

Function in Ecology: Concepts, Values, and Management

by

Katherine H. Morrow

BS Ecology, Seattle Pacific University, 2014

BA Philosophy, Seattle Pacific University, 2014

Submitted to the Graduate Faculty of the
Dietrich School of Arts and Sciences in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Katherine H. Morrow

It was defended on

April 4, 2022

and approved by

James Justus, Professor, Philosophy, Florida State University

Sandra Mitchell, Distinguished Professor, History and Philosophy of Science

Dissertation Co-Director: Colin Allen, Distinguished Professor, History and Philosophy of
Science

Dissertation Co-Director: Michael Dietrich, Professor, History and Philosophy of Science

Copyright © by Katherine H. Morrow

2022

Function in Ecology: Concepts, Values, and Management

Katherine H. Morrow, PhD

University of Pittsburgh, 2022

This dissertation develops accounts of functional explanations in ecosystem ecology. My arguments emphasize the following two themes: (1) Ecological change, whether “natural” or anthropogenic, is the status quo. Both philosophical and management work pertaining to ecosystems should reflect the ways in which ecosystems are expected to continue changing. (2) Human value judgments play important roles in ecological research, concepts, and explanations. As such, ecology presents many important and understudied cases for considering philosophical issues related to values in science and value-laden concepts.

I develop accounts of functional concepts at multiple scales. First, I develop a new causal-role theory of ecological role functions, which are the functions of species within ecosystems. Next, I characterize the concept of ecosystem functioning. I argue that conclusions within the biodiversity-ecosystem functioning research program are value laden because *ecosystem functioning* is a positive evaluative concept much like *wellbeing*. Then, I develop a constructivist account of *ecosystem health*. I argue that specific anthropocentric values, in conjunction with natural properties of ecosystems, inform contemporary judgments about the health of ecosystems. Finally, I consider the values of biodiversity at the ecosystem and global scales. I argue that global biodiversity may not have instrumental value to humans due to the lack of a causal connection between global-scale species richness and the functions of ecosystems. In view of this, I survey alternative types of value which might be attributed to global biodiversity.

Table of Contents

Preface.....	x
1.0 Introduction.....	1
2.0 A Causal-Role Account of Ecological Role Functions	7
2.1 Introduction	7
2.2 Role Functions	9
2.3 The Account	18
2.4 Aren't Role Functions Sensitive to the Environment?.....	24
2.5 Comparison with Other Theories	27
2.5.1 Selected Effect Accounts.....	27
2.5.2 The Propensity Theory	31
2.5.3 The Organizational Theory	34
2.6 Conclusion	38
3.0 Mixed Claims in the Biodiversity-Ecosystem Functioning Research Program	40
3.1 Introduction	40
3.2 Objections to Value Judgments.....	42
3.3 Mixed Claims and the Elimination Test.....	44
3.4 <i>Functioning</i> in the Biodiversity-Ecosystem Functioning Research Program	51
3.5 Does <i>Biodiversity</i> Also Generate Mixed Claims?	60
3.6 Assessing the Value Judgments.....	63
3.7 The Management Relevance of Biodiversity-Ecosystem Functioning Findings.....	66

3.8 As Biodiversity-Ecosystem Functioning Findings Become More Causally Detailed	73
4.0 A Constructivist Account of the Ecosystem Health Concept.....	76
4.1 Introduction	76
4.2 Clarifying the Question.....	78
4.3 Naturalist Views of Ecosystem Health.....	82
4.3.1 The Human Influence View	82
4.3.2 The Historical Benchmark View	85
4.3.3 Normal Function Views	90
4.3.4 Good-For-Ecosystem View.....	93
4.3.5 The Resilience View	95
4.3.6 Views with Multiple Criteria.....	98
4.4 Constructivism.....	99
4.5 Arguments for Eliminativism.....	110
4.5.1 The Ecosystem Antirealism Objection.....	110
4.5.2 The Organicism Objection	115
4.6 Arguments for Ecosystem Health	118
4.6.1 Argument from Intrinsic Value	118
4.6.2 Argument from Entrenchment	121
4.6.3 Argument for a Plurality of Conservation Goals	121
4.6.4 Argument from the Microbiome.....	122
4.7 Conclusion	123
5.0 Values of Biodiversity at the Ecosystem and Global Scales.....	125

5.1 Introduction	125
5.2 Preliminaries	128
5.3 The Ecosystem Services Argument.....	134
5.4 The Scale Problem	141
5.5 Alternative Services Arguments.....	146
5.5.1 Conjunctive Argument	147
5.5.2 Precautionary Conjunctive Argument.....	148
5.5.3 Option Value Argument	150
5.6 What is Global Biodiversity Good For?	153
5.6.1 The Tipping Point Argument.....	154
5.6.2 Future Options	158
5.6.3 Intrinsic Value	160
5.6.4 Experience Values	167
5.7 Conclusion	169
6.0 Conclusion: Major Themes and Future Directions	171
Appendix A Intrinsic Value	175
Bibliography	178

List of Tables

Table 1 Structural Comparision of C and CCl.....	23
---	-----------

List of Figures

Figure 1 A positive but saturating relationship between starting species richness and drought resistance.....	54
--	-----------

Preface

This dissertation represents the end of a process which I could not have completed without the help of many people.

First, thanks to my committee members, Sandy Mitchell and Jack Justus, and thanks especially to my co-advisors, Colin Allen and Mike Dietrich, for all of their constructive feedback and support of this project.

Many people in addition to my committee read parts of this dissertation. I especially benefitted from the 2020 Dissertation Work Group spearheaded by Dana Matthiessen and Tom Wysocki. I received helpful comments on material from Chapter 4 from the Dissertation Writers' Group organized by Colin Allen in 2022, which included Nuhu Osman Attah, Tim Elmo Feiten, Brendan Fleig-Goldstein, JP Gamboa, Seth Goldwasser, and Tom Wysocki. I also thank Marina DiMarco for feedback on material from Chapter 5.

During my time in Pittsburgh, I relied on the HPS graduate community for their advice and encouragement. Special thanks to the other members of my cohort: Michael Begun, Jennifer Whyte, and Tom Wysocki. In addition to those already named, thanks go to Siska De Baerdemaeker, Gal Ben-Porath, Katie Creel, Mahi Hardalupas, Dejan Makovec, Jacob Neal, Willy Penn, Morgan Thompson, and Zina Ward.

I am also grateful to all of the department administrators at Pittsburgh. I especially thank Juliet Nadershahi for all of her assistance with job market matters and technical support, and Jennifer Smoak for help arranging my defense and graduation paperwork.

Thanks are also due to my undergraduate professor, Pat McDonald, who set this in motion by introducing me to the philosophy of science, and for his ongoing support.

Finally, I could not have completed this program without the support of my family. I am thankful to my parents, Art and Deanna, for their understanding and encouragement throughout this lengthy process. And I am especially grateful to my brother, Jonathan. I have eaten many of Jonathan's cookies while writing and editing this dissertation. I especially thank Jonathan for the considerable time he has spent helping me to develop my academic website.

1.0 Introduction

There has recently been a surge of philosophical interest in functional concepts in ecology.¹ Both ecological research and environmental management work utilize a variety of functional concepts (Jax 2010), some of which seem to be unique to the context of ecology. Research on functions in ecology is vital for understanding our changing environments. For example, there is an experimental research program which studies how the diversity of *role functions* of species affects the *functioning* of ecosystems.² This research is thought to be a major source of evidence for how the ongoing biodiversity losses will affect our ecosystems—effects which may be salient both for ecosystem management and for human wellbeing (Díaz et al. 2019). Various questions have arisen about the interpretation and design of this type of research, questions on which philosophers might usefully comment. For instance, there are debates about whether the mentioned research program gives us trustworthy evidence about the behavior of ecosystems, due among other things to external validity concerns (discussed in Chapters 3 and 5). There are also questions about whether *ecosystem functioning* and *ecosystem health* are appropriate targets of empirical study, given that they are evaluative concepts (I argue for this in Chapters 3 and 4).

¹ McShane (2004), Jax (2005), Păslaru (2009), Jax (2010), Odenbaugh (2010), deLaplante and Picasso (2011), Nunes-Neto and El-Hani (2011), Nunes-Neto et al. (2014), Desjardins et al. (2015), Nunes-Neto et al. (2016), Dussault and Bouchard (2017), Dussault (2018), Desjardins et al. (2019), Odenbaugh (2019), Lean (2020), Millstein (2020b), E. Jones (2021), Frank (2022).

² E.g., Finegan et al. (2015), Gagic et al. (2015), Steudel et al. (2016), Mazel et al. (2018), Moi et al. (2021).

Getting a better philosophical grip on these functional concepts will allow us to better critique and evaluate ecological research; to make better-informed decisions about the management of the environment; and to better assess philosophical views about the nature of functions, values in science, and forms of scientific explanation. Philosophical discussions of the latter questions have to date considered relatively few examples from ecology.

This dissertation offers accounts of functional explanations in ecosystem ecology, with an emphasis on the following two themes. First, ecological change—both “natural” and anthropogenic—is a defining feature of the current environmental moment. On the side of ecology, research and management efforts center around trying to predict ecosystem changes and ameliorate environmental damage. In some cases, however, philosophical accounts of concepts central to ecological work have not placed due emphasis on the context of ecosystem change. For example, prior accounts of role functions have focused primarily on explaining how role functions contribute to the organization or persistence of ecosystems (Nunes-Neto et al. 2014; Dussault and Bouchard 2017). Such accounts may be extended to account for cases of ecological change. However, the account of role functions I develop (in Chapter 2) builds in ecosystem organizational changes as a central feature, which I argue results in explanatory benefits.

The second theme across the chapters of this dissertation is the extent to which anthropocentric, non-epistemic values influence how we study, conceptualize, and manage ecological systems. These applications of values have in some cases been the subject of publicized debates among ecologists³; in other cases, value judgments have been made implicitly and without recognition by ecological researchers (as I will argue in Chapter 3). The first three substantive

³ See Lackey (2001) on controversies about the ecosystem health concept.

chapters of this dissertation are concerned mainly with characterizing the presence, extent, and nature of value judgments influencing three major functional concepts in ecology: *role functions*, *ecosystem functioning*, and *ecosystem health*. Describing these value judgments in detail sets the stage for future work to criticize those values and to reflect further on whether they support sound management objectives. The final chapter engages directly with debates about our empirical and moral justifications for certain conservation measures—debates which appeal to assumptions about ecosystem functioning.

So, each chapter of this dissertation considers a different case involving functional explanation in ecology. The chapters may be read as standalone pieces, although later chapters will sometimes refer back to points established in earlier chapters.

In the Chapter 2, I develop a new account of ecological role functions, which are the functions of species within ecosystems. The account is informed by alternative regime phenomena in ecology, i.e. cases in which ecosystems undergo sudden discrete organizational changes. My account of role functions is a causal-role theory, a theory on which role functions cannot be understood to malfunction. My account includes a counterfactual sensitivity condition that tracks and explains the distinction ecologists make between functions and various behaviors that are not functions. This condition is specified in terms of regime shift phenomena. Thus, my account recognizes the widespread potential for regime shifts as a factor directly relevant to forms of functional explanation across ecology. I argue that the counterfactual sensitivity condition I develop resolves the liberality problem often attributed to causal-role theories of function, while also illuminating the explanatory breadth of role functions within ecology.

Chapter 3 considers values in the biodiversity-ecosystem functioning (BEF) research program, an ongoing experimental investigation into the relationships between changing

biodiversity levels and the behavior of ecosystems. I argue for a central role of non-epistemic value judgments within this ecological research, specifically pertaining to interpretation of the concept of *ecosystem functioning*. I argue that *functioning* is a positive evaluative concept, much like *health* or *wellbeing*. Thus, empirical claims about ecosystem functioning are value-laden. In support of this claim, I develop a new “elimination test” for the presence of certain concept-related values in scientific research findings. I argue that the value judgments made by ecologists should be more publicized and subject to discussion, but I argue that it is appropriate for ecologists to participate in making value judgments relevant to the management of ecosystems.

Chapter 4 develops a constructivist account of the ecosystem health concept, which has not received as much philosophical treatment as the human health concept. Ecosystem health has been tied to various functional and organizational properties of ecosystems—such as their trophic structure and resilience—but I argue that these natural properties cannot fully account for ecosystem health judgments. To argue for a constructivist position, I have made a comprehensive survey of potential naturalist views of ecosystem health, and argue that they each fail to capture ordinary judgments about ecosystem health. I argue that specific human values—particularly aesthetic value, authenticity, and human wellbeing concerns—help to track and explain contemporary judgments about ecosystem health, in conjunction with natural properties. I also engage with some arguments for eliminativism about the ecosystem health concept, including concerns pertaining to ecosystem realism. This chapter begins to develop my position that realist and antirealist views about ecosystems are not relevant to questions about their health, moral status, or related properties.

In Chapter 5, I develop an underexplored objection to the Ecosystem Services Argument for protecting biodiversity. This prominent argument states that global biodiversity should be

protected because it supports the functioning of ecosystems and the services (e.g., clean water, timber, food, and recreational opportunity) provided to humans by those ecosystems. The objection I develop to this argument form is a *scale problem*, which points to a scale mismatch between ecosystem services and global biodiversity. Ecosystem services are (at best) supported by local biodiversity levels, although the empirical support for this causal connection is weaker than often supposed. I argue based on both a conceptual and empirical reason that there is surprisingly no causal connection between global biodiversity losses and local biodiversity losses. As a result, the Ecosystem Services Argument is unsound. I consider other types of argument for the protection of global biodiversity—including arguments based on option value, intrinsic value, and aesthetic value—and urge that the more *prima facie* successful arguments among these should be given further consideration by environmental philosophers.

Finally, Appendix A contains an explanation of how I understand the distinction between intrinsic and instrumental value. It contains some detail on the nature of value which goes beyond that presented in the main text. The reader is encouraged to consult this Appendix while reading section 5.6.3, which argues that biodiversity may have intrinsic value of a certain type.

A word is in order about the relationships among the views developed across these chapters. First, it is intriguing that (as I will argue) role functions are relatively non-normative—they are not malfunction-permitting. In contrast, ecosystem functioning is a highly normative concept; ecologists who study ecosystem functioning estimate variables which are thought to indicate that an ecosystem is functioning *well*. The same goes for ecosystem health. Several commentators on my work have wondered how these two positions are consistent. Chapter 2 briefly addresses this question, and the accounts of ecosystem functioning and health in Chapters 3-4 will further clarify how we should think about this. The fact that we regard ecosystems'

behaviors more normatively than the behaviors of their parts is interestingly opposite of the situation in physiology, where parts have proper functions but whole systems (organisms) do not. Part of the moral is that functional concepts in ecology behave differently from those in other areas of biology. But these observations also have direct implications for how we should think about managing ecosystems. If ecosystems have “proper” ways of functioning, but their components’ functions are mere causal roles, this suggests that we ought to be managing ecosystems at the level of their functions rather than their composition (Callicott et al. 1999). This suggestion is further bolstered by my argument in Chapter 5 that species diversity as such may not have instrumental value to humans.

My arguments in this dissertation also raise some further complications for ecological management. It is often thought that managing ecosystems for their health represents a nonanthropocentric alternative to managing ecosystems for resources or processes defined in terms of human interests. But I show that our conception of ecosystem health is influenced by anthropocentric value judgments (Chapter 4). Thus, it may be difficult to disentangle efforts to promote environmental health from anthropocentric interests. This may not be a bad thing, but it certainly warrants further reflection. My hope is to stimulate further discussion of such issues.

2.0 A Causal-Role Account of Ecological Role Functions

2.1 Introduction

Role functions in ecology—the roles of species or other ecosystem components that help explain ecosystem behavior—have recently gained attention among philosophers of biology (Odenbaugh 2010; Nunes-Neto and El-Hani 2011; Saborido et al. 2011; Nunes-Neto et al. 2014; Nunes-Neto et al. 2016; Dussault and Bouchard 2017; Dussault 2018; Odenbaugh 2019; Lean 2020; Millstein 2020). In this chapter I develop a new account of ecological role functions. Because ecosystems can undergo sudden major organizational changes, or regime shifts, in response to disturbance, my account of role functions is informed by explanations having to do with alternative regime phenomena. The account thus captures the fact that species are parts of dynamical and changing ecosystems. This context has not been fully captured in prior accounts of ecological role functions, which have focused on how species’ activities⁴ contribute to the maintenance of ecosystem organization (Nunes-Neto et al. 2014) or to the tendency of ecosystem states to persist (Dussault and Bouchard 2017). Prior accounts have not considered regime changes or ecosystem degradation as central explanatory targets within ecology.

This chapter also illustrates how considering functional explanation in the context of ecology can generate new solutions to existing philosophical problems. I will show how, in this

⁴ Throughout this chapter, I use “activity” to refer to any of the things organisms do in ecosystems, such as metabolizing, decomposing, or invading. In some cases, “behavior” might be a more natural term to use; however, to avoid any ambiguity, I will reserve the term “behavior” to refer to ecosystem-level phenomena only.

context, we can solve the liberality problem for causal-role theories of function, i.e. the problem of discriminating between functions and non-functions. In the present account, I solve the liberality problem by including a counterfactual sensitivity condition, which states that role functions are much less counterfactually sensitive to changes in ecosystem organization than are activities which do not count as functions.

In the next section, I introduce role functions and explain why Robert Cummins' theory of function provides a promising partial account of role functions. Cummins' theory is partial because it does not by itself explain why ecologists distinguish functions from certain explanatorily important activities which are not functions. The distinction between functions and non-functional activities may arise uniquely within ecology, and goes beyond the conventional function versus mere effect distinction, for reasons which will be explained. I then present my account and argue that it accurately captures and explains the distinction between functions and non-functional activities. Later sections of the chapter compare my account with alternative views of role functions, including a consideration of the relationship between causal-role and normative accounts of function.

2.2 Role Functions

Role functions are the roles played by species and other relevant units in an ecosystem, including familiar activities such as pollination and browsing. Role functions are distinct from the functions or functioning of whole ecosystems.⁵

Some main goals of my account of role functions will be to (a) correctly identify role functions, while excluding activities that ecologists do not consider role functions; and (b) illuminate some of the implicit criteria by which ecologists identify role functions. Philosophical theories of function can be more faithful or revisionary in intent. I have opted for a mostly faithful-to-practice approach here for a few reasons. First, a clear understanding of how ecologists understand functions is needed before it is possible to develop criticism that usefully engages with ecological practice. Second, alternative philosophical theories of role functions seem to intend to be descriptively accurate, making accuracy a useful point of comparison. However, this does not mean that the account presented here cannot be applied prescriptively to cases where there is disagreement or unclarity about a possible function. I will leave consideration of such cases to future work.

