producing pathway, again with the caveat above. It’s true these results are at odds with what some people with CAIS and CAH would say about themselves, but that is not necessarily a reason to reject the conclusions. On the cluster account, meanwhile, ‘male pseudohermaphrodites’ and ‘female pseudohermaphrodites’ can potentially count as male or female, depending on how we collectively decide to weight the importance of external morphology over other characteristics in the female and male clusters. (Stock 2021: 57-58)

Accordingly, I take it that Stock does not deny *NON-EXHAUSTIVITY* when she argues that “binary sex exists”. Quite to the contrary, although her goal is to stress *DIMORPHISM*, she concedes *NON-EXHAUSTIVITY* too.

(d) *Byrne*

Alex Byrne (2018) suggests that there are two interpretations of the claim “Sex is binary”: (1) “there are only two sexes”; and (2) “everyone is either female or male, and no one is

³⁹ Stock favors the term ‘DSD’ (i.e. difference/disorder in sex development) for referring to intersex traits. Notice, however, that, although there is no consensus within the intersex community on the preferred terminology, many reject the terminology of ‘DSD’ as pathologizing (see <https://interactadvocates.org/interact-statement-on-intersex-terminology>). I will use ‘intersex’ in the remainder of the paper.

both”. Byrne commits only to the former and admits that, given the existence of some “unclear cases”, it might be more problematic to insist that sex is binary in the second sense. That is, when Byrne claims that sex *is* binary, he intends to convey the view that “there are only two sexes”. Byrne considers this to be in line with how sex is spoken about by biologists, who, indeed, distinguish the sexes (across species) by the relative size of the gametes they produce: males produce relatively small gametes and females produce relatively large gametes. That being so, Byrne claims: “Since there are no species with a third intermediate gamete size, there are only two sexes.”

Now, Byrne is certainly right to take this definition of sex to be the one that biologists employ when talking about sex as a cross-species phenomenon and as a reproductive strategy. Consider for instance, philosopher of biology Paul Griffiths, who says:

It’s uncontroversial among biologists that many species have two, distinct biological sexes. They’re distinguished by the way that they package their DNA into ‘gametes’, the sex cells that merge to make a new organism. Males produce small gametes, and females produce large gametes. Male and female gametes are very different in structure, as well as in size. This is familiar from human sperm and eggs, and the same is true in worms, flies, fish, molluscs, trees, grasses and so forth. (Griffiths 2020)

Griffiths refers to this as the “biological definition of sex”, but makes an important clarification that we introduced already at the beginning, and which is worth keeping in mind: there being two sexes does not entail that every organism must be either one or the other. Griffiths clarifies:

Many people assume that *if* there are only two sexes, that means everyone must fall into one of them. But the biological definition of sex doesn’t imply that at all. As well as simultaneous hermaphrodites, which are both male and female, sequential hermaphrodites are first one sex and then the other. There are also individual organisms that are neither male nor female. The biological definition of sex is not based on an essential quality that every organism is born with, but on two distinct strategies that organisms use to propagate their genes. (Griffiths 2020)

That is, there being two sexes is, according to Griffiths, perfectly compatible with *NON-EXHAUSTIVITY*.

Now, analyzing Byrne’s own position regarding *NON-EXHAUSTIVITY* demands some caution. For, even though Byrne acknowledges that sex characteristics do not always

strictly co-vary, his view on whether there are individuals that fall beyond the male and female categories is not completely clear. Indeed, on occasions he seems to concede that there are individuals that “are outside the binary”. He says:

The existence of some unclear cases shows that it would be incautious to announce that sex (in humans) is binary. By the same token, it is equally incautious to announce that it isn't—let alone that this is an established biological fact. And even if some people are outside the binary, they are a miniscule fraction of the population, nothing like the frequently cited 1–2 percent figure, which draws on Fausto-Sterling's earlier work. (Byrne 2018)

Although Byrne disagrees with the figure presented by Fausto-Sterling concerning the incidence of intersex cases,⁴⁰ he seems to concede here that there are some people, few as they may be, who fall outside the binary. In a footnote, however, he seems to deny that there are such cases and insists that “no one clearly falls beyond [the female/male binary]”.

I submit, though, that, despite this apparent unclarity, Byrne's own view regarding the non-strict co-variation of sex characteristics commits him to *NON-EXHAUSTIVITY*, as formulated above. Indeed, Byrne concedes, as he must, that the non-strict co-variation of sex traits results in there being people who instantiate properties of both typical males and typical females. He mentions, for instance, the case of Congenital Adrenal Hyperplasia, saying:

Consider, for example, the “intersex” condition Congenital Adrenal Hyperplasia—one of many “disorders of sex development” (DSDs). XX individuals with this rare condition can have an enlarged clitoris at birth (sometimes very penis-like), due to high levels of androgen hormones in the womb. They have progressed some considerable way down the developmental pathway that produces eggs (they have the usual ovaries and fallopian tubes), and have not even started down the (male) sperm-producing pathway. They are sometimes assigned male at birth, but are usually raised as girls, and indeed many of them go on to have children. Whether they are raised as girls or boys, the scientific literature correctly classifies them as female. (Byrne 2018)

⁴⁰ Stock also disagrees with Fausto-Sterling's numbers regarding the incidence of intersex cases. According to Stock (2021: 56), only 0.018% of individuals exhibit a trait in virtue of which it is difficult to locate them in the female or male categories. Byrne suggests that the actual frequency of intersex cases is 0.015% of the population.

As can be seen, though, even while acknowledging the non-strict co-variation of sex traits, and thereby conceding *NON-EXHAUSTIVITY* (as formulated above), Byrne insists that cases such as the above do not fall outside the binary, as the “scientific literature correctly classifies them as female”. But notice, however, that what allows Byrne to insist on this point is his commitment to the biological definition of sex, which, as we saw, relies only on the relative size of gametes to individuate the sexes. What Byrne fails to acknowledge, however, is that, useful as it may be for studying sex as a cross-species phenomenon, the biological definition of sex is not the *only* (scientifically) meaningful way of defining sex. Quite to the contrary, it is very likely that different purposes will require alternative concepts of MALE and FEMALE.⁴¹

This is not an isolated view. Griffiths, for instance, who supports the biological definition of sex, seems to adopt this pluralistic approach when he says:

While the biological definition of sex is needed to understand the diversity of life, that doesn't mean it's the best definition for ensuring fair competition in sport or adequate access to healthcare. We can't expect sporting codes, medical systems and family law to adopt a definition simply because biologists find it useful. Conversely, most institutional definitions of sex break down immediately in biology, because other species contradict human assumptions about sex. (Griffiths 2020)

Even Stock, as we saw earlier, exemplifies this pluralistic attitude towards sex by putting forward *three* different and compatible accounts of sex. Interestingly, she also discusses the case of people with Congenital Adrenal Hyperplasia (CAH), and explicitly acknowledges that not all of these three accounts will result in people with CAH being classified as females. She concedes, rather, that, depending on how the relevant properties are weighted, on some of these accounts people with CAH will potentially count as females or males. Contrary to what Byrnes seems to assume, then, the literature—even that part of it which is more sympathetic to his own views—is open to alternative definitions of sex.⁴²

⁴¹ To mention a salient example, notice that many sports institutions use testosterone levels as their main criterion for eligibility in the female category. (See, for instance, the IAAF's most recent regulation: <https://www.worldathletics.org/news/press-release/eligibility-regulations-for-female-classifica>.)

⁴² Given this, we should be wary of attempts to definitively set the frequency of individuals that do not fall within the binary, or, similarly, the frequency of intersex people. Indeed, notice that ‘intersex’ is not generally defined negatively as “those individuals that do not fall within the male and female categories”. Rather, ‘intersex’ is used as an umbrella term to refer to “a variety of conditions in which a person is born with a reproductive or sexual anatomy that doesn't seem to fit the typical definitions of female or male” (see https://isna.org/faq/what_is_intersex/). As such, the dispute over the frequency of intersex is not, despite appearances to the contrary, factual either. It involves, instead a terminological and conceptual

Having made this clarification, we are now in a position to see that Byrne also concedes *NON-EXHAUSTIVITY*, and thus does not deny either of the relevant biological facts. Additionally, our foregoing discussion also serves to vindicate the point made earlier regarding how important it is that we formulate the relevant biological facts (*DIMORPHISM* and *NON-EXHAUSTIVITY*) without employing contentious terms such as ‘sex’, ‘female’ or ‘male’. As the recent debate exemplifies, although everyone is likely to agree with the relevant biological facts, a terminological dispute around these facts could very easily arise, since not everyone seems to be willing to use them in the same way. Byrne, for instance, would perhaps want—even while conceding *NON-EXHAUSTIVITY*—to advance a *terminological* view according to which the terms ‘female’ and ‘male’ should *only* be used to refer to the kinds corresponding to the relative size of the gametes, and ‘sex’ *only* to refer to what is captured by the so-called biological definition of sex. Without going deeper into this issue, however, notice that this is a terminological dispute, which is compatible with there being no disagreement regarding the relevant biological facts.