⁵ Overviews of the several distinct functional concepts in ecology include Jax (2005), Jax (2010), Nunes-Neto et al. (2016). On occasion, the same process can be considered both a role function and an ecosystem function: for example, denitrification can be considered a function of bacteria or of a wetland ecosystem (e.g., Peralta et al. 2010). In contrast, browsing is exclusively a role function, while nitrogen cycling is exclusively an ecosystem function. The distinction is especially important for reading ecological research on how role functions affect ecosystem functioning (e.g., Gagic et al. 2015).

Later in the chapter, I will argue that my account more accurately captures ecological practice than some of the existing theories, at least for the types of example I consider. My focus here is on functions in community and ecosystem ecology at medium scales, e.g. the ecology of lakes, forests, and prairies. Of course, accuracy is not the only potential virtue of a theory of function. Some degree of inaccuracy could be intended by other theories if they are meant to provide a basis for revising usage of “function” in ecology. In my view, though, there ought to be some specific practical reason for wanting to revise term usage. I am working at primarily a descriptive level in this chapter because I am not aware of any clear *prima facie* reason to think that there is a flaw with the conception of role functions in ecology. But I do believe that further consideration of the appropriate or inappropriate uses of functional concepts in ecology will prove fruitful. Later chapters will raise further questions along these lines.

Having explained my approach, I will now present an overview of examples of role functions in ecology. Then I will consider the suggestion that Robert Cummins’ theory of function can provide the best account of role functions.

There are multiple areas of ecology that study role functions. What follows is an illustrative, but not exhaustive list. One way ecologists study the functional composition of communities is by characterizing food webs (Elton 1927), which help to explain associated phenomena such as trophic cascades and regime shifts (e.g. Wolf et al. 2007). Trophic activities including predation, grazing, parasitism, decomposition, and filter feeding are core examples of role functions. Additional examples of role functions have to do with reproductive strategies and roles pertaining to the reproduction and dispersal of other species. Examples include being a pollinator or being insect-pollinated (Ollerton 2017). A third class of examples are activities that contribute to the cycling of elements through ecosystems, including denitrification (Peralta et al.

2010) or nitrogen fixation. Additionally, species considered “ecosystem engineers” are species whose functional activities have major transformative effects on resource availability within their ecosystems (Jones et al. 1994; Coggan et al. 2018). A notable example of an engineering function is dam building by beavers.

Ecologists who study the functional composition of plant communities, or the effects of functional diversity on ecosystem properties, often study *functional types* (e.g., Saccone et al. 2017; Neuenkamp et al. 2019; Ma et al. 2020). Functional types are groups of species that exhibit relevantly similar role functions within their community for some explanatory purposes. Species belonging to the same functional type may be considered interchangeable within the system for causal or explanatory purposes with respect to a specific explanatory target. Examples of functional types from various studies include C3 and C4 grasses⁶; legumes; forbs⁷; warm-season forbs; trees; and early-succession plants (Díaz and Cabido 2001). Examples of the role functions of members of these functional groups include supporting primary production during drought (C4 grasses); introducing nitrogen into the soil via nitrogen fixation (legumes); or producing vertical structure and shade (trees). The diversity of functional types in a community or the existence of a particular functional type may help to explain ecosystem properties and processes such as biomass, nitrogen cycling, decomposition rates, invasibility, and so forth. To sum up, role functions include various activities of species which help to explain many ecological phenomena.

⁶ I.e., grasses that undergo C3 or C4 photosynthesis—different photosynthetic pathways.

⁷ Forbs are herbaceous flowering plants, excluding grasses and grass relatives. Most common garden vegetables and annual flowers are forbs.

In addition, role functions are not malfunction-permitting. Although species' activities can sometimes be undesirable—e.g., they can contribute to the degradation of ecosystems—it would be inappropriate to say that species have malfunctioned in these contexts. Ecologists do not ordinarily speak of species as having malfunctioned, and there are good reasons not to think this way. One reason is that species can act in exactly the same way in two different ecosystems; the activity can be desirable in one ecosystem and undesirable in another. Thus, undesirability of an activity is often due to the context in which it is performed. For example, nitrogen fixing plants are beneficial within many ecosystems but are undesirable in Hawai'i, where they have altered forest structure due to historically low levels of native nitrogen fixation. Nitrogen fixers which have invaded Hawai'i have harmed the native forests because of their role function of nitrogen fixing, not because of a malfunction.

Second, undesirable effects often result from changes in population size rather than from changes in the functions of individual organisms. Individual deer which contribute to overbrowsing a forest have not malfunctioned but have simply gone on performing their role of browsing. The problem is that the ecosystem contains too many deer—perhaps because there are too few predators. Thus, it is more appropriate to speak of dysfunction at the level of ecosystem structure than at the level of individual species.

Due to considerations like these, I hold that the concept of *role function* is relatively non-normative, in that it does not permit talk of malfunction or dysfunction. In contrast, I do hold that ecosystem-level concepts such as *ecosystem functioning* and *ecosystem health* are normative. We commonly think of whole ecosystems as being in either healthy or dysfunctional states, judgments which are clearly evaluative. One might wonder how a normative conception of ecosystem functioning is consistent with a causal-role view of role functions. In brief, norms enter at the level

of considering the desirability of the whole ecosystem state. A forest ecosystem with diverse flora and a regulated deer population is in a desirable, healthy state; a forest ecosystem lacking understory or tree regeneration due to an unregulated deer population is in an undesirable, dysfunctional state. The problem in the latter case is not that deer have malfunctioned with respect to their role of browsing, but that the whole trophic structure of the ecosystem is dysfunctional.

So, I will begin with the initial observation that role functions are a type of causal-role function, i.e. they are explanatory activities which do not permit malfunction. A few philosophers have previously suggested that ecological role functions can be understood roughly as Cummins' functions (Maclaurin and Sterelny 2008, chap. 6; Odenbaugh 2010, 2019). Robert Cummins' theory of function states:

(C) x functions as a ϕ in s (or: the function of x in s is to ϕ) relative to an analytical account A of s 's capacity to ψ just in case x is capable of ϕ -ing in s and A appropriately and adequately accounts for s 's capacity to ψ by, in part, appealing to the capacity of x to ϕ in s . (Cummins 1975, 762)

In plainer language, the theory states that a component of a system has a function ϕ just when the component's capacity to ϕ helps to explain a capacity of the system. Thus, for example, the capacity of plants for primary production helps to explain an ecosystem's ability to sequester carbon or generate biomass. Similarly, the capacity of sea otters to eat urchins—to engage in predation⁸—helps explain how kelp forests have been eliminated from some parts of the west coast of North America. Otters, urchins, and kelp form a classic trophic cascade. Declines in the otter population reduced their predation on urchins. An increase in urchin populations in turn eliminated

⁸ Or *depredation*, if you're a stickler.

the kelp forests through their increased browsing activity. This is an example of a regime shift, some examples of which will come up throughout the dissertation (Holling 1973; Beisner et al. 2003; Suding et al. 2004; Wolf et al. 2007). A regime shift is a sudden, discrete qualitative change in the physical and trophic properties of an ecosystem. The potential of marine kelp forests to undergo a regime shift is explained by the trophic structure of the ecosystem, which includes the trophic role functions of otters and urchins.

This would be a short chapter if Cummins' theory as stated provided a complete account of ecological role functions. However, Cummins' theory does not fully explain why ecologists consider only certain activities to be functions among those appealed to in explanations. There are many examples of activities of species which help to explain ecological outcomes and yet are not considered to be functions. Cummins' theory ties functions to explanatory aims, a program I am amenable to. The question we are left with is what explanatory aims or considerations help to delineate functions from non-functions in ecology. The theory developed in this chapter explicitly develops some criteria which track this distinction.

Some examples of activities which are explanatorily important but not functions include:

- Overbrowsing. As noted, overbrowsing is not a malfunction, and it is also not considered a function, in contrast with browsing.
- In eutrophic lakes, algae can form mats on top of lakes, covering plants and blocking sunlight. In the wake of algae blooms, the actions of decomposers can then deplete the water of oxygen, with resulting effects such as fish kills. These activities (mat-forming, oxygen depletion) are not considered functions, though they are explanatorily important.
- Bleaching is not a function of corals.

- Invading North America is not a function of dandelions.
- Eliminating urchins is not a function of otters, although eating urchins is.

Cummins' theory might seem to count all of these activities as functions, since they are capacities of organisms which can help to explain ecological outcomes.

An initial issue is that Cummins' theory is stated in terms of explaining system-level capacities. In the algae case, a relevant capacity would be the capacity of water bodies to kill their fish; in the otter and urchin case, a relevant capacity would be the potential to undergo a regime shift. Calling these “capacities” may sound like an odd way of talking, but I see nothing illegitimate in doing so, assuming “capacity” means roughly the potential to do something.

If you do not think these outcomes are appropriately characterized as capacities, then Cummins' theory simply will not apply to the cases. It will follow, however, that Cummins' theory has seriously restricted scope within ecology. Ecologists frequently appeal to role functions when explaining various kinds of ecological outcome or process which we would not ordinarily describe as a capacity—including single past events. Cummins' theory will not apply to these cases if we cannot treat diverse types of ecological outcome as capacities. In view of this, when I present my own account below, I will swap the term “behavior” for “capacity” to indicate that role functions explain various types of ecosystem process, property, outcome, or potential.

However, let us assume that we can redescribe all relevant ecosystem outcomes in terms of capacities. Then, in addition to capturing more examples of functional explanation, Cummins' theory will potentially regard various non-functional activities as functions, such as the examples described above. This creates another instance of the well-known liberality objection to Cummins' theory (Garson 2016, sec. 5.2). A defender of Cummins' theory might urge that it is not meant to be prescriptive but is only meant to capture the form of functional explanations, given scientists'

actual explanatory practices. Although this move side steps the liberality objection, it leaves us wanting an explanation of why ecologists regard only certain activities as functions. As stated, my account is meant to be a useful addition to Cummins' because it supplies this explanation.

At this point, I can address another question about the framing of my project. I am contrasting functions with non-functional activities, rather than drawing the more usual contrast between functions and mere effects or accidents. The reason for doing this is supplied by the above examples. In ecology, one can identify many activities of organisms which help to explain some system outcomes but which are not functions. So, a descriptively accurate theory should track this distinction. One can also draw a distinction between role functions and mere effects in ecology. For example, felling trees is a function of beavers, and producing gnawing sounds while doing so is a mere effect. The latter distinction is already addressed by Cummins' theory: felling trees helps to explain ecosystem capacities, while gnawing sounds do not contribute to relevant explanations. The further distinction between functions and non-functional activities may well be unique to ecology since it does not arise in the received literature on physiological functions.

An initial suggestion is that value judgments may have something to do with the example activities above not being considered functions. All of the example activities might be considered undesirable. So, you might think that an account of role functions could be given by Cummins' theory plus some value judgments (e.g., role functions are activities which fit Cummins' theory and are also desirable). However, there are some reasons not to immediately jump to value judgments to fix the liberality problem. A major reason is that not all functions are desirable or have desirable consequences. For example, some trees (such as eucalyptus) produce compounds that stimulate fire. Promoting fire is a functional role of these trees. However, forest fires are considered undesirable by many humans, and increasingly severe forest fires resulting in part from

past mismanagement can compromise forest regeneration. Even if it is undesirable for trees to contribute to severe forest fires, promoting fire is still a role function. Another example comes from invasion biology. Organisms that are successful invaders are sometimes thought to be invasive in virtue of some of their role functions. A nitrogen fixing plant which invades Hawai'i (a place with low levels of native nitrogen fixation, Vitousek and Walker 1989) still has the role function of nitrogen fixing, even though performing this function may have destructive consequences for native Hawaiian forests.

So, stipulating that role functions must be desirable activities would throw out too many actual functions in addition to the examples of non-functions. Still, I do not want to claim that values are irrelevant to thinking about role functions, even though I have characterized them as largely non-normative because they are not malfunction-permitting. One avenue for values to indirectly influence role function assignments will be discussed below.

Finally, notice there is a levels of description issue with some of the non-functions. For example, decomposition is a role function, but eliminating oxygen from lakes is not a function of decomposers, even though oxygen depletion results from rapid decomposition. Similarly, deer function as browsers when they are overbrowsing—where overbrowsing is intuitively the wrong level of description to count as a function. These distinctions might seem unprincipled or arbitrary, but they are a feature of how ecologists think about functions, so a descriptively accurate account should explain them. I will argue below that there actually is a surprisingly principled way to make these levels of description distinctions.

2.3 The Account

I will argue that the main difference between functional roles and capacities which are not functions is that functional roles exhibit less counterfactual sensitivity. Being a decomposer, pollinator, canopy-producing tree, or grazer is fairly consistent across changes to ecosystem structure or the environment. Mostly, primary producers do not change into predators when there is a change in ecosystem organization. In contrast, overbrowsing, mat forming by algae, or eliminating urchins only occurs under some ecological regimes. For instance, algae only form mats in eutrophic systems or systems where conditions permit a rapid algae bloom. Algae cannot form mats in oligotrophic (nutrient-poor, clearwater) lakes. In contrast, algae are always primary producers, regardless of their abundance or the overall ecological organization of the lake. Thus, mat forming is more counterfactually sensitive than primary production, in the sense that it is more likely to fail to occur under a range of potential changes to surrounding ecological conditions. Similarly, overbrowsing only occurs under conditions of high population density of deer or other browsers. Since the population density of browsers can be altered by manipulating ecosystem structure (e.g., by reintroducing more predators), overbrowsing is more counterfactually sensitive than browsing. This difference in counterfactual sensitivity provides the principled basis for the levels of description distinctions discussed just previously.

Notice that judgments about counterfactual sensitivity do not rely on any evaluative judgments about the desirability of the activity, but they do rely on assumptions about what range of possible ecological and environmental conditions are relevant. When (implicitly) recognizing counterfactual sensitivity, ecologists do not consider all possible environmental conditions that could be encountered by an organism. They instead consider how organisms behave under a range of salient ecological conditions—meaning conditions which are (a) considered relevant targets of

study in their area of ecology and (b) considered likely or plausible, if not actual, ecological conditions under recent past or projected future changes. However, relevant ecological states will not necessarily include the most prevalent conditions on earth, because much of earth's land surface is now covered by agricultural or urban areas which are not explanatory targets of most ecologists (aside from ecologists who specifically work on urban or agricultural ecology).

So, salient ecological conditions include actual and potential degraded or dysfunctional ecosystem states but may not include highly artificial or distantly counterfactual states. Thus, for instance, forest ecologists who are explaining forest community structure might typically ignore how organisms would behave if the forest were to convert into a shopping mall, because shopping mall ecosystems are outside the explanatory purview of most forest ecologists. Forest ecologists would also typically ignore how organisms would behave on another planet or if earth's orbit were substantially different, because these circumstances are too distant from their explanatory interests. Again, ecologists explaining the structure of temperate ecosystems do not usually consider what would happen if the region were to become re-glaciated, because this is not a likely possibility for the near future, though they might consider what would happen if the region were to become hotter, wetter or drier, or more invaded. Judgments about both (a), relevant explanatory targets, and (b), what constitute plausible ecological conditions, will both be sensitive to some non-epistemic values. In this way, my account recognizes a route for value judgments to influence ecological role function assignments.

One might wonder what happens, according to my view, as climate change results in changes to the range of "normal" ecological states on earth. My prediction is that role function attributions will change to reflect new views about what ecosystem states are typical. My view is not that role function attributions are fixed, but that they are relative to the explanatory and

counterfactual conditions, both of which can change over time due to alterations in ecological circumstances and ecologists' interests and values. Thus, an activity which is at one point considered novel or atypical may come to be described as a function at a later time.

Cummins emphasizes the explanation-relativity of functions, so what I have said so far may be consistent with Cummins' theory of function. As a pluralist about theories of function, my goal is not to defeat Cummins' theory but to provide an improved account for the context of ecology. In particular, by developing an explicit counterfactual sensitivity condition which tracks the distinction between functions and non-functions, my account avoids the liberality which is often attributed to Cummins functions. If you are of the view that Cummins can successfully respond to liberality concerns by appealing to explanatory aims, then my account can be read as an explication of the relevant explanatory circumstances in ecology.

I will call my account Cummins plus counterfactual insensitivity (CCI), since the structure is borrowed from Cummins' theory.

(CCI) x functions as a ϕ in ecosystem e (or: the function of x in e is to ϕ) just in case there are various explanatory accounts A_1, A_2, \dots, A_n of e 's behaviors⁹ $\psi_1, \psi_2, \dots, \psi_n$ across possible alternative regimes of e and x is capable of ϕ -ing in all alternative regimes in which it is present and some of A_1, A_2, \dots, A_n account for e 's behaviors by, in part, appealing to the capacity of x to ϕ in e .

⁹ I am using "behavior" as a catch-all for ecosystem phenomena which may be targets of explanation. These can include, among many other things, regime stability; changes in regime; processes such as the nitrogen cycle; population dynamics; or historical events.

The account states, in summary: Role functions are the activities of ecosystem components which are consistent across possible alternative regimes and which may be appealed to in explanations of ecosystem behaviors across those regimes. More concisely, role functions are those explanatory activities which would be counterfactually constant across (token) ecosystem reorganizations. Other kinds of activity can be explanatorily important in some contexts but are not functions. A presupposition of the account is that although role functions always explain ecosystem (or community) behaviors, ecosystem behaviors are not always explained in terms of role functions.

Here are two examples to show how this works. Consider beavers, which have the function of constructing dams. Beavers always construct dams, regardless of their abundance or the surrounding state of the ecosystem. And the fact that they construct dams helps to explain various ecosystem properties. For example, beavers can modify hydrology by slowing water flow and creating floodplains, which then support the regeneration of water-loving trees (Wolf et al. 2007). So, dam building counts as a function because it is explanatory and counterfactually constant. Deer overbrowsing does not count as a function because it cannot occur in some forest regimes (i.e. those with lower deer populations), although it does help explain ecosystem properties, such as changes in the tree community, when it occurs. Note that high deer populations and low beaver populations both have the potential to trigger a regime shift—not just hypothetically, but in real cases. The relevant difference is that beavers go on building dams regardless of the ecological regime, whereas overbrowsing is only possible within some regimes. This difference tracks the convention of considering dam building, but not overbrowsing, a function.