2.3 *The tension around ‘binary’*

Let us take stock. After this partial survey of the views of Fausto-Sterling, Ziemińska, Stock, and Byrne, we may conclude that the dispute over the binary nature of sex involves no factual disagreement. Hence it seems that ‘binary’ can convey different ideas when used to characterize sex. On the one hand, ‘binary’ can convey the idea that there are just two relevant kinds; that is, along the lines of *DIMORPHISM*. On the other hand, it can alternatively convey the idea that there are just two relevant groups of people; that is, along the lines of *EXHAUSTIVITY* (i.e. the negation of *NON-EXHAUSTIVITY*). Although sex is binary in the first sense, it clearly isn’t in the second.

We can see, then, how a defender of the “binary view” could characterize sex as “binary” to convey something along the lines of *DIMORPHISM* without thereby committing to *EXHAUSTIVITY*. A defender of the “non-binary view”, on the contrary, could take ‘binary’ to convey *EXHAUSTIVITY* and therefore, insist on characterizing sex as “non-

dispute regarding whether or not certain specific conditions should be counted as intersex. Bluntly put, everyone agrees that the frequency of people with late-onset CAH is 1.5% of the population; the disagreement involves deciding whether or not these cases should be counted as intersex.

binary”, precisely in order to emphasize *NON-EXHAUSTIVITY* and give visibility to those who fall outside the binary.⁴³

Now, if I am right and, contrary to appearances, there is no disagreement regarding the relevant biological facts (once appropriately formulated), a natural question arises: what are these authors disagreeing about? Interesting as it is, this is a difficult question which I won’t address here. Indeed, answering this question would require interpreting people’s specific intentions behind their characterization of sex as “binary” or “non-binary”, and the values that might be involved in that choice. Instead of attempting this seemingly difficult task, I turn now to advancing a normative claim according to which we should restrain from characterizing sex as “binary” and use the term ‘dimorphic’ instead. As I will argue, the fact that ‘binary’—on one of its plausible interpretations—conveys *EXHAUSTIVITY* is particularly undesirable, from both an epistemic and a moral standpoint. Replacing ‘binary’ with ‘dimorphic’, I’ll argue, preserves the important function that ‘binary’ fulfils, while avoiding the problems of its counterpart.

3 A normative proposal: doing without ‘binary’

By surveying part of the debate about the binary or non-binary nature of sex, we have seen that, contrary to appearances, this dispute does not involve any disagreement concerning relevant facts about human sexual biology—once appropriately formulated. This analysis, though, has exposed a tension around the term ‘binary’; a term which contenders in the debate seem to be using differently. Now, one of these usages—the one that conveys *DIMORPHISM*—is epistemically accurate and morally unproblematic. If things were to end there, then there would be no obvious reason not to characterize sex as “binary”. The other interpretation of ‘binary’, however, conveys *EXHAUSTIVITY*, which is both false and morally problematic. The falsity of *EXHAUSTIVITY* is uncontroversial. Indeed, all the authors discussed above agree—despite their discrepancies regarding the

⁴³ Both Byrne and Stock seem to be aware that ‘binary’, when applied to sex, leaves room for alternative interpretations. Byrne, for instance, considers two plausible interpretations of the claim “Sex is binary”. He says: “In this sense, sex *is* binary: there are only two sexes. However, the interpretation of ‘sex is binary’ relevant to the present debate is different: everyone is either female or male, and no one is both” (Byrne 2018). In a similar vein, Stock also distinguishes two possible meanings of ‘binary’ in relation to sex. She says: “So do what Fausto-Sterling calls ‘true hermaphrodites’ show that sex isn’t a binary? Only if ‘binary’ means that every entity in the world must clearly fall into one state or the other. Properly understood, the ‘sex binary’ requires only that the vast majority of people fall into one category or the other” (Stock 2021: 59). Although Stock suggests that only the latter is the “proper” understanding of ‘binary’, she does not provide any evidence for this.

exact numbers—on the fact that the non-strict co-variation of sex characteristics results in there being a minority of cases which do not fall neatly into the categories of female and male. Let us turn, then, to illustrating in what sense it might be *morally* problematic to convey *EXHAUSTIVITY* when characterizing sex as “binary”.

3.1 *The harm of conveying EXHAUSTIVITY: the importance of visibility for stigmatized identities*

Conveying *EXHAUSTIVITY* amounts to denying the existence of intersex people. This can be particularly harmful, as a lack of awareness of the existence of people with intersex traits can accentuate the stigma around intersex identities. In her paper “For the Sake of the Children: Destigmatizing Intersexuality”, Sharon E. Preeves explains:

[... T]he social expectation is that babies are born as one of two clearly delineated anatomical types—female or male—as ascertained by genital presentation at birth. [...] The specific response to intersexual “deviance” is so strong that we have developed institutional means of covering up or erasing the violation, so that the initial social expectation of sex binarism may be upheld. More specifically, we have created medical means of surgically and hormonally engineering bodies that adhere to a two-sex social system.

Why would a cultural institution go to such great lengths to uphold a two-sex system when there are clearly consistent exceptions to this norm? One reason is that, because intersex is incongruent with the predominant, binary understanding of sex and gender, it generates the potential for social stigma and identity confusion. (Preeves 1999: 52-53)

Along similar lines, Cheryl Chase, an intersex activist and the founder of the Intersex Society of North America, claims:

[P]oor surgical outcomes are not the only—or even the primary—reason former patients feel harmed. The primary source of harm described by former patients is not surgery per se, but the underlying attitude that intersexuality is so shameful that it must be erased before the child can have any say in what will be done to his or her body. Early surgery is one means by which that message is conveyed to parents and to intersexed children. (Chase 1999: 147)

More generally, a lot of work has been carried out in order to emphasize the crucial role of political activism and *visibility* as a mean for stigmatized individuals to gain pride in their identities. This strategy to combat stigma, which has been studied in relation to other minorities (see Anspach 1979; Becker 1981; Cass 1979), also fits the intersex narrative. Preeves, again, in sharing the results of more than 40 interviews carried out with adults born with intersex traits, reports:

The first half of most interviews was laden with tales of pain, sorrow, bewilderment, and anger; the second half encompassed accounts of empowerment, identification, and reappropriation of intersexuality as a positive aspect of self. Though their association with various intersex support and/or advocacy organizations, all participants related narratives of coping with the stigma of difference through “coming out” rather than assimilating to the norm. (Preeves 1999: 59)

Intersex groups and organizations, of course, are very aware of the importance of visibility and thus coordinate activities and events such as Intersex Awareness Day⁴⁴—recognized by the U.S Department of State and the United Nations, among other institutions—in order to address this important gap.

This hopefully makes clear how visibility and awareness are fundamental for dealing with stigmatized identities such as intersex. That being so, we are now in a better position to understand in what sense it might be, not only epistemically, but also morally problematic to characterize sex in a way that conveys *EXHAUSTIVITY*, and thereby denies the existence of intersex people.

Accordingly, I suggest that it is preferable not to characterize sex as “binary” and to characterize it as “dimorphic” instead.⁴⁵

3.2 Drop ‘binary’, use ‘dimorphic’ instead

I have argued that characterizing sex as “binary” can convey *EXHAUSTIVITY*, and therefore be both epistemically and morally problematic. That being so, I suggest that we drop ‘binary’ and characterize sex as “dimorphic” instead. The main reason in favor of this

⁴⁴ <https://interactadvocates.org/intersex-awareness-day/>

⁴⁵ Notice that, by similar reasoning, characterizing sex as “non-binary” will also be problematic. Indeed, if ‘binary’ is often used, as we have seen, to convey *DIMORPHISM*, on at least one plausible interpretation ‘non-binary’ will convey *NON-DIMORPHISM*. As this is clearly false, we have epistemic reasons to avoid characterizing sex as “non-binary”.

terminological move is that ‘dimorphic’ preserves the important function that ‘binary’ presumably fulfils (i.e. conveying *DIMORPHISM*), while avoiding the problems of its counterpart. More precisely, ‘dimorphic’ seems to allow for more nuanced qualifications that might serve to cancel the unwanted interpretation of ‘binary’ that conveys *EXHAUSTIVITY*.

Notice that although it is grammatically possible to apply some such adverbs to the term ‘binary’, there is a significant sociolinguistic difference between ‘binary’ and ‘dimorphic’ in that the latter is *de facto* often qualified in potentially beneficial ways. This, I suggest, makes ‘dimorphic’ preferable. Indeed, notice that, in contrast with ‘binary’, ‘dimorphic’ is often qualified with adverbs for greater precision. These qualifications are of two different sorts. First, biologists commonly qualify ‘dimorphic’ and ‘dimorphism’ in order to be more precise about the *degree* of dimorphism that the two morphs of a given species exhibit, that is, the level of differentiation between the females and males of a particular species. For instance, it is common to refer to certain species as exhibiting “extreme” dimorphism (Magnus 2012: 160) when their males and females are very different and have very few features in common, or “low” dimorphism (Nikitovic 2018), instead, when their males and females are very similar. Second, and more importantly for our purposes, such qualifying expressions can be used to *cancel* any interpretation of ‘dimorphic’ that might convey *EXHAUSTIVITY*. In this sense, it is usual to further qualify ‘dimorphic’ with adverbs such as ‘nearly’ (Fausto-Sterling 2016: 191), ‘strictly’ (Lebacqz 1997: 217; Fausto-Sterling 2016: 191), ‘almost’ (Fausto-Sterling 2016:190), or ‘absolute’ (Ayala & Vasilyeva 2015: 727).