In practice, most ecosystems have a limited number of regimes which are of explanatory interest to ecologists—often just two or three. For example, lakes can switch between eutrophic and oligotrophic regimes (discussed further in Chapter 4), and the rivers in Yellowstone are subject

to two different hydrological regimes (those with and without beaver dams). So, in practice, role functions are those activities which occur in both or a few regimes of interest, while non-functional activities are typically ones which occur in a single regime. However, I intentionally did not stipulate in CCI that role functions occur in all possible regimes of an ecosystem. As I stated previously, ecosystems might have additional possible regimes which are not of explanatory interest within relevant areas of ecology. Function assignments should not be taken to depend on all physically or metaphysically possible ecosystem regimes, but just those which are of explanatory interest to ecologists.

The distinctive features of ecological role functions can be seen based on the differences between the accounts (Table 1). First, Cummins specifies that functions are relative to “analytical accounts,” by which he means something roughly like a mechanistic explanation.¹⁰ In my account of ecological role functions, I change “analytical account” to “explanatory account” in order to remain more clearly neutral about the form of the explanations. There has been discussion about whether and under what circumstances explanations in ecology are mechanistic (e.g., Pâslaru 2009; Raerinne 2011). My sense is that ecologists sometimes give mechanistic and often give non-mechanistic explanations,¹¹ but the CCI account is phrased to remain neutral about this.

¹⁰ Cummins may also intend to include certain types of computational explanation. In larger-scale ecology, however, computational explanations are uncommon.

¹¹ N.B. that what ecologists call “mechanisms” don’t always correspond to “mechanistic explanations” as characterized by philosophers of science.

Table 1 Structural Comparison of C and CCI

	C (Cummins 1975, 762)	CCI (ecological role functions)
1	x functions as a ϕ in s (or: the function of x in s is to ϕ)	x functions as a ϕ in ecosystem e (or: the function of x in e is to ϕ)
2	relative to an analytical account A of s 's capacity to ψ	just in case there are various explanatory accounts A_1, A_2, \dots, A_n of e 's behaviors $\psi_1, \psi_2, \dots, \psi_n$ across possible alternative regimes of e and
3	just in case x is capable of ϕ -ing in s and	x is capable of ϕ -ing in all alternative regimes in which it is present and
4	A appropriately and adequately accounts for s 's capacity to ψ by, in part, appealing to the capacity of x to ϕ in s .	some of A_1, A_2, \dots, A_n account for e 's behaviors by, in part, appealing to the capacity of x to ϕ in e .

The second main difference between the accounts is that in CCI, functions feature in multiple explanations of multiple behaviors of the ecosystem across different regimes. When applied to the examples I have discussed, Cummins' original theory cannot fully illuminate how ecologists differentiate between functions and activities which are non-functions. In contrast, CCI explains this by appealing to the counterfactual stability of certain activities. Activities which will usually count as functions under CCI include trophic activities, engineering activities, growth forms, and salient metabolic properties, since these processes are relatively counterfactually stable. This is consistent with the kinds of activities which are regarded as functions in ordinary ecological thought.

Much of this chapter argues that my account accurately captures the way role functions are assigned. However, there is an independent reason to think counterfactual stability is an important property within ecological explanations. Because ecosystems are quite organizationally flexible,

creating stable models of ecosystems relies on identifying features of the causal or organizational structure which are relatively insensitive to surrounding changes. One of the early developments in community ecology was the description of food webs (Elton 1927). Food webs are useful because their general structure can remain relatively constant within a token ecosystem, even when ecosystem's other features change substantially. The reason for this is that trophic roles of organisms are less sensitive to surrounding changes than are other ecosystem properties (e.g., nutrient levels or population sizes). So, it is unsurprising that these comparatively insensitive properties have explanatory centrality in many parts of ecology. A property or activity which changes rapidly as the environment changes can still be explanatorily important, but only under a particular circumstance; these types of activity do not have the distinction of being regarded as functions.

Although I have emphasized the relative counterfactual stability of role functions, it is not my claim that role functions are rigid. In fact, there are many types of role function which are sensitive to the environment. In the next section, I explain how environmental sensitivity is compatible with counterfactual stability. After that, I will discuss how the account compares with existing theories of role functions.

2.4 Aren't Role Functions Sensitive to the Environment?

Ecological role functions have often been characterized as highly context-dependent and sensitive to environmental conditions (Jax, 2005; Jax 2010, p. 79). The reader may wonder how CCI handles examples of functions that are flexible or environmentally variable. In this section I will further clarify what sense of counterfactual stability is intended by the account. CCI is

consistent with role functions being sensitive to spatial and temporal environmental fluctuations, as well as to changes in ecosystem structure which do not constitute a regime shift. Explaining how this works should resolve concerns about the handling of plasticity.

Environmentally sensitive role functions include functions which occur during certain times of the day or year (e.g. photosynthesis); functions which occur in response to “special” conditions (e.g. feeding only during high tide); functions which are dependent on local physical conditions such as the amount of sunlight or wind exposure (e.g. growth form in some plants); and functions which are responsive to changes in populations sizes (e.g. prey switching). Other functions may occur only during certain life history stages (e.g. feeding on different sources in different insect life stages) or certain stages of succession (e.g. some trees can only occur later in forest succession).

A key point is that an ecosystem regime is an abstract concept pertaining to the overall organizational state of the ecosystem. Ecosystems within a certain regime—such as an oligotrophic lake, or an overbrowsed forest—will exhibit both spatial variation and temporal changes such as succession and seasonal cycles. Therefore, the requirement that role functions be counterfactually insensitive to regime changes does not require that they be insensitive to spatial and temporal variations within an ecosystem.

Prey switching—preferentially hunting different prey depending on prey population densities—is one interesting example of a flexible role function. Suppose that coyotes preferentially hunt rabbits only in certain seasons or during times when rabbits are especially abundant. Suppose also that coyotes’ behavioral hunting strategies are sensitive to the local abundance of other predators, e.g. wolves. These are not problems for CCI, since the concept of an ecosystem regime accommodates both spatial and temporal fluctuations in population densities.

For *rabbit predation* to count as a role function, the requirement is that coyotes be able to exhibit rabbit-hunting activity in any salient ecosystem regime, not that they must actually hunt rabbits constantly. Put differently, hunting activity may occur in a spatially or temporally patchy manner within a given ecosystem regime (in response to external factors, such as the density of other populations). The still counts as a role function as long as its occurrence would not be disrupted by a counterfactual regime change, e.g., changing the system from a forest to a prairie.

Similar points can be made about organisms that are only periodically active, such as adult cicadas, organisms that experience dormancy, or organisms that only occur in certain stages of succession. CCI requires that their functional capacities not be disrupted by a regime shift, but it does not require that the organism be constantly active over time. Thus, deciduous trees need not photosynthesize in all seasons in order to have the function of primary production.

Another point of clarification is that CCI is primarily meant to assess functions at the level of individual token ecosystems. Species with large geographic ranges or species that migrate seasonally might play somewhat different functional roles in different parts of their ranges, in response to environmental gradients or differences in local community composition. For example, various species which are trees under some conditions will grow only into small shrubs in response to high latitude or elevation. Similarly, animals of the same species may exhibit physiological and behavioral differences in milder versus more hot and arid climates. Predators with large ranges may hunt different species in different regions. These species may then have different role functions in different token ecosystems. This accords with commonsense ways of describing role functions. For instance, you would not say that the wolves in Minnesota, USA are (function as) predators of ptarmigan or bighorn sheep, since neither ptarmigan nor bighorn sheep are present in Minnesota, although wolves do hunt these species in more northern and western parts of their

current range. Thus, functions are commonly described such that they are indexed to token ecosystems, which is consistent with the intended application of CCI.

In summary, CCI allows that species' functions are exhibited in spatially or temporally patchy ways within ecosystems, and that species' role functions can differ across ecosystems. What CCI does assert is that, to count as function, a capacity must be maintained across potential regime changes to a token ecosystem, assuming that the species is still present. This condition is met by many activities that are variable, seasonal or context-dependent, so CCI does count such activities as functions.

2.5 Comparison with Other Theories

2.5.1 Selected Effect Accounts

So far, I have considered whether ecological role functions might be Cummins functions. I argued that my account, which adds a counterfactual insensitivity condition, performs better than the original Cummins theory both in terms of accuracy and explanatory value (i.e., illuminating what considerations are relevant role function assignments).

Having laid out my account, the rest of the chapter will consider how it compares to other views about role functions. Before I turn to these alternate accounts, I should note that I endorse pluralism about function in biology. I take pluralism about function to mean that there are multiple distinct conceptions of “function” in biology and that these distinct uses may warrant separate philosophical accounts (see Allen and Neal 2020). Distinct accounts of function—assuming that they are otherwise successful based on desiderata including accuracy, explanatory value,

consistency, etc.—need not be considered competing or conflicting. For instance, biologists utilize both causal-role and etiological functional explanations in different contexts (Amundson and Lauder 1994). So, causal-role and etiological accounts of function should not be seen as inherently in conflict, but rather but rather as illuminating different features of functional explanation.

The account I give in this chapter is a broadly causal-role account of role functions, i.e. it identifies functions with some actual activities of system components which do not permit malfunction. We can now consider whether it is possible to give an etiological account of role functions. Etiological theories of function in biology are derived from the work of Larry Wright (Wright 1973). The general idea of etiological accounts is that a trait or component's function is the activity which explains why it is present in the system. In a biological context, this is usually interpreted to mean that the function is the activity or consequence for which the component was selected in recent evolutionary history. This etiological theory—also known as a selected effect account—has provided a successful account of functional explanation in several areas of biology, particularly when considering physiological functions from an evolutionary perspective (Millikan 1989; Neander 1991; Godfrey-Smith 1994; Mitchell 1995).

The standard etiological theory describes proper functions, i.e. functions which permit malfunction. For example, my heart has the proper function of circulating my blood, a function which it could (but hopefully will not) fail to perform. I have already stated that role functions do not permit malfunction. Although it is common to speak of whole ecosystems as being dysfunctional (Chapters 3-4), ecologists do not talk about individual organisms, species, or functional types as having malfunctioned—even when they act in ways that are considered detrimental to the ecosystem.

An additional problem for the selected effect account in this context is that there is not a consistent explanatory connection between which traits are adaptive and which traits are considered role functions. Because role functions are actual activities of species in ecosystems, ecologists often do not make presumptions about the evolutionary history of the traits underlying those activities. A role function can explain ecosystem behaviors regardless of its evolutionary origins; and not all adaptive traits have explanatory relevance in ecosystem ecology. So, it seems clear that the standard selected effect theory of function is not relevant to ecological role functions. This has been a matter of general agreement among philosophers of ecology (Nunes-Neto et al. 2014; Dussault and Bouchard 2017; Dussault 2018; Odenbaugh 2019).

Recently, Roberta Millstein has developed an evolutionary account of role functions that takes a slightly different approach. She argues that role functions can be understood as products of coevolution because:

what makes the functional role claim, “The blister beetle is a parasite,” true is that there was coevolution between the Mojave Desert blister beetle and the white-faced bee as well as coevolution between the Oregon blister beetle and the dune silver bee. In other words, both populations of blister beetle underwent reciprocal natural selection to become parasites to their respective hosts, underwriting the functional role claim(s)... (Millstein 2020b, 1113)

Millstein’s account of role functions is structured slightly differently from the standard selected effect account. On this view, organisms have role functions when those organism’s traits are products of the species coevolving with other species. Although ascribing a selected function to a whole organism rather than to the trait is nonstandard, I agree with Millstein that nothing clearly prevents thinking in this way, particularly because ecologists do ascribe functions to whole organisms based on those organisms’ traits.

This coevolutionary account is more plausible than the standard selected effect account for the ecological context because it explains why evolutionary history is relevant. The claim is that

evolutionary history is relevant when role functions are products of coevolution, since coevolution pertains directly to the ecological interactions between species. The history of tight ecological interactions between two species creates the selective pressure which helps to explain why certain activities are role functions. This is especially plausible for species with mutualistic relationships such as figs pollinated by wasps, or legumes and rhizobial bacteria; and for highly specialized organisms like the parasites discussed by Millstein, or organisms which rely on a specific food source like monarch caterpillars and milkweed.

The coevolutionary account may be able to illuminate some features of role functional explanation, especially for the ecology of specialist organisms and cases of coevolved mutualism. As I said previously, I do not think that causal-role and evolutionary accounts of function are in conflict with each other, to the extent that they capture different explanatory contexts. However, in the event that not all role functions are products of coevolution, my CCI account may capture a broader range of functions than does the coevolutionary account. Future work might further consider why evolutionary history is relevant to some ecological explanatory contexts and not others.¹²

¹² Malaterre et al. (2019) make the interesting suggestion that *functional response traits* can be captured by some type of evolutionary account of function, while *functional effect traits* should be captured by a causal-role account. Functional response traits are traits pertaining to an organism's responses to the environment (e.g. drought tolerance), while functional effect traits pertain to an organism's effects on ecosystem processes (e.g. nitrogen fixing). Prima facie, both drought tolerance and nitrogen fixing may be selected for *and* they both help to explain ecosystem behaviors. So, although the suggestion merits further consideration, I am doubtful that traits should be given different functional accounts just in virtue of their classification as response or effect traits.

2.5.2 The Propensity Theory

The classic source for the propensity theory of function is Bigelow and Pargetter (1987). This propensity theory is a forward-looking evolutionary theory of function. On this view, a component has a function when it increases the fitness of the individual of which it is part. This theory has recently been applied to ecology by Antoine Dussault and Frédéric Bouchard. In order to apply this to ecology, Dussault and Bouchard have argued that ecosystems are evolutionary units.

In their words, “[u]nder the ecological interpretation of the propensity theory of fitness, a trait increases the fitness of an entity insofar as it contributes to that entity’s ability to persist in the face of the challenges posed by its environment” (Dussault and Bouchard 2017, 1131). Thus, what they mean by the “fitness” of an ecosystem is its tendency to persist. Species have functions within an ecosystem on their view when those species support the ecosystem’s tendency to persist. They call this the persistence enhancing propensity (PEP) account:

(PEP) The function of x in an ecosystem E is to F if, and only if, x is capable of doing F and x ’s capacity to F contributes to E ’s propensity to persist. (Dussault and Bouchard 2017, 1122)

Dussault and Bouchard further explain that on their view, the persistence of an ecosystem has to do with its resilience rather than with stability in the sense of constancy or equilibrium. So, on their view, species have functions in ecosystems if and only if they contribute to the ecosystem’s resilience. Part of the reason the authors build resilience into their theory of role functions is to exclude “destructive” activities from counting as functions. For example, they believe that invasive species responsible for ecological degradation do not have a function of causing the degradation.

As noted, Dussault and Bouchard believe that ecosystems can be treated as evolving under this framework. However, it seems this is not a necessary interpretation of their theory of function. One could accept that species have functions when they contribute to an ecosystem's resilience and ability to persist, even if one disputes that ecosystems thereby count as evolving. So, I will discuss only the theory of function as stated and set aside the evolutionary interpretation here.

I agree that causing or maintaining degradation per se is not an ecological role function. My theory CCI also provides a natural explanation of this. Degraded states of ecosystems contrast with alternative healthy regimes. The activity "helping to maintain a degraded state" of a given ecosystem by definition fails to occur in some alternative regimes of that ecosystem, and therefore does not count as a function. However, we should be careful not to assume that species which are considered harmful, invasive, etc. do not have role functions at all. Explanations of invasion dynamics often do appeal to invasive species' functional traits. We should also not assume that all role functions are beneficial to ecosystems, as argued previously.

The major shortcoming of PEP is that it links all role functions of ecosystem components to ecosystem resilience. In practice, role functions are relevant to explaining many ecological phenomena which can have little to do with persistence or which involve the breakdown of resilience, including trophic cascades and regime changes. So, PEP does not capture the explanatory scope of role functions within ecology. A similar criticism is made by Jay Odenbaugh, who argues that this theory problematically "takes one type of effect of the constituents of ecosystems [i.e., promoting ecosystem persistence] and assumes all ecosystem [role] functions consist in just this effect" (Odenbaugh 2019, 174).

Another problem is that resilience is a high-level ecosystem property which seems to be driven by other higher-level properties, such as connectivity¹³ (Adam et al. 2011; Olds et al. 2012) and functional trait diversity (Mori et al. 2013; Sakschewski et al. 2016). In most cases, activities of individual species are explanatorily distant from ecosystem resilience. Therefore, it would be difficult to assign functions to most species on this theory. This problem has both a factual and an epistemic component. Factually, many species' functions would not be recognized by PEP if they are not explanatorily relevant to ecosystem resilience. Epistemically, even supposing that some species' particular activities do contribute in meaningful ways to ecosystem resilience, it would be extremely difficult to discover most of these functions. This would require infeasibly detailed knowledge of how individual species contribute to high-level ecosystem properties.

A final problem is that some ecosystems are apparently not very resilient in the face of climate change, including especially coral reefs and some high-latitude and high-elevation systems. Many such ecosystems are undergoing a regime shift or may be in imminent danger of collapse. Yet role functions are still assigned to coral reef and high-latitude species. Since PEP assigns role functions only within persisting, resilient ecosystem states, this feature of the theory is in tension with ecological explanatory practice. Thus, PEP is limited in scope as compared to CCI.

As suggested, these counterintuitive results are an effect of the theory having linked all role functions to a single ecosystem property, resilience. Because of this, I argued that PEP does not capture the full range of explanatory applications of role functions.

¹³ Connectivity refers to the rate at which organisms move between ecosystems.

2.5.3 The Organizational Theory

The organizational theory has been developed by a group of authors in philosophy and biology including Charbel El-Hani, Alvaro Moreno, Matteo Mossio, Nei Nunes-Neto and Cristian Saborido. The organizational theory holds that biological functions can be characterized in terms of how parts contribute to the self-maintenance of a larger organization. A statement of the theory is as follows. The statement refers to traits, but their other work clarifies that “components” could be substituted for “traits” to reflect the intended generality of the theory.