By using ‘dimorphic’ in this way one can convey the biologically relevant *DIMORPHISM*, without thereby further conveying the epistemically and morally problematic *EXHAUSTIVITY*. For instance, by qualifying sex as “almost dimorphic” or “nearly dimorphic” one can still recognize the existence of two main morphs while making the crucial point that not every individual belongs to one of them.

An additional advantage is that these qualifications (‘almost’, ‘strictly’, etc.) can be put to work differently depending on the needs of each context. For instance, in a context which focuses on sex from an evolutionary stance, and which aims to study sex as a cross-species phenomenon, characterizing it simply as “dimorphic” could be appropriate.⁴⁶ In a context, however, in which sex is used as a criterion to segregate people

⁴⁶ One could worry that it is inconsistent to call sex ‘dimorphic’ and ‘almost dimorphic’ at once, since ‘x is almost F’ generally implies ‘x is not F’. Appealing to the idea of context-sensitive terms (e.g. “flat”),

into two groups—for legal or other purposes—it will be crucial to emphasize *NON-EXHAUSTIVITY*. This can be done by characterizing sex as “almost dimorphic” or even by explicitly stating that “*absolute* dimorphism is false”.⁴⁷

4 Conclusion

In this paper I have analyzed the current debate on the binary or non-binary nature of human biological sex. I have reviewed the views of some of the most prominent authors involved and have argued that none of them disagrees with the relevant facts about human sexual biology. I have thus concluded that, contrary to appearances, this dispute does not involve any factual disagreement. The second part of the paper has been devoted to arguing for a terminological normative claim regarding the use of ‘binary’ when characterizing sex. I have argued that this term can be both epistemically and morally defective and that it would accordingly be better to replace it with the less problematic ‘dimorphic’.

however, can be helpful here. Indeed, whether an imperfectly flat surface counts as “flat” arguably depends on specific contextual standards. Similarly, one can think of “dimorphic” as being context-sensitive, in which case whether sex counts as “dimorphic” or not will depend on the contextual standards (For more on context-sensitive terms see Lewis 1979). I thank John Horden for raising this worry and helping me with the solution.

⁴⁷ Let me clarify that although I take this terminological move to be potentially beneficial for increasing the visibility of intersex people, it is a rather small step towards the overarching goals of the intersex community, which go far beyond what I have addressed in this paper. For one, the problematic portrayal of the intersex community does not amount, exclusively, to a *lack* of representation. Crucially, as some of the above authors suggest, the problematic portrayal of the intersex community also derives from associating intersexuality with *abnormality*.

Conclusion

In this PhD dissertation I have explored various topics in relation with metaphysical naturalness, conceptual amelioration, and sex categories. I have done so through three papers.

In the first paper I have engaged the current natural kind literature and I have put forward an original account of natural kinds according to which the naturalness of kinds is to be identified with their degree of projectibility. As I have shown, although projectibility has traditionally been given a prominent role in the characterization of natural kinds, it has often been considered insufficient for characterizing naturalness. My proposal, instead, has departed from these theories in singling out no other additional condition for naturalness. In light of this, a distinctive characteristic of my view is that, by identifying naturalness with a gradual property such as projectibility, the notion of naturalness itself becomes one of degree. I have argued that, rather than constituting a shortcoming of the view, understanding naturalness in a gradual way not only appropriately counters the relevant notion of arbitrariness but, moreover, brings important advantages over dichotomic alternatives. Additionally, in the last part of this paper I have defended, *contra* Magnus (2012) and Spencer (2015), that projectibility is a necessary condition for naturalness.

In the second paper, I have argued against Mona Simion's (2017) Epistemic Limiting Procedure for conceptual amelioration and presented a more permissive normative framework. More precisely, I have argued that epistemic losses (i.e. losses in naturalness) do not systematically prevent amelioration but merely provide reasons against it. According to my view, thus, conceptual revisions involving epistemic losses are permissible provided that they are compensated by the non-epistemic reasons that motivate the amelioration. To illustrate and defend this proposal I have considered two specific case studies involving the amelioration of MARRIAGE and SPECIES.

In the third paper, I have analyzed the current debate on the binary or non-binary nature of human biological sex and have argued that, appearances notwithstanding, this dispute does not involve a factual disagreement. To do so, I have identified the relevant biological facts and I have formulated them without using contested terms such as 'sex', 'female', or 'male'. This has enabled me to show that, although there might be a significant terminological dispute involving the use these terms, none of the authors considered disagrees with the relevant biological facts. Additionally, in this paper I have

argued in favor of a terminological normative proposal. I have defended that, when used to characterize sex, 'binary' may be both epistemically and morally problematic. In light of this, I have suggested to avoid its usage and replace it with the less problematic 'dimorphic' instead.