(OT) [A] trait T has a function if and only if:

C1. T contributes to the maintenance of the organization O of S;

C2. T is produced and maintained under some constraints exerted by O;

C3. S realizes organizational closure. (Saborido et al. 2011, 594; also Moreno and Mossio 2015, 73)¹⁴

Further work extends this approach to ecological systems. For example, Nunes-Neto et al. (2016) explain that in ecology, “The producers, consumers, and decomposers play functions that contribute to the self-maintenance of the ecological system and these functions correspond to their constraining action on the flow of matter and energy within the ecosystem closure of constraints” (271). Nunes-Neto et al. (2014) claim that “the ascription of function to the biodiversity or to its components (such as the traits, populations, functional groups, etc.) aims at explaining the

¹⁴ Organizational closure refers to a certain type of mutual dependence among constraints, where constraints are certain entities which influence matter/energy flow within a larger organization. The reader is referred to the cited works for a more detailed explanation of the authors’ sense of “organizational closure” and “closure of constraints.” Mossio et al. (2009) contains a prior version of the cited theory.

maintenance of ecosystem properties (nutrient cycling, primary productivity, etc.)” (125) under circumstances in which “the items of biodiversity harness (constrain) the matter in the ecosystem in a way that the coordinated action of all these items makes ecosystem properties possible.” (131) Thus, the organizational theorists focus on the roles of components in maintaining certain organizational properties of the system, where the systemic properties are also conditions for the maintenance of the components.¹⁵

So, the OT as applied to ecology can be summarized as follows. Organisms constrain the movement of matter and energy through ecosystems via their metabolic activities. These activities of organisms make possible the properties and processes of ecosystems, including their nutrient cycles. These properties in turn allow for the continued existence of the component organisms, which depend on the presence of (e.g.) certain nutrient sources. Thus, the organisms’ activities help explain the ecological organization, which in turn helps explain the existence of the organisms. On this view, role functions are those activities which are involved in constraining matter and energy flow in certain ways, as in the case of organisms’ trophic activities.

The organizational framework may be helpful for understanding the self-maintenance of biological entities like cells and organisms. However, there are some questions about how the framework applies to ecology and to the assignment of role functions.

First, this framework assumes that ecosystems can be treated as self-maintaining units along the same lines as organisms (though it need not assume that these entities realize the same

¹⁵ This kind of theory has some historical precedent. Wimsatt describes a similar view as the “maintenance” view of function which he suggests traces to thought about medicine from the nineteenth and earlier twentieth centuries (Wimsatt 1972, 9). However, to my knowledge the theory has not previously been developed for application to ecology.

degree of complexity or stability). It has been a matter of controversy for around a century whether ecological units can be treated as relatively unified, self-maintaining entities or whether they should be treated as more accidental aggregations of populations (e.g., Lean 2018). Interestingly, the organizational framework may have resources to address some of these debates:

Constraints subject to closure *constitute* the biological organisation and, accordingly, make an essential contribution to determining the identity of the system.... [O]ne may conjecture that closure in fact defines biological individuality. (Moreno and Mossio 2015, 23, emphasis in original)

Thus, on this framework, ecosystem individuality is a matter of whether component populations and abiotic structures collectively realize closure (roughly, mutual dependence). Functions are assigned within individual ecosystems in such a state.

As noted, however, it is bound to be controversial whether ecosystems should be modelled as individuals at all. More work needs to be done to see whether different ecosystems do realize closure in the sense of the organizational theory. In view of this, a theory of function which does not make ontological assumptions about ecosystems may be desirable. CCI provides an account of role functions which is ontologically lightweight by comparison. Although functions are indexed to ecosystems, CCI leaves open that ecosystem boundaries could be purely stipulative or conventional; that ecosystems may not be self-maintaining; and that ecosystems may not be “real” (meaning, roughly, practice-independent) individuals (see Garcia and Newman 2016). Without intending to endorse all of the preceding ideas, CCI is a useful theory of function given that it leaves open the possibility of multiple ontological and empirical views about the nature of the ecosystem. In this respect, it tracks explanatory practice because ecologists disagree among themselves about the nature of ecological units (Jax 2006), yet they tend to agree about the assignment of role functions.

Second, there are plenty of examples of role functions which, speaking loosely, destabilize one regime in an ecosystem and eventually stabilize another regime. Real examples arise from the introduction of nitrogen fixing plants to an area without native nitrogen fixers (Vitousek and Walker 1989); the introduction of earthworms to an area without native earthworms (Frelich et al. 2019); or the introduction of beavers to an area without native beavers (Jax 2010). In each of these cases, the presence of a new role function in the ecosystem causes major changes to matter and energy cycling within the system with cascading effects that may lead to a new regime.

The organizational framework may be able to characterize this type of change as a change in ecosystem identity, and can recognize nitrogen fixing, leaf litter consumption, and tree felling as role functions which contribute to the self-maintenance of the post-invasion ecosystem regimes. However, an issue is what the OT says about these functions prior to and during the course of a regime shift. If earthworms have recently been introduced to a part of central North America and are in the process of causing functional changes to the local ecosystem, their litter eating acts against the maintenance of the current organization. The OT framework seems to lack the resources to call this activity a function until after the ecosystem has settled into a new regime. This is counterintuitive, since litter eating is a clear example of a role function, and its status as a function should not depend on whether an ecosystem is temporally before, during, or after a regime shift. One reason for this is epistemic: we knew that litter eating is a function before we knew that introducing this function could cause a regime shift. CCI recognizes these causal roles as functions regardless whether they disrupt or stabilize the present state of an ecosystem.

Moreover, suppose that introduced earthworm populations are managed at a low level such that they do not cause a regime shift. Although the earthworm population is too small to disrupt the functional state of the ecosystem, the activities of the earthworms have the tendency or capacity

to alter the ecosystem's organizational state. In this circumstance, it appears that the OT will also not recognize the earthworms' activities as functions. I have argued that this misrepresents ordinary thinking about ecological role functions. Role function assignments are maintained across ecosystem regimes, and whether an activity is a function does not depend on whether it tends to stabilize or disrupt a current regime.

Finally, there are some general questions about the consistency of the organizational approach with explanatory aims in ecology. For instance, Moreno and Mossio (2015, ch. 3) write that the organizational approach to function can be developed to explain various features of functions including (a) the normativity of functions, (b) the fact that attributing a function to an entity explains its existence, and (c) the nature of malfunction. These features make the OT well-suited for examining physiological functions within organisms, which are proper functions that permit malfunction. However, these features are in tension with the largely non-normative understanding of role functions argued for in this chapter. The cited authors claim that the OT can provide a single unified account of biological function, but for reasons stated I am among those who think that we require multiple accounts of function to accommodate normative and non-normative functional concepts.

2.6 Conclusion

I am a pluralist about theories of function, and I believe that different theories which capture different explanatory uses of functional concepts can peacefully coexist. Nevertheless, I have argued in previous sections that my CCI account performs better than some existing accounts of role functions according to some basic criteria, such as accuracy at identifying functions and

consistency with ecological explanatory practices. In addition to achieving a level of accuracy, I have argued that my theory

- a. regards the potential for ecological change as central rather than anomalous;
- b. resolves the liberality problem facing Cummins-type theories;
- c. explains the principled distinction between functions and non-functional activities in terms of the explanatory centrality of activities which are relatively counterfactually stable across organizational changes; and
- d. does not rely on any ontological assumptions about ecosystems, while clarifying how non-epistemic judgments about ecosystem boundaries and relevant ecosystem regimes may influence role function assignments.

These points are meant to show why my theory is a useful addition to the many existing discussions of role functions.

A defender of the other theories can say that their theory is not only meant to achieve some level of accuracy, but is also meant to guide or even revise usage of “function” in ecology. While I would support such a project undertaken for epistemic or ethical reasons, it is not a project I have undertaken in this chapter. My sense is that the literature needs more reflection on reasons for revising function terminology in ecology. Some such considerations will arise in the following chapters. For now, absent a compelling reason to think there is a practical flaw with the role function concept, I offer an account which is primarily meant to faithfully capture existing usage.

3.0 Mixed Claims in the Biodiversity-Ecosystem Functioning Research Program

3.1 Introduction

The biodiversity-ecosystem functioning (BEF) research program is an ongoing experimental investigation into the relationships between changing biodiversity levels and the behavior of ecosystems. In this chapter I establish a central role of non-epistemic value¹⁶ judgments within this ecological research. The values I consider here are implicit within the concepts studied by the researchers, in contrast to values that guide the evaluation of evidence, the choice of research focus, or the formation of policies.¹⁷ Philosophers have recently argued for the importance of value-laden concepts in certain scientific practices (Alexandrova 2018; Jones 2021). Research findings about concepts which are partly evaluative in nature have been described as *mixed claims* (Alexandrova 2018).

Drawing from the work of Anna Alexandrova, I argue that the findings of the BEF research are mixed claims because *functioning* is a thick concept which evaluates some ecosystem states as better or more desirable. Thick concepts are concepts with both evaluative and descriptive content, such as *generous* or *rude* (Väyrynen 2011). I argue that the value judgments pertaining to functioning within the BEF research ought to be publicly acknowledged and discussed by stakeholders, particularly when BEF findings are being applied to guide management decisions. I also argue that it is appropriate for ecologists to participate in making value judgments due to their

¹⁶ Subsequent uses of “value” should be understood to mean “non-epistemic value.”

¹⁷ See the excellent paper by David Frank (2022) on several additional roles of values in BEF research.

expertise, and that the judgments which have been made so far are not highly controversial. Finally, however, I caution that the BEF research at present has limited relevance to ecosystem management due to issues with the external validity of the research. These limitations have been widely acknowledged by ecologists conducting the research but have sometimes been overlooked by environmental philosophers when appealing to this research to justify arguments about conservation. The latter problem will be further developed in Chapter 5.

There are two major upshots of my evaluation of the BEF program in this chapter. First, the BEF program presents a new (to philosophers) example of a case where value judgments are integral to a major scientific practice. The case gives further reason to be skeptical of a value-free ideal for science, and illustrates why it continues to be important to focus attention on the nature of value judgments being made within scientific practice.

Second, environmental philosophers should be more cautious when interpreting BEF research in support of management objectives. More attention should be paid to the value judgments being made implicitly within this research, since these have the potential to affect decision-making in unexpected ways if the values are not explicitly recognized. In addition to this, we should recognize the limitations of conventional BEF research with respect to causal interpretation.

In the following, I will first explain why one might find it concerning that BEF researchers are making implicit value judgments. I then present Alexandrova's characterization of mixed claims, and introduce a new test for determining whether a given area of science is generating mixed claims. I present an overview of the BEF research program and apply my test to show that BEF researchers are making mixed claims. Later in the chapter I evaluate these value judgments and the extent to which BEF findings are relevant to management decision-making.

3.2 Objections to Value Judgments

Many environmental thinkers assume that empirical evidence, in part coming from the biodiversity-ecosystem functioning (BEF) research program, supports a causal connection between biodiversity and ecosystem functioning. This putative causal connection is then taken to support certain conservation efforts.¹⁸ In this chapter, I address some preliminary concerns about the BEF literature. A major focus will be the concern that the BEF research is an inappropriate guide to management decision-making on account of the roles that value judgments play within this research. There are several potential concerns one might have:

- A. Ecologists (qua scientists) should not be making value judgments at all (as part of their assessment or formulation of scientific claims). Because the BEF researchers make value judgments, this research is not reliable.¹⁹
- B. Ecologists conducting the BEF research have made inappropriate value judgments. As a result, this research is not reliable.
- C. Ecologists conducting the BEF research have made value judgments unilaterally without consulting the public or stakeholders. These values must be scrutinized before the research can be considered reliable.

Concern (A) has been raised for discussion but rejected by ecologist Kurt Jax (Jax 2010, 134 ff). Jax claims (though without citation) that ecology “generally is considered as a classical

¹⁸ Arguments of this type will be directly evaluated in Chapter 5.

¹⁹ I will focus specifically on the question whether BEF results are reliable as guides to ecosystem management. For more discussion of the epistemic implications of the use of value-laden concepts in science, see Ahn (2020) and the works cited by Ahn.

value-free science by most of its practitioners” (Jax 2010, 136). This suggests that, first, some ecologists may not be aware of the extent to which they are making implicit value judgments in their research and, second, the value-free ideal remains a salient concern among ecologists. This makes it all the more important to explicitly discuss the value judgments occurring in the BEF context. However, I will mostly set aside discussion of item (A) because philosophers of science (along with Jax and other philosophically inclined ecologists) have provided compelling arguments that it is impossible and even undesirable for scientists to avoid making value judgments while conducting epistemic work.²⁰

Concerns similar to (B) and (C) have been characterized by Gillette et al. (2021) as creating a “crowding out” problem: some values are crowded out in environmental science either because the scientists neglect some potential value-related positions or because the role of values in their work is obscured. A related worry is “that inattention to the normative character of [some scientific] claims presents [the claims] *as if* the underlying normative judgments are uncontroversial or undebatable” (Gillette et al. 2021, 87, emphasis in original). Gillette et al. focus on ecosystem health among other case studies rather than ecosystem functioning, but these concerns can be applied directly to the case of research on ecosystem functioning. The cited authors conclude that “these worries motivate a lesson regarding the importance of transparency about the value-based decisions that inform practice” (2021, 88). I agree with this position: even if value judgments made by scientists turn out to be acceptable, transparency is important to prevent the possibility or appearance of crowding out some stakeholders’ values.

²⁰ Philosophical discussions of this issue include Rudner (1953), Jeffrey (1956), Douglas (2000, 2009), Winsberg (2012), Parker (2014), Ludwig (2016), and Ahn (2020).

This chapter considers in detail the case of ecosystem functioning research. The role of values in this research has not been recognized as widely as the role of values in work on ecosystem health (the distinction between ecosystem functioning and ecosystem health will be further considered in Chapter 4). In this chapter, I consider whether “crowding out” is a serious problem in the BEF research program and why more public transparency about value judgments might be important. I will argue that crowding out (in the sense of concern B) may not be a serious threat in the case of ecosystem functioning research, because the relevant value judgments are less controversial than in other areas of environmental research. However, I remain sympathetic to concern (C) and to potential negative effects of failing to publicly disclose and consider the value judgments.

The next step in my argument is to show that the BEF researchers are in fact making value judgments and to clarify the nature of these judgments. In the next section I show this by drawing from the recent work of Anna Alexandrova on the nature of “mixed claims” in science. I argue that the BEF findings are important examples of mixed claims.

3.3 Mixed Claims and the Elimination Test

Anna Alexandrova has pioneered work on the status of mixed claims in science. Mixed claims are empirical causal or correlational claims involving at least one variable which is value-laden, e.g., human *wellbeing*. Concepts like this with both evaluative and non-evaluative content are sometimes called thick concepts. To clarify the nature of mixed claims, I will summarize the way Alexandrova interprets the case of research on human wellbeing. Then I will summarize a different interpretation of the same case, which holds that wellbeing researchers are not making

mixed claims. Explaining the different interpretations of this case will help to clarify the nature of mixed claims. I will then introduce a test for whether scientists are making genuine mixed claims which will later be used to show that biodiversity-ecosystem functioning (BEF) researchers' claims about ecosystem functioning are mixed.

Alexandrova discusses the case of social scientists doing research on human wellbeing. To conduct a study, these scientists must choose one among various ways to conceptualize wellbeing, including as happiness, life satisfaction, or flourishing. Since wellbeing is an evaluative concept, the judgment is not merely a pragmatic or methodological decision about how to operationalize or measure wellbeing. On her view, the judgment also involves a commitment to a substantive metaethical view about the nature of wellbeing or the circumstances under which humans should be said to be doing well. This value judgment will suggest different ways of assessing wellbeing and will therefore potentially affect the empirical conclusions drawn by a study. The conclusions about human wellbeing from such studies are therefore mixed claims which have an underlying commitment to a certain interpretation of *wellbeing* (Alexandrova 2018).

To clarify when mixed claims are being made, I will now consider a competing interpretation of the same research program given by Uwe Peters. Peters (2020) objects to Alexandrova's view as follows:

[I]f, in their project to investigate well-being, scientists choose, say, life-satisfaction as the most plausible conception of well-being, they therewith select one phenomenon (i.e. life-satisfaction) for investigation over another one (say, happiness). Since the values at issue direct scientists' attention toward one phenomenon rather than another, they are in a straightforward way determining scientists' research agenda. Their functional role is thus not special compared to that of the values already explored in philosophy of science [i.e. agenda-setting values.] (Peters 2020, 6)

So, Peters holds that the differing conceptions of human wellbeing pick out distinct phenomena, and that when social scientists choose to investigate one of these conceptions, they are simply

choosing a research focus. By contrast, Alexandrova thinks that when social scientists pick a conception of wellbeing to investigate, they are taking a stance about the nature of wellbeing.

In principle, either Alexandrova's or Peters' interpretation could accurately describe a given field of science. The difference is important because in a Peters-style case, the scientific conclusions should not be interpreted as committed to a normative view about the nature of a concept like *wellbeing* but rather as conclusions about a more limited or circumscribed phenomenon. The main concern in this case is that scientists choose justified research targets. In an Alexandrova-style case, where the researchers are generating mixed claims, we must also be concerned about whether research teams have correctly or appropriately interpreted the target thick concept.²¹ Put differently, in a Peters-style case, the main concern about a project investigating self-reported life satisfaction is whether life satisfaction is a beneficial research target. In an Alexandrova-style case, we must also consider whether self-reported life satisfaction represents a correct (or alternatively, fit-for-purpose) way to interpret *wellbeing*. Thus, the scope for error in an Alexandrova case is greater. This is why mixed claims merit their own treatment.

I propose the following test for which interpretation better describes a given area of science. That is, the test distinguishes between a case where there are (only) agenda-setting values at work from a case where a value judgment is internal to the interpretation of thick concepts within the research. I will sometimes call the latter type of value judgment a "concept-interpretation" value judgment. Note that there are likely some agenda-setting values at work in all scientific programs, so the question is whether there are also concept-interpretation values at work. I call it the

²¹ I am hedging between "correct" and "appropriate" in order not to take a stance on whether there is a single right way to interpret *wellbeing*.

“elimination test” for whether a scientific conclusion, study or program produces genuine mixed claims. If the scientific program *fails* the elimination test then it is producing mixed claims. Note that my language of “failing” or “passing” the test is not meant as a judgment about the appropriateness of mixed claims.

To reiterate, mixed claims arise when scientific findings are stated in terms of a general concept whose interpretation requires appealing to one or more values or ethical stances. Mixed claims may arise at several different levels within a research program: for example, in the statement of a single study’s finding or in a statement summarizing the general consensus in a field. The elimination test functions in the same way regardless of the level or scope of the claim which is being tested. The general idea behind the test is that if a value-laden concept can be eliminated from the statement of a conclusion or finding without any significant loss of content, then it is possible to translate apparently value-laden claims into non-mixed claims. In this case, the research only appears to be producing mixed claims, but the relevant value judgments actually only pertain to choice of research focus. If it is *not* possible to eliminate the value-laden concept from statements of the conclusion or finding, then the research is producing genuine mixed claims and we can conclude that the researchers are committed to some concept-interpretation value judgments or ethical stances.