References

- Ainsworth, C. 2015. "Sex Redefined." *Nature*, 518(19), pp. 288-291.
- Anspach, R. R. 1979. "From Stigma to Identity Politics: Political Activism Among the Physically Disabled and Former Mental Patients." *Social Science & Medicine*, 13A, pp. 756-773.
- Ayala, S., & Vasilyeva, N. 2015. "Extended Sex: An Account of Sex for a More Just Society." *Hypatia*, 30(4), pp. 725-742.
- Becker, G. 1981. "Coping with Stigma: Lifelong Adaptation of Deaf People." *Social Science & Medicine*, 15B, pp. 21-24.
- Bird, A. 2009. "Essences and Natural Kinds." In *The Routledge Companion to Metaphysics*, ed. R. Le Poidevin, P. Simons, A. McGonigal, & R. P. Cameron. London & New York: Routledge, pp. 497-506.
- Bonchev, D., & Kamenska, V. 1981. "Predicting the Properties of the 113-120 Transactinide Elements." *Journal of Physical Chemistry*, 85(9), pp. 1177-1186.
- Boyd, R. 1991. "Realism, Anti-Foundationalism, and the Enthusiasm for Natural Kinds." *Philosophical Studies*, 61, pp. 127-48.
- _____. 1999. "Homeostasis, Species, and Higher Taxa." In *Species: New Interdisciplinary Essays*, ed. R. A. Wilson, Cambridge, MA: MIT Press, pp. 141-185.
- Byrne, A. 2018. "Is Sex Binary?" *Arc Digital*, 2 November 2018, <https://arcdigital.media/is-sex-binary-16bec97d161e>
- Cappelen, H. 2017. *Fixing Language*. Oxford: Oxford University Press.
- Cass, V. C. 1979. "Homosexual Identity Formation: A Theoretical Model." *Journal of Homosexuality*, 4, pp. 91-104.
- Chakravartty, A. 2007. *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge: Cambridge University Press.
- _____. [forthcoming]. "Last Chance Saloons for Natural Kind Realism." *American Philosophical Quarterly*. Available at: <https://anjanchakravartty.files.wordpress.com/2020/10/kind-realism.pdf>
- Chalmers, D. 2011. "Verbal Disputes." *Philosophical Review* 120: 515-566.
- Chang, Z., Li, J., & Dong, C. 2010. "Ionization Potentials, Electron Affinities, Resonance Excitation Energies, Oscillator Strengths, and Ionic Radii of Element Uus (Z) 117 and Astatine." *Journal of Physical Chemistry*, 114, pp. 13388-13394.

- Chase, C. 1999. "Surgical Progress is Not the Answer to Intersexuality." In *Intersex in the Age of Ethics*, ed. Alice Domurat Dreger. Hagerstown, MD: University Publishing Group.
- Churchland, P. 1985. "Conceptual Progress and Word/Word relations: In Search of the Essence of Natural kinds". *Canadian Journal of Philosophy* 15 (1): 1-17
- Clark, A., and D. J. Chalmers. 1998. "The Extended Mind". *Analysis* 58 (1): 7–19.
- Conix, S. 2017. *Radical Pluralism, Ontological Underdetermination, and the Role of Values in Species Classification*. PhD dissertation, University of Cambridge.
- Devitt, M. 2008. "Resurrecting Biological Essentialism." *Philosophy of Science*, 75(3), pp. 344-382.
- Dorr, C. 2019. "Natural Properties." In *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), ed. E. N. Zalta. Available at: <https://plato.stanford.edu/archives/fall2019/entries/natural-properties/>.
- Dupré, J. 1981. "Natural Kinds and Biological Taxa." *Philosophical Review*, 90(1), pp. 66-90.
- _____. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- _____. 1999. "Are Whales Fish?" In *Folkbiology*, ed. D. L. Medin & S. Atran. Cambridge, MA: MIT Press, pp. 461-476.
- Ellis, B. 2001. *Scientific Essentialism*. Cambridge: Cambridge University Press.
- _____. 2008. "Essentialism and Natural Kinds". In *The Routledge Companion to Philosophy of Science*. Ed. by Stathis Philon and Martin Curd. New York: Routledge. pp. 139-149.
- Ereshefsky, M. 2001. *The Poverty of the Linnaean Hierarchy*. Cambridge: Cambridge University Press.
- _____. 2007. "Species, Taxonomy, and Systematics." In *Philosophy of Biology*, ed. M. Matthen & C. Stephens. Oxford: North-Holland, pp. 403-427.
- Ereshefsky, M., & Matthen, M. 2005. "Taxonomy, Polymorphism and History: An Introduction to Population Structure Theory." *Philosophy of Science*, 72(1), pp. 1-21.
- Ereshefsky, M., & Reydon, T. 2015. "Scientific Kinds." *Philosophical Studies*, 172, pp. 969-986.

- _____. [forthcoming] “The Grounded Functionality Account of Natural Kinds.” In *From Biological Practice to Scientific Metaphysics: Minnesota Studies in the Philosophy of Science*, ed. W. Bausman, J. Baxter, O. Lean, A. Love, & C. K. Waters. Minneapolis: University of Minnesota Press.
- Fausto-Sterling, A. 1993. “The Five Sexes: Why Male and Female Are Not Enough.” *Science*, 33, pp. 20-24
- _____. 2000. “The Five Sexes, Revisited.” *The Sciences*. July/August 2000.
- _____. 2016. “On the Critiques of the Concept of Sex: An Interview with Anne Fausto-Sterling.” *Journal of Feminist Cultural Studies*, 27(1), pp. 189-205.
- _____. 2018. “Why Sex Is Not Binary.” *New York Times*, 25 October 2018, <https://www.nytimes.com/2018/10/25/opinion/sex-biologybinary.html>
- _____. 2020. *Sexing the Body: Gender Politics and the Construction of Sexuality*. New York: Basic Books. (First edition 2000. New York: Basic Books.)
- Franklin-Hall, L. R. 2015. “Natural Kinds as Categorical Bottlenecks.” *Philosophical Studies*, 172, pp. 925-948.
- Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Goodman, N. 1973. *Fact, Fiction, and Forecast*, 3rd edn. Indianapolis & New York: Bobbs Merrill.
- Griffiths, P. E. 1999. “Squaring the Circle: Natural Kinds with Historical Essences.” In *Species: New Interdisciplinary Essays*, ed. R. A. Wilson. Cambridge, MA: MIT Press, pp. 209–228.
- _____. 2020. “Sex Is Real: Yes, There Are Just Two Biological Sexes. No, This Doesn’t Mean Every Living Thing Is Either One or the Other.” *Aeon*. Accessed at: <https://aeon.co/essays/the-existence-of-biological-sex-is-no-constraint-on-human-diversity>
- Hacking, I. 2007. “Natural Kinds: Rosy Dawn, Scholastic Twilight”. *Royal Institute of Philosophy Supplement* 61: 203–39.
- Häggqvist, S. 2005. “Kinds, Projectibility and Explanation.” *Croatian Journal of Philosophy*, 13, pp. 71-87.
- Hoffman, D. C., Lee, D. M., & Pershina, V. 2010. “Transactinides and the Future Elements.” In *The Chemistry of the Actinide and Transactinide*

- Elements*, 4th edn., ed. L. R. Morss, N. M. Edelstein, J. Fuger. Dordrecht: Springer.
- Ilkka, N. 2019. "Scientific Progress." In *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), ed. E. N. Zalta. Available at: <<https://plato.stanford.edu/archives/win2019/entries/scientific-progress/>>.
- Kendig, C. & Grey, J. 2019. "Can the Epistemic Value of Natural Kinds Be Explained Independently of Their Metaphysics?" *The British Journal for Philosophy of Science*, 72: 2, pp. 359-376.
- Khalidi, M. 2013. *Natural Categories and Human Kinds: Classification in the Natural and Social Sciences*. Cambridge: Cambridge University Press.
- _____. 2018. "Natural Kinds as Nodes in Causal Networks." *Synthese*, 195, pp. 1379–1396.
- _____. 2021. "Etiological Kinds." *Philosophy of Science*, 88(1), pp. 1-21.
- Kitcher, P. 1984. "Species." *Philosophy of Science* 51: 308 – 333.
- _____. 1987. Ghostly Whispers: Mayr, Ghiselin, and the "Philosophers" on the Ontological Status of Species. *Biology and Philosophy* 2: 184 – 192.
- Lebacqz, K. 1997. "Difference or Defect? Intersexuality and the Politics of Difference." *Annual of the Society of Christian Ethics*, 7, pp. 213-229.
- Lemeire, O. 2021. "The Causal Structure of Natural Kinds." *Studies in History and Philosophy of Science*, 85, pp. 200-207.
- López de Sa, D. (MS). "The Construction of the Sexes."
- Lewis, D. 1983. "New Work for a Theory of Universals". *Australasian Journal of Philosophy* 61: 343–77.
- _____. 1979. "Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1), pp. 339-359.
- Ludlow, P. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford University Press.
- Ludwig, D. 2018. "Letting Go of 'Natural Kind': Toward a Multidimensional Framework of Nonarbitrary Classification". *Philosophy of Science* 85: 31–52.
- MacLeod, M. 2014. "Following Through on Naturalistic Approaches to Natural Kinds" (book review). *Metascience*, 23, pp. 335-338.