The way a value-laden concept might be eliminated from statements of findings is by substituting a related concept which is not a thick or evaluative concept—or at least is not subject to multiple substantively differing evaluative interpretations. In the case of human wellbeing research, for instance, we can consider whether a conclusion (e.g., *wealth does* or *does not correlate to wellbeing*) can be rewritten, replacing the thick concept “wellbeing” with a descriptive concept such as “self-reported life satisfaction” or whatever other property was measured in the

study. If this substitution can be made without distorting the authors' intended conclusion, then those researchers have not made a substantive interpretive judgment when they decided to measure self-reported life satisfaction. They merely made a choice of research focus, and their studies' conclusions can be read as conclusions about self-reported life satisfaction rather than about wellbeing per se. So in this hypothetical case, even if researchers use the word "wellbeing" when summarizing their findings, their work does not commit them to a CI judgment about the nature of wellbeing in general.

In the alternative, it may be that trying to replace "wellbeing" with "self-reported life satisfaction" results in a distortion or loss of content with respect to the conclusion intended by the researchers. This might happen, for instance, in a meta-analysis or review which summarizes the results of multiple studies, not all of which assessed wellbeing in terms of life satisfaction. In this hypothetical review, the intended conclusion is about wellbeing per se, and trying to rewrite the conclusion in terms of life satisfaction would result in a misleading statement. The review's conclusion cannot just be about life satisfaction if it incorporates evidence from some studies which measured wellbeing in terms of a different property. Thus, this review would fail the elimination test, and we would conclude that the reviewers have made a substantive value judgment about the nature of wellbeing when they decided to include studies of life satisfaction as well as some other properties in their review.²²

²² Notice that the judgment doesn't need to be that a single measure of wellbeing is the correct one; it could be the view that a few measures are good ones, or that some measures are better than others, depending on the design of the review.

Clearly, in some cases it is difficult to tell whether the research should be thought to pass or fail the elimination test. I believe this is why Alexandrova and Peters differ in their interpretation of actual wellbeing research. Because my main aim is to assess the BEF literature, I will not take a stance about the right interpretation of this case. However, I will share some further considerations that may be relevant to resolving this type of debate.

When looking at an individual study, it may be difficult to tell whether the conclusion is meant to be about an evaluative or descriptive concept. One potential indication comes from how the authors cite other studies. If the authors cite studies which measure wellbeing differently and either (a) take those studies to have evidentiary relevance to their own conclusions or (b) assert that those studies measure wellbeing incorrectly, then the authors likely intend to be studying wellbeing *per se*. It would be a misconstrual of their intent to rewrite their conclusion in terms of a non-evaluative concept. If the authors ignore studies which measure wellbeing differently or if they cite those studies and claim that they are not relevant to the present study's conclusions, then the authors intend to be making a choice-of-focus judgment only and are likely not committed to mixed claims.

At this point I can respond to two potential concerns about the elimination test. The first is whether too many instances of measurement choice will result in failing the elimination test. Any research targeting an interesting concept requires some judgments about how to operationalize and measure it. Even for apparently very simple concepts—such as the number of individuals in a population—researchers have to choose among multiple complex methods for estimating population size when “direct” counting is not feasible. If two teams of ecologists choose different methods to estimate a population size, are they now making claims which are mixed with respect to that concept?

In this case they are probably not making mixed claims because decisions about how to estimate population size often do not involve evaluative judgments about the nature of *population size*. At least within nonhuman ecology, there are no moral or interestingly evaluative positions associated with the meaning of “population size” so claims about population size are not candidates to be mixed claims to begin with. There are differing methods of estimating population size, and values might be relevant to deciding what method to use (e.g., whether the team values greater efficiency or greater accuracy), but these values do not pertain to the nature or interpretation of *population size* itself.²³

The second worry is that concepts might be “evaluative all the way down” such that it could be impossible to find a purely descriptive underlying concept in order to implement the elimination test. For example, suppose that some wellbeing researchers believe that wellbeing consists in *flourishing*, and decide to assess flourishing in their study about wellbeing. Because flourishing is also a thick concept, substituting “flourishing” for “wellbeing” will not eliminate mixed claims. Suppose that the researchers decide to assess flourishing in part by estimating the quality of a person’s relationships. *Relationship quality* too is an evaluative concept. So, it is possible that their target concept, wellbeing, does not “bottom out” in a non-evaluative concept.

To be clear, my elimination test does not presuppose that all concepts have underlying non-evaluative measures. In a case where no suitable descriptive concept can be identified, the statements will automatically fail the elimination test. However, from considering examples I do

²³ As pointed out by Sandra Mitchell, studies of human populations make judgments about who should and should not be included in a census (e.g., citizens versus noncitizens). These decisions may involve substantive value judgments. So, part of the lesson is that similar concepts can generate mixed claims in one context and not in another.

not believe that this will happen often enough to make failure of the elimination test trivial or automatic. For instance, in the following discussion, I will claim that one ecological concept (*functioning*) fails the test—resulting in mixed claims—and another concept (*biodiversity*) passes the test in one specific research context.

To make this argument, I will next present a brief overview of the BEF literature which studies the relationship between biodiversity and ecosystem functioning.

3.4 *Functioning* in the Biodiversity-Ecosystem Functioning Research Program

The biodiversity-ecosystem functioning (BEF) program arose from inherited controversies about the role of diversity in regulating community or ecosystem stability. In mid-twentieth century ecology, it was proposed that there is a positive relationship between community complexity and stability. This hypothesis seemed to be supported by observation and theoretical prediction, but some early mathematical models of communities, such as the one developed by Robert May (1972), contradicted the hypothesis.²⁴ In the 1980s, David Tilman and colleagues began long-term manipulative field experiments at the Cedar Creek Ecosystem Science Reserve (Minnesota, USA) to investigate the relationships among grassland community diversity and other variables reflecting the functioning of the ecosystem.²⁵ Additional well-known field studies were

²⁴ Detailed discussions of the diversity-stability debate include McCann (2000), Justus (2008), Jax (2010, chap. 3), deLaplante and Picasso (2011), and Frank (2022).

²⁵ Cedar Creek studies include Tilman and Downing (1994), Tilman (1996), Tilman et al. (1996), Tilman et al. (2001), Tilman et al. (2006), Hautier et al. (2015).

conducted by the BIODDEPTH project at various sites in Europe.²⁶ This research and subsequent experiments have generally reported a positive relationship between diversity and various measures of ecosystem functioning, although there exist studies which report the opposite finding (e.g., Sasaki and Lauenroth 2011).

In the classic BEF studies, ecologists manipulated the number of species present in experimental plots. Plot sizes are given in square meters, to give a sense of the spatial scale. Temporally, the studies have spanned years to decades. The earliest studies were performed on grasslands, though subsequent research has investigated other kinds of community. In these long-term field studies, ecologists have tracked various properties of the plots such as the biomass at the end of each growing season; decomposition rates; or the number of species present each year. Based on these studies, ecologists have attempted to draw conclusions about whether variations in biodiversity levels are causally connected to variations in the functioning of the ecosystem.

To show that the resulting conclusions of this research are mixed claims, I will need to show that:

(1) *Ecosystem functioning* is a thick concept.

(2) Ecologists are drawing conclusions about ecosystem functioning per se, rather than about non-evaluative measures of functioning. (Elimination test)

You may notice that the same claims can be assessed with “biodiversity” substituting for “ecosystem functioning.” Philosophers have written somewhat more on biodiversity than on ecosystem functioning, and it has been generally accepted that *biodiversity* is a complex idea which

²⁶ BIODDEPTH studies include Spehn et al. (2005), Hector et al. (2007), Hector and Bagchi (2007).

is interpreted differently in various contexts (see Maclaurin and Sterelny 2008; Justus 2011). However, in the context of the BEF literature, it is *not* the case that

(2') Ecologists are drawing conclusions about biodiversity *per se*, rather than about non-evaluative measures of biodiversity.

At least currently, many BEF researchers are drawing general conclusions about the relationships between *specific measures of biodiversity* and *ecosystem functioning per se*. Thus, although the research may involve two evaluative concepts, the conclusions are mixed claims only in virtue of the treatment of *ecosystem functioning*. To reiterate, the issue is not solely whether the concepts involved are evaluative, but whether the conclusions drawn by the research program are implicitly committed to one evaluative stance over others.

I will begin by explaining why (2) holds, and then I will return to the issue of (1), whether functioning is an evaluative concept. After that I will explain why I reject (2').

BEF researchers tend to state their conclusions in terms of ecosystem functioning, even though they never measure functioning directly (an impossible task). To illustrate, the much-cited Tilman and Downing (1994) paper reports some results of measuring plant biomass in experimental plots over several years. They specifically report how plant biomass changed after two years of drought. They find in brief that there is a positive but saturating relationship between the species richness²⁷ in plots and the amount of biomass retained following the drought (i.e., plots with the fewest species were less drought resistant; fig. 1). The authors conclude that their findings support the hypothesis that “ecosystem functioning is sensitive to biodiversity,” implicitly taking

²⁷ “Species richness” refers to the number of species present. This is the simplest way to estimate species diversity. Other diversity indices aim to balance the number of species against the differences in abundance of species.

higher drought resilience to be an indicator of functioning (Tilman and Downing 1994, 365). This move is typical throughout the BEF literature, including studies up to the present.




Figure 1 A positive but saturating relationship between starting species richness and drought resistance. Drought resistance was estimated from the ratio between the total biomass of a plot prior to the drought and the biomass after two years of drought. This shows the extent to which biomass was maintained during an extended drought, with lower numbers indicating a greater loss of biomass. Reprinted with permission from Nature Springer, *Nature*, “Biodiversity and Stability in Grasslands,” David Tilman and John A. Downing, copyright 1994.

Aside from examining the language used in publications, some observations about the research support the assessment that what BEF researchers really care about is functioning, not particular indicators of functioning. BEF researchers often construct artificial experimental

systems. For logistical reasons, these systems usually contain fewer species total than one would observe in a natural system. Clearly, nobody cares directly about the measured variables like the biomass of the simplified artificial plots. The aim is to be able to extrapolate from the experiments to properties of communities in general.

This extrapolation is only possible if ecologists assume that the features measured in different studies all get at some shared property of ecosystems (viz., functioning). This is because, depending on the study design and the type of ecosystem, ecologists measure different indicators of functioning which are difficult to compare directly. One study might measure biomass, another decomposition rates, another how species richness fluctuates over time. Notice that at least for some ecosystems, high biomass but low richness fluctuation would indicate good functioning. This is why the measured values cannot be directly compared without interpreting them qua indicators of functioning. In addition, different ecosystem types vary in their structure and behavior to an extent that makes it difficult to make detailed inferences from the behavior of one ecosystem type to another (for instance, between terrestrial and aquatic systems: Giller et al. 2004). Thus, it would be hard to make a prediction about aquatic decomposition rates directly from findings about terrestrial decomposition rates. Instead, ecologists must interpret these findings in terms of ecosystem functioning in order to generalize and compare results. The concept of *functioning* acts as an intermediary which allows specific findings about diverse ecosystems to come into contact in this literature. There are a many reviews and meta-analyses which compare studies of different community types in this manner (e.g., Raffard et al. 2019).

I have shown that the BEF research fails the elimination test: researchers are drawing conclusions about ecosystem functioning per se, rather than about specific measures of functioning. In order to conclude that BEF researchers are making mixed claims, I also need to

show that *functioning* is a thick or evaluative concept. The reason for thinking this has already been alluded to. Studies of ecosystem functioning implicitly take certain ecosystem characteristics to be good or bad for an ecosystem of a given type, much like we take certain traits to be good or bad for humans in studies of wellbeing. (There is a further question about what it means for certain properties to be good for ecosystems. I address this issue further in Chapter 4, where I argue that we should interpret such claims as reflecting anthropocentric preferences about the environment.)

As noted, Tilman and Downing assume that a functioning prairie will maintain comparatively high productivity during a drought. On this assumption, the correlation between species richness and functioning is positive. But in principle, they could have instead assumed that prairie plots which are functioning well should show a greater decrease in productivity during a drought. On this assumption, the same data would exhibit a negative relationship between species richness and functioning. This second option seems “off” because we are accustomed to think that resistance to environmental fluctuations is a feature of a functioning prairie. But this cannot be decided by any natural facts; instead, it relies on the judgment that greater drought resistance is a good, preferable, or appropriate feature of a prairie. Because resilience is intuitively a good feature of a prairie ecosystem, higher resilience is taken to be an indicator that the ecosystem is functioning well. Similarly, ecologists assume that it is good (all else being equal) for prairies to have comparatively high primary production, so they take plots with higher production to exhibit better functioning. These assumptions show that *functioning* is treated as a positive evaluative concept. This has gone unnoticed because the evaluation is widely shared and seems common-sense.

One might wonder if ecosystem functioning can be defined as resilience or stability over time, and if doing this removes the evaluative component, assuming there is a purely descriptive sense of “stability.” There are various reasons why this cannot be done. First, ecologists have

moved away from studying ecosystems in terms of their stability, for reasons which are outside the scope of this chapter to discuss in detail. Suffice it to say, BEF studies track various properties such as primary production, decomposition rates, or community composition as indicators of functioning. So ecologists do not take functioning to consist just of ecosystem stability.

Resilience is the capacity of ecosystems to bounce back after disturbance, which is distinct from stability in the sense of constancy over time. Many ecologists do believe that resilience is connected to ecosystem functioning, as discussed above. But strictly, ecosystems are not resilient tout court; rather, certain properties or states of ecosystems are more or less resilient to certain perturbations. For example, Tilman and Downing studied the extent to which prairie productivity is resilient to drought. This type of resilience is connected to the concept of a well-functioning prairie because of the presupposition that prairies should be highly productive. Put differently, the resilience of productivity is relevant because it means the prairie is more likely to remain productive in the future. This depends on the evaluative judgment that it is good for a prairie to be productive.²⁸

Another way to show that a concept is thick is to ask whether it makes sense to assign either a positive or negative valence to having the property under typical circumstances. For example, *being above average in height* is a largely descriptive concept because it can make sense to either positively or negatively evaluate being taller than average. Being tall might be beneficial for some purposes and detrimental for others. In contrast, *wellbeing* is not purely descriptive because it does not usually make sense to think that wellbeing is a detrimental property. Thinking someone is worse off because of their wellbeing seems to involve a conceptual confusion, because

²⁸ I will return to some more problems with resilience in Chapter 4.

wellbeing is inherently good.²⁹ *Functioning* is more like *wellbeing* than *being tall* here. It would be inherently confused to think that a functioning ecosystem is in a worse state in virtue of its functioning, or that an ecosystem experiencing a breakdown of functioning is thereby in a better state.³⁰

Thus, *functioning* is a positive evaluative concept much like *wellbeing*, *beauty* or *kindness*. Note that there is a standard distinction between “ecosystem functioning” and “ecosystem function.” “Ecosystem function” is purely descriptive—it refers to an ecosystem process—while “ecosystem functioning” is evaluative and describes an ecosystem whose functions are good, normal, healthy, natural, desirable, or sustainable.

Because *functioning* is a positive evaluative concept, the BEF researchers must make assumptions about what properties indicate an ecosystem is doing well—much like human wellbeing researchers must make assumptions about what properties indicate that a human is doing well. A complication about ecosystem research is that different properties (or values of variables) are taken to be relevant to functioning in different ecosystem types. Different ecosystem types will have differing characteristic nutrient levels, decomposition rates, rates of biomass production,

²⁹ Philosophers who study thick concepts have debated the extent to which the evaluative component of thick concepts is separable from the descriptive component (e.g. the work of Pekka Väyrynen (2008; 2011; 2013)). The evaluative valence of a thick concept can be context-sensitive, e.g., we can dream up circumstances where it makes sense to say that someone’s wellbeing is a bad thing (it is bad for an evil dictator to be healthy), though presumably these scenarios are not representative or typical. My argument in this paragraph skates over these nuances for the sake of expository clarity, and my claim that *functioning* is generally taken to be a good feature of ecosystems does not rest on any position about the nature of thick concepts.

³⁰ This paragraph is due to a conversation with Dana Matthiessen and Tom Wysocki.

trophic structures, seasonal cycles, and so forth. Therefore, there is not a single set of properties which can characterize functioning ecosystems in general, but functioning has to be characterized separately for each ecosystem type. These assumptions about what properties indicate that an ecosystem of a given type is functioning well are ubiquitous in the BEF research, to the extent that they rarely are marked as assumptions. Most ecologists seem to take it as obvious that decreased productivity indicates a loss of functioning in a prairie, for instance. This again contrasts with the human wellbeing case, where there is a well-developed set of philosophical positions about the nature of human wellbeing, and the positions have been subject to controversy.

The fact that there are developed philosophical positions about human wellbeing makes it comparatively easy to detect when wellbeing researchers are making substantive normative or metaethical assumptions. In contrast, it may not be as apparent that ecosystem functioning researchers are making substantive assumptions because of (a) the comparatively underdeveloped state of philosophical positions about ecosystem functioning and (b) the fact that the specific assumptions researchers have made about what properties track functioning have not usually been subject to controversy.

Ecologists recognize that properties like biomass are merely stand-ins for the functioning of the ecosystem. Various works by ecologists explicitly discuss ecosystem properties that studies might investigate as a way of assessing functioning (e.g., Giller et al. 2004; Tett et al. 2013). So, ecologists clearly recognize that they are making judgments about how to measure or operationalize *functioning*. However, it has sometimes gone unnoticed that it is not enough to choose a variable to measure; those conducting a study must also decide which values of the variable correspond to functioning well, e.g., whether higher or lower rates of a process correspond

to better or worse functioning. This latter judgment is what I claim is often made implicitly and without recognizing that it involves a commitment to normative views about ecosystems.³¹

What I have said should make it clear that *functioning* is an evaluative concept which involves judgments about what ecosystems are like when they function well. Along with the fact that ecologists aim to make general claims about the drivers of ecosystem functioning, this shows that the BEF researchers are making mixed claims à la Alexandrova.

I will now comment on why I am focusing on *functioning* rather than *biodiversity*. Then, I will return to the question whether it is appropriate for ecologists to be making the relevant evaluative judgments, and whether this is in any way a threat to the status of the research.

3.5 Does *Biodiversity* Also Generate Mixed Claims?

I will now briefly consider whether biodiversity-ecosystem functioning (BEF) findings can be considered doubly mixed in virtue of the use of the biodiversity concept. We can assume that *biodiversity* counts as an evaluative concept in the right sense, and set this issue aside. I will argue that the BEF claims do not consistently fail the elimination test with respect to the use of “biodiversity,” so they are not mixed in this respect.

It is true that some BEF papers, especially the earlier ones, phrase their findings in terms of a relationship between biodiversity (in general) and functioning. However, there has been increasing appreciation that the exact biodiversity index used can make a big predictive difference in the BEF context. More recent reviews and studies have in fact directly investigated which

³¹ This point in this paragraph has previously been made by ecologist Kurt Jax (Jax 2010, 73 ff).

biodiversity indices best predict functioning (e.g., Gagic et al. 2015). The cited authors find that indices of functional diversity³² are better predictors than indices of species diversity. They do not conclude that this undermines a relationship between diversity and functioning. Instead, they conclude that species richness is a worse candidate to explain ecosystem functional differences. This conclusion suggests that subsequent BEF studies should consider estimating functional diversity rather than species richness. Thus, BEF researchers are not aiming at conclusions which are mixed with respect to biodiversity per se, but are willing to explain differences in functioning in terms of a particular type or interpretation of biodiversity.

Recently, David Frank has come to a seemingly different conclusion about the centrality of biodiversity per se within the BEF literature. Frank focuses on a series of influential review papers which published between 1999 and 2012. Frank argues—I think persuasively—that in some of these reviews, the putative causal role of “diversity is emphasized and composition³³ deemphasized, even though the relatively significant effects of composition and functional traits are [also] explicitly recognized” (Frank 2022, 18). Frank argues that “changes in conservation ethics and politics [i.e., the buzz around biodiversity conservation] created incentives for ecologists to construct and frame scientific hypotheses, models, and results in terms of *biodiversity*” (*Ibid.*, 18, emphasis in original).

I see this as a point of difference among ecologists. Some researchers, as Frank argues, tended to emphasize the interpretation of BEF results in terms of a causal role of biodiversity per

³² Functional diversity refers to the diversity of role functions or functional types in the community. This is distinct from taxonomic diversity, i.e. the diversity of species, genera, or other taxonomic units.

³³ “Composition” effects refer to the effects of the presence of particular species, as distinguished from the effects of species diversity.

se, especially in the “classic” literature. I agree with Frank that this has in some cases been unwarranted (see my discussion in section 3.7 below and later in Chapter 5). However, there are other ecologists who are willing to state their conclusions in terms of functional traits or composition, and who have thought directly about how to select among diversity indices. This contrasts with the case of functioning where ecologists have largely converged on what they take to be the appropriate explanatory target.

Although Frank emphasizes reasons for ecologists to be concerned about the causal role of diversity *per se*, Frank’s discussion actually supports my claim that the BEF literature passes the elimination test with respect to “biodiversity.” Frank argues that the conclusions of some review papers should actually be reinterpreted as conclusions about composition and functional traits rather than about biodiversity *per se*. If this is right, it shows that “biodiversity” not only can but should be eliminated from statements of conclusions. Put differently, ecologists are speaking loosely, perhaps too loosely, when they state BEF results as pertaining to causal roles of diversity *per se*. The studies more accurately support conclusions about specific indices of functional diversity, functional traits, or species composition.

So, the situation with respect to biodiversity is more similar to the case described by Peters, although with a major difference. In the case Peters describes, he seems to assume there could be multiple legitimate ways to study human wellbeing, and the choice is just a matter of preference. In the case at hand, it is a partly empirical matter which biodiversity indices are predictive of functioning, so at some point it would become useless to continue studying the indices which are poor predictors. Nevertheless, I agree with Frank (2022) that the choice to focus on diversity or diversity indices is partially influenced by values, including concerns about biodiversity

conservation. This is an important role of values within science, but is distinct from the role of concept-interpretation values I have described with respect to the *functioning* concept.

To sum up, ecologists have had explicit discussions and done studies about what measures of biodiversity are candidates to explain observed differences in ecosystem functioning. This shows that they are not (necessarily) wedded to making general claims about biodiversity per se can explain outcomes in terms of a specific component of biodiversity. Thus, as said previously, the current literature is best read as investigating candidate relationships between specific biodiversity measures and ecosystem functioning per se.³⁴

3.6 Assessing the Value Judgments

I have now argued that when ecologists do biodiversity-ecosystem functioning (BEF) studies, they implicitly make judgments about what an ecosystem should be like when it is functioning well. Recall that there are three potential objections to the resulting mixed claims: that ecologists should not be making value judgments at all; that ecologists may have made the wrong value judgments; or that the value judgments require further public scrutiny. These concerns are particularly salient given that this research is often appealed to in justifications of management strategies. If there are concerns about the roles of values within the primary research, then we should be all the more concerned about how this research is being applied—especially given that environmental management is already an ethically fraught enterprise.

³⁴ My claims apply specifically to the BEF program. Biodiversity might generate mixed claims in other parts of biology.

I have set aside the first concern, since there are good reasons not to endorse a value-free ideal about science. So, let us examine whether ecologists are likely to have made “wrong” value judgments about what ecosystems are like when they function well.

To begin with, the sorts of properties tracked in BEF studies are comparatively simple to estimate, such as primary production (measured as dry biomass). The studies are primarily concerned with the direction or shape of the correlation between biodiversity and these properties. Thus, the judgments are relatively straightforward assumptions of the form “biomass increases with functioning in prairies” or “temporal variability of biomass decreases with functioning in prairies.” As stated previously, these assumptions are subject to widespread agreement, so much so that they have not been a matter of dispute among contemporary ecologists. The lack of controversy is noteworthy given that there have been heated debates among ecologists about many other features of the BEF research.

So, ecologists agree (and even take it as obvious) that a decrease in productivity indicates a loss of functioning in many terrestrial systems. A question is whether other stakeholders might reasonably want to dispute this kind of judgment. If so, then we should be concerned about ecologists’ judgments crowding out public preferences. The nature of public preferences is an empirical question. However, I believe it is plausible that most people would find ecologists’ consensus judgments acceptable, to the extent of their familiarity with the relevant ecosystem types.

Some might prefer a model where all value judgments are outsourced to the public, but particularly in this context, it is appropriate for ecologists to participate in making value judgments relevant to ecosystem management. Ecologists can make important contributions to judgments about ecosystems because of their expertise. Ecologists have the best familiarity with the

organization and causal structure of ecosystems, and therefore they can be best positioned to make more detailed judgments about what properties of the ecosystem are linked to undesirable changes in functioning. In addition, many ecologists are stakeholders because they care about the conservation of ecosystems which they study, and rely on those ecosystems professionally as sources of their scientific data as well as personally as sources of recreation, nature experience, and so forth. This precludes any suggestion that value judgments ought to be outsourced entirely, or that it is inappropriate for ecologists to participate in making value judgments. At the same time, this does not mean that the judgments which have so far been made implicitly should not be scrutinized or subject to public comment.

For one thing, there may be harms or lost benefits associated with failing to involve diverse parties in discussing value judgments. For instance, involving the public in the discussion could have associated benefits (such as increasing participation in, knowledge of, or trust in the scientific process) even if the resulting value judgments were not altered at all. These benefits are lost when scientists make the value judgments without consulting the public.

It is especially important that the presence of value judgments within this research be made explicit and discussed in a public manner where implications for management are under consideration. Any discovered correlation between biodiversity and ecosystem functioning is not merely an empirical finding but also depends on the assumption that ecosystems with certain properties are in a good or preferable state.

In the following section, I directly assess the management relevance of current BEF findings and conclude that the findings have less relevance than is generally assumed. However, BEF researchers have been actively improving study design in an effort to make findings more suited to ecosystem management. In the final section I will argue that as BEF findings become

more causally detailed, BEF researchers will need to be more cautious about the value judgments they are making, and should focus more effort on consulting relevant experts and interested members of the public prior to designing studies.

3.7 The Management Relevance of Biodiversity-Ecosystem Functioning Findings

Someone might respond to the preceding discussion by arguing that the biodiversity-ecosystem functioning (BEF) findings currently have limited relevance to ecosystem management. If this is the case, then the value judgments made within the literature are fairly isolated, and we need not be concerned that they will affect management decisions in a way that runs contrary to public preferences. This might be taken to contrast with the case of climate change, where the results of climate models have direct implications for climate policy, so we ought to be very concerned about the value structures implicit in climate model design. There are three points to make in response to this.

First, in my view, it is correct to note that the BEF findings so far have limited management relevance.³⁵ I will explain why in what follows, since this point has not always been fully appreciated. Second, however, since the BEF research has been taken to support empirical generalizations with management implications, it is relevant to worry about the value judgments affecting these putative findings. Third, it is reasonable to expect that BEF findings will become more relevant to management as the research becomes more causally sophisticated, which provides

³⁵ Similar conclusions have been drawn by Newman, Varner and Lindquist (2017, chap. 2). My assessment of the BEF literature was arrived at independently from, but generally agrees with, that of the cited authors.

further reason to want to discuss the value judgments. This section will explain why the BEF research cannot closely guide ecosystem management for empirical and inferential reasons. The next section explains why the implicit value judgments will still be a matter of concern as the research moves forward.

Recognizing current limits of scientific understanding is important so that environmental philosophers do not come to hasty, incorrect, or overly simple conclusions about how the environment ought to be managed to promote certain goals. It is also important to recognize deficits in our knowledge so that we can work to fill them in. For instance, it is well-understood that greenhouse gas emission will contribute to hotter temperatures on average and to various other large-scale changes in climate patterns, and it is well-understood that climate change and human land uses are contributing to an accelerated rate of global species losses, but there is not a consensus understanding of how modest global species losses will generally affect the behavior of ecosystems (a position I argue for further in Chapter 5). The latter fact has various implications about future research. A positive implication is that further research about the connection between species losses and ecosystem functions ought to be a high priority, both among scientists and philosophers of science. A negative implication is that environmental philosophers cannot argue purely on the grounds of empirical findings that stopping global species losses ought to be a priority based on the assumption that modest species losses are going to cause a loss of ecosystem functions.³⁶ This very common argument will be examined in detail in Chapter 5.

³⁶ Arguments similar to this can be made but must utilize a precautionary principle—species losses *might* cause ecosystem collapse and this possibility ought to be avoided.

For now I will show that the classic BEF findings cannot guide management decisions because it is currently poorly understood how to generalize from the experimental findings to the behavior of real ecosystems. Put concisely, there is an external validity issue. The BEF research program is still evolving and has so far generated various important experimental and conceptual innovations. BEF researchers usually do not purport to have established claims about the behaviors of entire ecosystems. So, my discussion here is a caution to environmental thinkers who sometimes appeal to BEF findings (or to general causal assumptions they take to be supported by BEF findings) to support positions about environmental management.

First, BEF experiments tend to focus on comparatively small scales and on single trophic levels—most commonly on terrestrial primary producers—which prevents the generalization of these results to real, multi-trophic ecosystems (Hooper et al. 2005; Brose and Hillebrand 2016). More recently there has been a push to study more realistic systems or systems with multiple trophic levels in experiments (Little and Altermatt 2018; Moi et al. 2021), but it is also still common for studies to focus on a single level. Because the classic experiments exclude important factors such as herbivory, soil microbes, longer-timescale disturbance regimes, or realistic community assembly processes, they cannot predict how a real ecosystem would respond to the sorts of changes investigated.³⁷

Second, BEF findings have been challenging to interpret in terms of mechanisms (see Huston and McBride 2002; Frank 2022). This means that it is difficult to infer specific management implications from the experimental findings. A relatively mundane point about this

³⁷ “Exclude” can mean either that the factor is eliminated, that some aspects of study design attempt to control for the factor, or that the factor is not considered at all as part of the study.

research is that it “directly” finds correlations between variables, and causal relationships between the variables must then be inferred. This is particularly difficult because there is a large number of hypothesized mechanisms which might account for similar observed correlations in experiments (see Downing et al. 2014). Some of these potential mechanisms do not vindicate a causal relationship between biodiversity and functioning. To illustrate, the *sampling effect* mechanism assumes that a small number of species have a disproportionate effect on overall functioning. Experimental plots with few plant species are less likely to have one of the highly productive species, and therefore likelier to have lower productivity, merely as a result of the species present.³⁸ If this mechanism is at work, then an observed positive correlation between biodiversity levels and functioning would be an artifact of the experimental design (Huston et al. 2000; see also Hooper et al. 2005; Frank 2022).

To be clear, the sampling effect is just one of many proposed mechanisms. Other mechanisms involve a causal role for some form of diversity. For example, one hypothesis states that in systems with high functional diversity, different functional types will respond asynchronously to environmental fluctuations, which stabilizes properties of the community. Research attempting to disentangle which mechanisms are important in which systems is ongoing (e.g., Hector et al. 2010). But there currently appears to be no consensus about the relative

³⁸ In part for logistical reasons, BEF experiments tend to include a large number of plots with few total species (e.g., fewer than 10), including many with just one or two species. Presumably the most impoverished plots (which consistently exhibit a huge amount of inter-plot variance) are heavily influenced by the sampling effect, regardless whether a sampling effect is important in the surrounding ecosystem.

importance of the many proposed mechanisms.³⁹ It is likely that different mechanisms are more important in different types of communities (Hallett et al. 2014). This means that even if we have good evidence about the mechanisms at work in grasslands, it would be difficult to guess from this research what mechanisms are at work in coral reefs or boreal forests. Recall further that even if we have good evidence about the mechanisms at work in an experimental system, this may not constitute good evidence about the causal structure of a real multi-trophic ecosystem.

Mechanisms are far more important than correlations, from a management perspective, because only the mechanistic details can predict what interventions will achieve a given management aim. To illustrate, if we want to increase primary production in a grassland area, we need to know whether primary production is driven by the presence of a few important species; the presence of some important functional groups; functional diversity; or species richness. Each of these options would suggest a different management intervention (e.g., introducing a lot of individuals of a single species versus trying to maintain as many species as possible).

BEF researchers have recognized all of these points and are making strides towards a better understanding of the mechanisms connecting species and diversity to ecosystem behavior. However, the current gaps in our knowledge means that there is little basis to justify particular ecosystem interventions on the grounds of BEF evidence alone.⁴⁰ Certainly BEF studies consistently support the conclusions that (a) major declines in species richness and (b) changes in

³⁹ Philosophers of biology won't be surprised to hear that the controversy about BEF mechanisms appears to be turning into a relative significance debate (see Beatty 1997; Dietrich 2020).

⁴⁰ Of course there are other sources of evidence for how to intervene in ecosystems, including people's acquaintance with particular ecosystems; understanding of trophic dynamics or other relevant theory; and knowledge of the effects of past interventions.

taxonomic or functional composition of ecosystems can both alter ecosystem functioning. But these claims are too general to govern *specific* interventions targeting particular ecosystems (e.g., decisions about reserve design, management of an invasive species, or promoting sustainable fisheries).

Third, BEF findings actually cut against the common assumption that species diversity *per se* is the primary driver of ecosystem functions. There is evidence that for some functional properties of some ecosystems, dominant species, rather than diversity of any kind, best explains functional differences (Sasaki and Lauenroth 2011; van der Plas 2019). Where diversity does have causal force, there seems to be agreement that functional diversity rather than taxonomic diversity does most of the explanatory work (Díaz and Cabido 2001; Mori et al. 2013; Gagic et al. 2015; Sakschewski et al. 2016; but cf. Hector et al. 2010). In addition, many ecologists believe that ecosystems exhibit functional redundancy, which means that systems can absorb a certain amount of species losses without changes in functioning (Hooper et al. 2005). If these points are right, then BEF evidence suggests that species richness is a poor candidate for explaining ecosystem functional properties.

This conclusion does support one general management implication, which is that we ought to focus on functional diversity and important functional roles rather than merely species richness when managing communities for functional properties. Chapter 5 will investigate the relationships between biodiversity management and ecosystem management in far more detail. For now, however, I have argued that the BEF research program fails to clearly support the claim that modest arbitrary species richness losses will negatively affect the functioning of communities. Of course, the functioning of many communities is changing. The causes of this that are well-supported by

empirical evidence include human actions and climate change,⁴¹ as well as changes in biodiversity composition⁴² (i.e., what species are present and in what abundances, in contrast to species richness, which is simply the number of species present).

Finally, we should do a common-sense check about how the BEF results relate to global-scale patterns. There is a robust pattern of decreasing richness of taxa with both latitude and altitude (Rahbek 1995; Hillebrand 2004; Hawkins et al. 2011). Thus, temperate forests generally have lower taxonomic richness than tropical forests. In addition to this global gradient, some types of ecosystem are “normally” characterized by a comparatively low diversity of plant species (e.g., many mangrove forests⁴³ and forests on some remote islands⁴⁴). This means that explanations of ecosystem functioning cannot appeal only to species diversity. For instance, one study reports that “despite the extremely low species richness of the Hawai‘i [forest plots], some structural variables, particularly those for the [montane] wet forest, were well within the range of values for the world’s most diverse tropical forests” (Ostertag et al. 2014, 14–16). So, any general relationship between diversity and ecosystem properties must be heavily modulated by other local factors such as climate variables and functional properties of individual species. This observation further supports the expectation that different ecosystems will respond in qualitatively different ways to biodiversity losses.

⁴¹ Díaz et al. (2019).

⁴² See Hillebrand et al. (2018).

⁴³ Ricklefs & Latham (1993).

⁴⁴ Ostertag et al. (2014).

3.8 As Biodiversity-Ecosystem Functioning Findings Become More Causally Detailed

More recently, there has been a concerted attempt to make the biodiversity-ecosystem functioning (BEF) research more relevant to understanding real ecosystems by, e.g., conducting research on multi-trophic systems (e.g., Moi et al. 2021), on larger spatial scales (e.g., Patrick et al. 2021), and doing non-experimental studies (van der Plas 2019). It is likely that in the future, specific BEF findings will be increasingly relevant to management decision-making. As this happens, it will become even more vital to scrutinize the value judgments and assumptions being made in the course of study design with an eye to potential management implications.

To illustrate, consider my view that for the purpose of management we should be concerned about the sustainability of ecosystems into the future under the expected conditions of climate change, rather than trying to “dial back the clock” or preserve ecosystems in a static manner. For the sake of discussion, make the admittedly unlikely assumption that all interested parties agree with me about this. If so, then studies ought to be designed in a way that helps us understand what will promote the sustainability of ecosystem functions rather than the temporal stability of ecosystem composition.