- Magnus, P. D. 2012. *Scientific Enquiry and Natural Kinds: From Planets to Mallards*. New York: Palgrave Macmillan.
- Martínez, M. 2020. “Synergic Kinds.” *Synthese*, 197(5): 1931-1946.
- Mayr, E., and P. D. Ashlock. 1991. *Principles of Systematic Zoology*, 2nd ed. Columbus: McGraw-Hill.
- Mill, J. S. 1843/1974. *A System of Logic in The Collected Works of John Stuart Mill*, vol. VII, ed. J. M. Robson. Toronto: University of Toronto Press.
- Nikitovic, D. 2018. “Sexual Dimorphism (humans).” In *The International Encyclopedia of Biological Anthropology*, ed. W. Trevathan, M. Cartmill, D. Dufour, C. Larsen, D. O’Rourke, K. Rosenberg and K. Strier. John Wiley & Sons, Inc.
- Okasha, S. 2002. “Darwinian Metaphysics: Species and the Question of Essentialism.” *Synthese*, 131(2), pp. 191–213.
- Pietsch, T. W. 2009. *Oceanic Anglerfishes: Extraordinary Biodiversity in the Deep Sea*. Berkeley, CA: University of California Press.
- Podosky, P. 2018. “Ideology and Normativity: Constraints on Conceptual Engineering”, *Inquiry*, DOI: 10.1080/0020174X.2018.1562374
- Preeves, S. E. 1999. “For the Sake of the Children: Destigmatizing Intersexuality.” In *Intersex in the Age of Ethics*, ed. Alice Domurat Dreger. Hagerstown, MD: University Publishing Group.
- Sax, L. 2002. “How Common Is Intersex? A Response to Anne Fausto-Sterling.” *Journal of Sex Research*, 39(33), pp. 174-178.
- Schaffer, J. 2009. “On What Grounds What”. In *Metametaphysics: New Essays on the Foundations of Ontology*. Ed. by D. Chalmers, D. Manley and R. Wasserman. Oxford University Press, pp. 347–83.
- Simion, M. 2017. “The ‘Should’ in Conceptual Engineering”. *Inquiry*, 61:8, pp. 914-928. doi:10.1080/0020174X.2017.1392894.
- Slater, M. 2013. *Are Species Real? An Essay on the Metaphysics of Species*. New York: Palgrave Macmillan.
- _____. 2015. “Natural Kindness.” *British Journal of Philosophy of Science*, 66, pp. 315-411.
- Souza, S. P., & Begossi, A. 2007. “Whales, Dolphins or Fishes? The Ethnotaxonomy of Cetaceans in São Sebastião, Brazil.” *Journal of Ethnobiology and Ethnomedicine*, 3, 9.

- Spencer, Q. 2015. "Genuine Kinds and Scientific Reality." In *Natural Kinds and Classification in Scientific Practice*, ed. C. Kendig. London & New York: Routledge, pp. 157-172.
- Steel, D. 2010. "Epistemic Values and the Argument from Inductive Risk." *Philosophy of Science*, 77: 1, pp. 14-34.
- Stock, K. 2021. *Material Girls: Why Reality Matters for Feminism*. London: Fleet.
- Sun, S. 2019. "Stop Using Phony Science to Justify Transphobia." *Scientific American*, June 13, 2019. Accessed September 15 2021: <https://blogs.scientificamerican.com/voices/stop-using-phony-science-to-justify-transphobia/>
- Takahashi, N. 2002. "Boiling Points of the Superheavy Elements 117 and 118." *Journal of Radioanalytical and Nuclear Chemistry*, 251(2), pp. 299-301.
- Quine, W. V. 1969. "Natural Kinds." In *Ontological Relativity and Other Essays*. New York: Columbia University Press, pp. 114-138.
- Wilkerson, T. E. 1988. "Natural Kinds." *Philosophy*, 63(243), pp. 29-42.
- Woodward, J. 2000. "Explanation and Invariance in the Special Sciences." *British Journal for the Philosophy of Science*, 51, pp. 197-254.
- Ziemińska, R. 2018. "Beyond Dimorphism: Intersex Persons and the Continuum of Sex Characteristics." *International Journal of Gender and Women's Studies*, 6(1), pp. 176-184.
- _____. 2020. "The Epistemic Injustice Expressed in Normalizing Surgery on Children with Intersex Traits." *Diametros*, 17 (66), pp. 52-65.