Let me further show why this matters with a hypothetical example involving different research proposals. Suppose research team A wants to investigate how whole ecosystems with different biodiversity levels differ with respect to temporal stability. They will estimate this by tracking how much year-to-year variance there is in the population sizes in communities with differing diversity levels under relatively constant environmental conditions. Research team B wants to investigate how whole ecosystems with different biodiversity levels differ in their resilience with increasing average temperatures. They will estimate this by tracking how much year-to-year variance there is in the population sizes under some different treatments which

increase temperature over time. Research team C also wants to investigate how whole ecosystems with different biodiversity levels differ in their resilience with increasing average temperatures. They will estimate this differently, by tracking how much year-to-year variance there is in the total biomass of experimental plots under some different treatments which increase temperature over time.

Assume that these studies have solved the various current problems with external validity (including problems of statistical design and causal inference). Each of these studies would intuitively tell us something having to do with ecosystem functioning. However, the project of team C is most relevant to management under the stipulated conditions, because team C will track the resilience of ecosystem functions under changing climatic conditions.

This is important because is very plausible that studies A, B and C would come up with divergent conclusions. For example, team B might find that temporal stability of populations decreases with diversity, while team C might find that temporal stability of biomass increases with diversity. Existing studies have reported similar results (Gross et al. 2014). If this were the case, we would have to make a judgment about whether the property studied by team B or C better tracks ecosystem functioning. There are no naturalistic facts which can tell us whether ecosystems “should” exhibit more stability of populations versus more stability of biomass. We have to make a value judgment about how we prefer ecosystems to behave, which in turn will affect how we manage ecosystems into the future. Under the stated public preferences, the researchers should take stability of productivity to be a better indicator of functioning. Hypothetical study B’s results would be misleading if the disconnect between the study design and public preferences was not noticed (an instance of the “crowding out” problem cited above). This illustration shows how

unexamined value judgments within research design could unexpectedly influence management decisions down the road.

I have argued that at present, the potential for such disconnects is limited because the BEF results can tell us less than one might assume about how to manage ecosystems. But as we move forward, it is vital that the presence of values within this research program continue to be recognized and discussed.

4.0 A Constructivist Account of the Ecosystem Health Concept

4.1 Introduction

The concept of ecosystem health is frequently used to frame management programs, ecological studies, and popular discussions about environmentalism, yet there have been few recent discussions of the concept in the philosophy of science. In this chapter I will develop a constructivist view about ecosystem health—i.e. a position on which collective human values help to distinguish healthy from unhealthy ecosystem states. I argue that a constructivist view better captures common judgments about ecological health than any plausible naturalist view. I also propose some reasons to reject eliminativism about the ecological health concept.

By explicitly characterizing the values at work in ecosystem health judgments—something which has been done infrequently even by other constructivists—I open the doors to further assessment and evaluation of the health concept. My constructivist position also illustrates the ubiquity of anthropocentric value judgments in assessments of environmental health. The major values I identify pertain to aesthetics, authenticity, and human wellbeing. Although this chapter primarily aims to characterize ecosystem health as it is currently conceptualized, I emphasize that this conception of health is not set in stone and that the values at work should be subject to further evaluation.

To reiterate, the position I develop is not meant to be an “ideal” normative view of the ecosystem health concept. It is primarily meant to capture common intuitions about the health status of familiar types of ecosystem. However, I will at some points also make some normative arguments that certain properties of ecosystems should not be thought to be relevant to their health.

In some of these cases, others have thought that a certain property (such as the extent of human influence on an ecosystem) tracks its intuitive health status, but I argue that this is incorrect. I will also sometimes argue that health should be distinguished from other ecological properties for management reasons—i.e., to promote clear reasoning about divergent management objectives. So, as with many accounts of health, I aim to capture intuitions while balancing this against select normative considerations.

Currently, there are both too many and too few different accounts of ecosystem health. There are many loosely related studies of ecosystem health in ecology and environmental ethics which may or may not utilize an explicit definition of “health”; yet there are few general and up-to-date philosophical discussions of the nature of the concept. This chapter aims to unify the discussion by surveying the relevant literature across disciplines. My final account of ecosystem health draws insights from scientific and philosophical work and aims to be one of the most comprehensive and general (i.e., applicable across ecosystem types) accounts which has been developed so far.

In the next section I clarify what we should be asking about when we discuss the nature of ecosystem health in a philosophical context. Then I present an overview of potential ways to naturalize the ecosystem health concept. I argue that none of these naturalistic strategies work, which leaves us with a constructivist view. On the view I develop, ecological health assessments are based on natural properties of ecosystems, including resilience and trophic structure, together with judgments about aesthetic value, authenticity, and human wellbeing. At the end of the chapter, I consider arguments for and against eliminativism about *ecosystem health*.

4.2 Clarifying the Question

The existing literature on ecological health is quite decentralized, so I begin by mentioning the relevant sources which I take as a starting point.

- (1) Some earlier environmental ethics literature (c. 1990s) includes discussion of ecosystem health, often from the perspective of whether we have moral obligations towards ecosystems.⁴⁵
- (2) A number of ecologists and philosophers have written criticisms of the ecosystem health concept, often in scientific publications.⁴⁶
- (3) Various teams of ecologists have developed frameworks for assessing the health of particular ecosystem types (c. 2000-present).⁴⁷ A few ecologists, notably Costanza and Mageau (1999), have also presented more general accounts of the ecosystem health concept.

A majority of the above sources are published in either ecology journals or works in environmental ethics, which may partly explain why the topic of ecological health has not been as discussed among philosophers of science. By surveying and responding to some of the existing work, I hope to encourage further discussion of this topic.

⁴⁵ Sources include Callicott (1992; 1995); Sagoff (1992); McShane (2004).

⁴⁶ Ehrenfeld (1992); Suter (1993); Jamieson (1995); Russow (1995); Wicklum and Davies (1995); Lancaster (2000); Lackey (2001); Odenbaugh (2010); and Newman, Varner and Lindquist (2017, chap. 9); Rapport (1998) responds to some criticisms.

⁴⁷ Examples include Rapport et al. (1998); Tett et al. (2013); and Harwell et al. (2019).

It will be helpful to frame the discussion about ecosystem health in a way that parallels discussions about human health, since the latter issue has been more philosophically developed. (However, we should not assume that the two concepts will behave similarly.) To begin with, I will adopt the convention of distinguishing between the “constitution question” and the “status question” about health and disease or dysfunction (Glackin 2019). The constitution question asks what natural properties constitute a particular state of health or disease. The status question asks what makes the state count as either healthy or unhealthy. The status question is the philosophically interesting question in this context.⁴⁸ Our question is, among ecosystems’ many scientifically characterizable states, what makes some of those states healthy and other states dysfunctional?

The constitution question is important for management purposes, and has been addressed in detail by some teams of ecologists for particular kinds of ecosystems (Rapport et al. 1998; Tett et al. 2013; Harwell et al. 2019). These frameworks identify properties of ecosystems that indicate or characterize a healthy state, including productivity, trophic structure, nutrient cycling regimes, species functional diversity, and other properties relevant to specific ecosystem types (e.g., salinity or fire frequency). Ecologists also have identified factors which cause changes in these properties and thereby threaten the health of the ecosystem. The ecologists developing these frameworks assume as a starting point that certain states of an ecosystem are healthy or unhealthy, presumably based on common consensus. Although these works are important resources for understanding ecological health, they often do not directly provide answers to the status question.

⁴⁸ In the literature on human disease—particularly on mental disorders—there have been disputes about the constitution question (e.g., whether there is an appropriate sort of biological state underlying recognized mental disorders). These questions are philosophically interesting, but there are not parallel questions about ecosystems.

There are three approaches to answering the status question: naturalism, constructivism, and eliminativism. Naturalism is the position that the distinction between healthy and unhealthy states can be made on purely naturalistic grounds. In the human case, proposals for naturalizing health include appeals to population statistics or biological functions (e.g., Boorse 1977; Hausman 2012). Below I will survey potential strategies for naturalizing ecosystem health, which will tend to differ from strategies for human health.

Constructivists hold that the distinction relies in part on social values, e.g., in the human case, collective views about what is “normal” or beliefs about what states help to promote a good life.⁴⁹ “Constructivism” here is a blanket term encompassing views which emphasize the roles of norms or values in answering the status question. Constructivists need not hold that disease or dysfunction states are recognized arbitrarily, and they can accept naturalistic answers to constitution questions.

Finally, eliminativists argue for the rejection of health or disease concepts. Note that there is a difference between typical eliminativists about human disease versus ecosystem health. Typical eliminativists in the human disease literature argue that particular conditions should not be considered diseases.⁵⁰ In contrast, typical eliminativists in the ecosystem health literature argue for the wholesale rejection of discourse about ecological health. This position will be addressed later in the chapter.

Someone might argue from the presence of human values in judgments about ecosystem health to the conclusion that such talk should be eliminated. So, the three views are not all mutually

⁴⁹ As I have characterized constructivism, it will include both normativism and “hybrid” views about human disorders.

⁵⁰ E.g., Szasz (1960) on mental illness.

exclusive. However, constructivism is distinct from and does not immediately entail eliminativism—the move from one to the other requires an argument. In addition, there are some arguments for eliminativism about ecosystem health which are unrelated to constructivist claims. So, it is important to clearly distinguish these two positions.

The dialectical situation for the constructivist is as follows. I will treat naturalism about health as the default position.⁵¹ To support constructivism, I will show in the following section that strategies for naturalizing ecosystem health fail. The reason for making this kind of negative argument is that the philosophical literature about the nature of the ecosystem health concept is relatively undeveloped. So, if I begin my argument by identifying cases where values seem to be at work in an application of the concept, someone might claim that the concept is being used unclearly, metaphorically, or inappropriately in the example or that the example is not representative. What we need is therefore an argument that the concept cannot be naturalized. To this end, I will consider all of the potential naturalistic views of which I am aware and argue that they each fail to capture ordinary judgments about ecosystem health. However, it might be noted that other commentators on the ecosystem health concept have simply taken for granted that the concept is value laden (Gillette et al. 2021, sec. 2.3).

⁵¹ Although my anecdotal impression is that naturalism is the least popular view among scholars who have written about the ecosystem health concept.

4.3 Naturalist Views of Ecosystem Health

Previous philosophical discussions about ecosystem health have not usually been framed in terms of the naturalism/constructivism distinction. Instead they have often jumped directly to arguments for eliminativism or to positions about how ecosystem health is related to our moral obligations. While these latter issues are important, a failure to first clarify the nature of the health concept has sometimes invited confusions, often resulting from the failure to distinguish constructivism from eliminativism.

Naturalist views about ecosystem health hold that the status question can be answered by appealing solely to physical, biological, or historical properties or facts. What follows is a list of potential natural properties which could serve the function of demarcating healthy from unhealthy states of ecosystems. Some of the disanalogies between human and ecosystem health will be exhibited in the following discussions, since properties which have been advocated for accounts of human disease are very implausible for the ecosystem case, and some properties which have been proposed to pick out healthy ecosystems are irrelevant to the human case.

4.3.1 The Human Influence View

The human influence view about ecosystem health has been historically and popularly important and so it is worth discussing, even though it is subject to a number of well-known objections. In the early twentieth century, ecology was developing as a discipline at the same time as there was growing concern about ecological disasters caused by (among other problems) poor management of agricultural practices and rangeland. Early ecologists recognized the roles of humans in creating ecological disasters. They were focused on the theory of succession, the

process by which ecosystems change structure and composition over time following a disturbance. So, early ecologists such as Frederic Clements thought that ecological disasters like the Dust Bowl occurred when human influences altered the trajectory of succession.⁵²

We now have more ecological theory than just succession for describing how communities and ecosystems behave. The Clementsian view that succession generally has a single “target” endpoint—albeit an unstable endpoint that can be altered by climactic changes and human influences—has evolved into a more complex dynamical view of ecosystems in which they can follow multiple trajectories, can shift between stable regimes, and are influenced by surrounding ecosystems via processes at the landscape⁵³ level. There is growing recognition of the ways both conventional ecological processes and stochastic effects may jointly influence ecosystem properties. Although our ways of explaining ecosystem behavior have gotten more complex, the Clementsian way of thinking about ecological degradation remains plausible for many people. In my experience, it is common for people to observe that humans are responsible for our present ecological problems and to hold that ecosystems remain healthy to the extent that they have escaped human influence.⁵⁴

⁵² Clements’ views about succession are helpfully summarized in the first several pages of his “Nature and Structure of the Climax” (Clements 1936).

⁵³ In ecology, the term “landscape” has a meaning distinct from the colloquial sense. It refers to the level of organization above that of the ecosystem. Landscapes are often conceptualized as a patchwork of different ecosystem types, ecosystems at different stages of succession, etc.

⁵⁴ There is an interesting longer history of views about the environment shifting from thinking of nature as a danger to humans, to thinking of nature as an inexhaustible material and spiritual resource, to the present view of humans as a danger to nature. See Merchant (2007).

It is true that human activities have resulted in widespread ecological changes, and that we could have avoided various ecological disasters had humans collectively modified some of our actions. However, there are well-known problems with the view that healthy ecosystems can be identified by appealing solely to the absence of human influence.

William Cronon (1996) argues against the view that conservation ought to prioritize protecting wilderness to the exclusion of other aims, where wilderness means “nonhuman nature.” Arguments of the same form can also be used to target the human influence view about ecosystem health, so I will paraphrase two of Cronon’s points here, slightly rephrased to target the health concept.

First, human activities are not inherently distinct from natural processes. Humans were part of ecosystems and exerted influences on ecosystems long before recent environmental disasters. The view that humans are inherently destructive of nature contributed to morally objectionable policies, such as forcible removal of native people from US national parks. So, there are conceptual, factual, and moral difficulties with the assumption that humans are inherently at odds with the health of nature.

Second, if the concept of ecological health only pertains to relatively “natural” ecosystems, then we can have no sensible concept of the health of urban ecosystems or agricultural ecosystems. But we should particularly care about the health of these “artificial” systems due to their direct effects on human wellbeing.

In addition to these two arguments, I will add a third objection pertaining to management strategies. It is widely accepted that, under our circumstances (climate change, accelerated movement of species around the globe, widespread use of fertilizers and pesticides, the history of fire suppression, etc.) the best way to manage ecosystems is not to leave them alone. Instead, active

management is needed to ameliorate the effects of past human actions. But if we take human influence per se to delimit healthy from unhealthy ecosystems, then there is no way to distinguish destructive from protective human influences. Thus, we would end up with the unintuitive consequence that ecosystems which are intensively and apparently successfully managed may count as less healthy than ecosystems which experienced past degradation and have been left unmanaged. So, the human influence view is at odds with how we presently manage ecosystems. Therefore, without denying that humans are causally responsible for much environmental harm, we should reject the view that the mere presence or absence of human influence is what distinguishes healthy from unhealthy ecosystems.⁵⁵

4.3.2 The Historical Benchmark View

Someone who retains the intuition that ecological harm is closely related to recent human actions might instead appeal to a historical benchmark view—the view that ecosystems are unhealthy to the extent that they have recently changed from their historical states. This view has the benefit of not categorically separating human activities from ecological processes. Instead, one might think, it is only relatively recent and high-impact kinds of human activities, including

⁵⁵ Human influence is not a sufficient condition for environmental degradation, but could it be a necessary condition? This question is likely not practically important, because humans have influenced most ecosystems, so the condition obtains in all contemporary cases of interest. Since I am a constructivist about health, I think there is not a naturalistically-fixed answer to this question, but it would depend on further thought about our reasons for considering certain ecological states unhealthy.

industrial activities, mass transportation, and mechanized agriculture, which are ecologically harmful because they have the ability to substantially and rapidly alter ecosystems.

Historical benchmarks have sometimes been used and advocated for identifying “intact” ecosystems for conservation priority. “Intact” ecosystems are variously defined as ecosystems which have been subject to less human influence or ecosystems with low rates of species extirpation. The authors of one recent study (which finds that under 3% of earth’s land surface is “faunally intact”) chose to “focus on changes that have occurred since the year 1500 AD, because this is the baseline date for assessing species extinctions within the IUCN Red List of Threatened species” (Plumptre et al. 2021, 2). So, this study defines faunally intact ecosystems as those whose animal species composition has changed little since 1500.

This benchmark choice makes methodological sense for the purpose of the study given the nature of IUCN data. However, I have reservations about the importance of this historical benchmark for conservation practice. The most intact areas identified by this study (i.e., areas with the fewest animal extirpations) are unsurprisingly areas that many humans would consider remote, e.g. northern Canada, Siberia, the Sahara desert, and parts of the Amazon basin. Areas with higher human population densities or more intensive agricultural practices will of course have higher rates of species extirpation.

The fact that the identified sites have changed less in faunal composition over the past 500 years does not by itself mean that they ought to be prioritized for conservation or replicated elsewhere. Briefly consider various reasons why these sites should be prioritized. I have already argued that preventing human influence per se is not an appropriate conservation goal or health metric. The authors of the study suggest that the intact areas constitute “wilderness” and that protecting wilderness is an important conservation goal. I am inclined to agree with the latter claim,

but the justification for protecting wilderness is often that it permits recreation or nature experiences, or holds cultural significance. To the extent that the identified sites are very remote, this justification for protection is less likely to apply to them.

If the goal is instead to protect global biodiversity,⁵⁶ then presumably we want to know a lot more than just how “intact” the site is. For example, we would want to know about the number of species present at a location and the current risk to those species. Some locations which are faunally intact may have low diversity to begin with or contain species which are not in danger of extinction, in which case the intact sites may have low importance for biodiversity conservation.

If the goal is to protect the functioning of the ecosystems, then the number of recent species extirpations also has little immediate relevance. As discussed in Chapters 2-3, the loss of a single functionally important species can have a major effect on ecosystem functioning, so an area could count as relatively compositionally intact and yet be in a completely different functional state from 500 years ago. On the flip side, ecosystems can withstand losses of a number of species without suffering functional changes.

In short, various plausible conservation aims provide no clear justification for prioritizing or attempting to replicate sites which count as intact based on the historical benchmark. Some of these sites may be important for conservation, but the mere fact that the sites have changed less in composition since 1500 cannot provide the reason why.

A final comment is that, if we should really conclude that just under 3% of earth’s land surface is intact,⁵⁷ then perhaps we ought to give up on intactness as a conservation goal. Instead,

⁵⁶ This conservation goal is discussed in Chapter 5.

⁵⁷ Other ways of estimating intactness will come up with higher percentages.

we should focus on goals and assessment methods which are applicable to the remaining majority of earth's land surface.

The above discussion illustrates how historical benchmarks have recently been applied in research which aims to inform conservation decision-making.⁵⁸ I expressed reservations about this use of historical benchmarks. The authors of the cited work do not explicitly tie “intactness” to health, but this seems like an obvious possible interpretation. However, similar problems will also arise for attempting to evaluate health based on a historical benchmark.

First, there is the problem of justifying the historical benchmark date for the healthy state of an ecosystem. The study I cited above chose 1500 for practical and methodological reasons. This may be justified for the purpose of a research study, but for the purpose of deciding what ecosystems to consider healthy, the choice of date is too arbitrary to be helpful. There are two distinct problems here. One has to do with the relevance of benchmarks as such. Even if a benchmark date *tracks* health assessments—i.e. some healthy ecosystems have in fact changed less since a certain date—this does not really illuminate why we think those ecosystems are healthy or unhealthy.⁵⁹ A definition of health in terms of a benchmark date can be extensionally adequate at best. Notice that in order to even know whether the definition is extensionally adequate, we need to know exactly which ecosystems are healthy and unhealthy to begin with. This is why benchmark-based definitions fail to really illuminate our conception of health.

⁵⁸ On historical fidelity as a management goal, see Garson (2014) and Desjardins (2015).

⁵⁹ Unless our conception of health amounts to *it hasn't changed much recently*. This conception of health is probably not useful.

The second problem pertains to the choice of benchmark date. If an ecosystem would count as healthy based on one plausible benchmark and unhealthy based on another, there is no way to choose between the dates, and therefore no way to resolve the question whether it is healthy. This might actually happen for some ecosystem types which switch between different regimes somewhat regularly (such as many aquatic systems).

In addition, there may be a further issue regarding the justification for the choice of a benchmark date. Going back to the intactness study, Plumptre et al. quote IUCN biodiversity assessment guidelines suggesting that a reason for choosing 1500 as a benchmark is that it is before major industrial activity (IUCN 2016). So, they seem to be trying to capture the state of the ecosystem before major human influence. There are a couple of problems with this reason for choosing a date. First, humans have had major influences on ecosystems well prior to 1500 CE or the development of industrial technology.⁶⁰ Thus, the choice of 1500 CE rather than 1500 BCE or 15,000 BCE appears unprincipled based on the given reason. Second, if you choose a benchmark date based on the belief that it represents a pre-industrial date or a pre-human influence date, then you are assuming that human-caused changes are what determine the health of ecosystems. This is not equivalent to the human influence view, since it allows for human restoration activities (i.e., humans can improve the health of an ecosystem by cancelling out our past influences). However, it still relies on an assumption that anthropogenic changes are detrimental to ecological health. I have already discussed some reasons for rejecting this viewpoint. Our view of ecosystem health

⁶⁰ Agriculture seems to have begun by 11,500 BP, but there is evidence of humans modifying plant communities earlier than this (Zeder 2011). There is also evidence that the arrival of humans in North America contributed to declines in some megafauna populations starting around 13,000 BP (Broughton and Weitzel 2018). So, humans have been affecting the composition of ecosystems for at least this long.

should be able to distinguish between helpful and harmful human changes, and should allow for the fact that humans are (and have historically been) components of many ecosystems.

Focusing on past benchmarks bodes poorly for the future. Earth's land surface has been extensively altered by human activities, and climate change is causing rapid further ecological changes. If we continue to focus on the composition of ecosystems in 1500, we will only find that ecosystems are becoming less and less healthy. But ecosystems are not static entities, and even in the absence of humans they undergo major changes in functioning and composition over time. Humans are causing ecosystems to change more rapidly than they otherwise would, but even in an identical world without humans, no ecosystem would remain frozen in its year 1500 state for the rest of time. What we need is a forward-looking view that can help to distinguish healthy from unhealthy ecosystems even as all ecosystems undergo major changes.

Finally, the historical view gives us little basis for assessing the health of urban, intensive agricultural, or artificial ecosystems, since some of these ecosystems appeared more recently than any plausible benchmark date. Presumably some agricultural systems are healthier than others; it is not the case that all recent ecosystems are unhealthy just in virtue of being recent.

4.3.3 Normal Function Views

It is commonly thought that ecosystems are healthy when they are functioning properly, and ecological degradation occurs when ecosystem functions break down. This leads one to wonder whether an account of health could be developed in terms of normal functions.⁶¹ Since

⁶¹ E.g., Callicott (1995), although his view also includes a role for values.

ecosystems support two levels of functional description (Chapters 2-3), there are two ways one might try to construct this account. The first is in terms of normal functions of ecosystem components, and the second is in terms of normal functioning of the ecosystem as a whole.

First consider an account of ecosystem health based on the functions of components. This would require components of ecosystems to have normal or proper functions from which they can deviate or malfunction. However, as I argued previously, the functions of ecosystem components are not proper functions but are a type of causal role function, where causal role functions do not permit malfunction (Chapter 2). For example, even when a large population of deer overbrowse a forest, the deer themselves have not malfunctioned but have simply carried on their role of browsing. Even if you accept a coevolutionary account of some role functions (Millstein 2020b), you should not conclude that organisms which depart from a selected ecological function are malfunctioning. This would require a teleological view of ecosystems on which organisms “should” perform certain functions for the good of the ecosystem or in order to support certain ecosystem functions—a view which is widely rejected because it lacks a basis in evolution or other biological theory. While I grant that some role functions have probably been selected for, they were selected for because they provided a fitness benefit to the organism or other evolutionary unit, not because they provided a benefit to an entire ecosystem. Because role functions are not proper functions in this sense, they cannot provide any basis for evaluating the health of ecosystems.⁶²

⁶² Notice the contrast with the human health case, where it is *prima facie* plausible that normal functions of components have something to do with health and certain diseases.

Next consider building an account of health in terms of the functioning of the ecosystem as a whole. A first go at such an account might say that ecosystems are healthy when they are functioning (well), and less healthy when they experience dysfunction. Potentially a lot more could be said about the relationship between the concepts of ecosystem functioning and ecosystem health. However, for my present argument the relevant issue is whether or not ecosystem functioning can provide a naturalistic basis for answering the status question. Ecosystem functioning cannot do this, because ecosystem functioning is not a naturalistic concept. As I argued previously (Chapter 3), it is a thick or evaluative concept. So, although functioning is connected to ecosystem health, this connection does not provide a basis to evaluate health in a fully naturalistic manner.

An alternative would be to appeal to ecosystem functions rather than ecosystem functioning. *Ecosystem function* is a non-evaluative concept, since it simply refers to any ecosystem process such as nutrient cycling, primary production, and decomposition. However, ecosystem functions alone cannot provide an answer to the status question, which asks why certain ecosystem states are considered healthy or unhealthy. Describing the functions of an ecosystem in a healthy state answers the constitution question, not the status question.

This is not to say that ecosystem functions are not relevant to thinking about ecosystem health. Instead, what we need is an explanation why certain ecosystem functions characterize healthy ecosystems of some kinds (but perhaps not others). Notice, for instance, that nitrogen fixation is healthy for many terrestrial ecosystem types, but nitrogen fixers have become invasive and caused major ecological changes in places like Hawai'i with little native nitrogen fixation (Vitousek and Walker 1989). So, the same function can be healthy for one ecosystem type but

destructive of another. This illustrates why the difference between healthy and unhealthy ecosystems cannot be explained in terms of functions alone.

4.3.4 Good-For-Ecosystem View

Katie McShane has a view on ecosystem health which focuses on ecosystem functions (and structure) but provides an answer to the status question. Her view is as follows:

[E]cosystem health [is] a matter of maintaining the structure and functions that are good for the ecosystem.

She goes on to appeal to a view about what it means for a state to be good for an entity:

In order to determine which structure and functions are good for the ecosystem, we should ask what it would make sense for someone who cared for the ecosystem to want for it for its sake. (McShane 2004, 245)

I read this account as providing answers to both the constitution and status questions. McShane identifies certain structures and functions as the properties which constitute states of health. She then answers the status question by the appealing to the good of the ecosystem—i.e., the structures and functions which are good for the ecosystem are the healthy ones.

To be clear, it would require a significant detour to decide whether this view should be considered naturalistic or non-naturalistic. I have included McShane's view in this section because it is representative of a type of view that is clearly distinct from the sort of constructivist view I argue in support of. The present view states that there are certain (health) states that are inherently good for ecosystems. This way of thinking about ecosystem health is distinct from a historical or human-influence view, which takes ecosystem health to be determined by extrinsic processes. It is also distinct from my constructivist view, on which ecosystem health is partly determined by

specific human values and preferences, a view which does not presuppose that certain states are inherently better for ecosystems.

There is a major epistemic concern about views that treat health states as inherent features of ecosystems: how can we tell which states of ecosystems are good for them? McShane seems to assume that people who care for ecosystems can either intuit or discover which states are the good ones. But it is not apparent how this can be done. Most people working in ecosystem management or environmental ethics presumably care about ecosystems, but they frequently have disagreements about the appropriate goals of ecosystem management. Moreover, an eliminativist about ecosystem health can reasonably claim to care about ecosystems, and yet not to believe that health is one of their properties. So, appealing merely to what states are good for ecosystems—analyzed in terms of what those who care for ecosystems should want for them—may seriously underdetermine which states of ecosystems are the healthy ones.

A related problem with an account like McShane's is that it does not single out the health property. People who care about ecosystems may want various things for those ecosystems—for them to remain biodiverse, for them to be sustainable, for them to be resilient, and so forth—and presumably these properties are non-identical with health. The account therefore fails to clarify the nature of ecosystem health in contrast to other related properties.

For these reasons, it is not enough to argue that some functions of ecosystems are good for them. For practical applications we need a more detailed account of what considerations humans have in mind to when they believe an ecosystem is in a healthy state. I will lay out some of these considerations below.

4.3.5 The Resilience View

Perhaps the most popular view currently is to think about ecosystem health in terms of the system's resilience, sustainability, and similar properties.⁶³ The idea is that ecosystems are healthy when they are able to withstand environmental changes, perturbations, invading species, and other external processes which tend to alter or disrupt ecosystem functioning. This way of thinking about health is said to be inspired by Aldo Leopold's view that health is the "capacity for internal self-renewal," as described in his 1941 essay "Wilderness as a Land Laboratory" (Leopold 1991).⁶⁴

A more recent example that has generated some discussion is the view developed by Costanza and Mageau (1999) who claim that healthy ecosystems are *sustainable*, where they define sustainability as the resilience of both the system's organization and its "vigor" (i.e., productivity). Thus, in effect, their view holds that ecosystems are healthy when they are able to maintain their organization and productivity in spite of perturbations.

There are many potential variants on views like this. One could place more emphasis on the resilience of ecosystem structures versus processes. One could also distinguish between an ecosystem's ability to maintain its current state versus the ability to adapt to environmental changes. However, I treat all these potential views together because they are all subject to the same kind of objection. Although the following objection is based on commonly-known facts about

⁶³ Published sources include Rapport et al. (1998), Costanza and Mageau (1999), and Döring et al. (2015). In my experience, resilience is among the first things to come up in discussions of environmental health with other philosophers.

⁶⁴ A more detailed reconstruction of Leopold's view on land health is provided by Millstein (2020b) and will be discussed below.

ecosystems, I am not aware of any previous sources which have discussed these facts as a problem for resilience accounts of health.

The problem is that degraded or unhealthy states of ecosystems can be very resilient, in the sense that they resist a regime change. A well-known example of this is eutrophication in aquatic systems. It can be difficult to get the system to revert from a eutrophic to an oligotrophic state due to feedbacks stabilizing the eutrophic state. Some processes which contribute to ecosystem degradation, including erosion and invasion, can also generate positive feedback cycles which encourage further erosion and invasion. Ecologists think that some kinds of regime following a regime shift are so stable they are effectively permanent and cannot be reversed (Folke et al. 2004, 567). But the fact that degraded states can be very persistent, stable or resilient does not make them healthy.

One might wonder whether healthy states of ecosystems are in general more resilient than unhealthy states with respect to specific properties—for example, whether they are more likely to retain their level of productivity or diversity when the climate changes. Answers to these kinds of question will be complicated and probably vary across ecosystem types. My view is that we should not rely on ecological science to produce desirable answers to these questions. Some otherwise healthy ecosystems might in fact be unable to resist or recover from certain kinds of environmental change.

An issue to be aware of is that empirical assessments of resilience are generally influenced by judgments about what properties or states of systems “should” be resilient. Since resilience is

often operationalized as a measure's time to or ability to return to an initial state,⁶⁵ the assessed resilience of a system depends on what properties the investigators choose to measure and what they take to be the beginning state of the system. Investigators usually assume that we are interested in the resilience of a relatively healthy state of an ecosystem. But conceptually, nothing precludes one from measuring the resilience of a degraded ecosystem in the same manner, taking the degraded state as the starting point.

What authors like Costanza and Mageau seem to be thinking when they link health to resilience is that, e.g., an oligotrophic lake is healthier when it is resilient, meaning among other things that it resists becoming eutrophic. A lake which is already in a eutrophic state is presumably not healthier if that state is highly resilient. Similarly, one might think that a coral reef is healthier if it is more resilient in the face of changes to ocean temperature and acidity. A reef that has already experienced a regime shift is not considered healthier in virtue of the degraded state being difficult to alter. Therefore, resilience functions as an amplifier. An ecosystem in a healthy state is healthier in virtue of being resilient, while an ecosystem in a degraded state is less healthy if that state is highly resilient.

So, accounts that link health to properties like resilience rely on healthy states of ecosystems already having been identified. Although knowing about resilience is important to understanding what factors may compromise ecosystem health, resilience by itself cannot provide

⁶⁵ See Holling (1973); Westman (1978); Carpenter et al. (2001); and Desjardins et al. (2015) on resilience. Resilience is alternatively operationalized as the magnitude of disturbance a system can absorb before experiencing a regime shift; the same argument applies to this formulation. For the purpose of my discussion here it is not important to distinguish the various related senses of "resilience."

a basis to answer the status question. This is because, both conceptually and empirically, resilience is not a property that is unique to the healthiest states of ecosystems.

4.3.6 Views with Multiple Criteria

To conclude the section on naturalism, consider the possibility of appealing to several of the above properties jointly to answer the status question. There are several ways one might include multiple criteria in an account. I will first consider conjunctive and disjunctive accounts.

First, given the way the above properties individually fail to accurately delimit healthy from unhealthy ecosystem states, it is unlikely that a conjunction of the properties will fare better. Asking ecosystems to fulfill several of these criteria at once will likely increase rather than decrease the number of errors.

A disjunctive account is no more promising. As I have shown, many candidate demarcation properties fail in both directions, i.e., they count some unhealthy ecosystems as healthy and some healthy ecosystems as unhealthy. Consider the disjunctive view that an ecosystem is healthy when it is close to its historical state *or* when it is resilient. This view will potentially decrease the number of false negative results—i.e., it will reduce the number of cases where a healthy ecosystem is judged to be unhealthy. However, adding disjuncts will not eliminate the false positives associated with either of the criteria. Assuming the criteria are independent, the total number of false positives will probably increase.

A remaining alternative is to reject the idea of providing explicit necessary-and-sufficient conditions to answer to the status question, but to provide a list of relevant factors such that healthy ecosystems generally exhibit several of the factors, and unhealthy ecosystems generally fail to exhibit many of the factors. This approach is warranted since it is more likely to remain true to

practice and to allow for the fact that the concept is somewhat flexible. However, recall that we are trying to justify a naturalistic answer to the status question. With the approach under consideration, if value judgments seem relevant in practice to some judgments about ecosystem health, then it will be very hard not to include these value judgments in the account. Put differently, if someone declines to give necessary-and-sufficient conditions, then it becomes hard to justify the categorical exclusion of non-naturalistic considerations from the account of ecological health.

This concludes my discussion of possible naturalist accounts of ecological health. The failure of the above criteria, both individually and jointly, to provide an accurate and naturalistic account of the consensus distinction between healthy and unhealthy ecosystems strongly suggests that some other factors must be at work. I claim that the additional factors are certain kinds of value that are relevant to our judgments about ecosystem health. The following section considers some of the relevant values.

4.4 Constructivism

The view that values are relevant to ecosystem health judgments has previously been endorsed, e.g. by Ehrenfeld (1992), Callicott (1995), Rapport (1995), and Rapport, Costanza and McMichael (1998). However, there has not been much recent discussion of what specific norms or value judgments are at work.⁶⁶ This section will discuss some of the values which are relevant to ecological health judgments: aesthetic value, authenticity, and human wellbeing. These values

⁶⁶ Though see Jones (2021) for a related discussion of how specific values are relevant to the concepts of *regeneration* and *degradation* as applied to coral reefs.

have been discussed a fair bit in environmental ethics as candidate reasons to justify conservation, but they have been less discussed as internal components of our judgments about the nature of ecological health.

Recall that these values are supposed to help answer the status question: what distinguishes healthy from unhealthy states of ecosystems. Accepting the relevance of these values to the status question is consistent with affirming that we can describe the composition of healthy states of ecosystems in a naturalistic manner. I also do not intend to give a blanket endorsement to these values as the best way to assess ecological health. My claim here is that these values are relevant to health assessment and reliably track and helpfully illuminate many intuitive judgments about health. Further consideration of these values is important since they have the potential to influence empirical research findings and management work on ecosystems. We should especially consider whether it is appropriate that anthropocentric values clearly influence how we assess the health of nonhuman entities.

I will present the values by walking through some examples. Suppose there is a small local pond and you wish to make a judgment about its health. The pond has multiple potential regimes:

- (1) An oligotrophic state, which is characterized by clear water, a lot of macrophytes (large aquatic plants), and multiple trophic levels of animals including fish, frogs, and large insects.
 - (1a) Incoming water is relatively nutrient-poor.
 - (1b) Incoming water is runoff from heavily fertilized lawns and contains elevated levels of nitrogen and phosphorus compounds.
- (2) A eutrophic state, which is characterized by turbid water, large amounts of algae, and fewer macroscopic plants and animals.

- (3) A farmed state where the pond is heavily stocked with a species of fish for human consumption (to the exclusion of other large animals) and requires more intensive management to feed the fish and maintain water quality.

The oligotrophic state is the healthy state of the pond. 1a, with nutrient-poor runoff, is intuitively healthier than 1b, in part because the nutrient-rich runoff makes the pond more susceptible to shifting to the eutrophic state. The eutrophic state is considered unhealthy. The farmed state is plausibly less healthy than the oligotrophic state qua ecosystem, although it might be more or less healthy qua aquacultural system.

We will consider why the eutrophic state is considered less healthy than the oligotrophic state. As I argued for at length above, natural properties of the pond cannot fully explain this distinction. Briefly consider the following natural properties.

Functions: Both the oligotrophic and eutrophic states experience ecological functions such as primary production and nutrient cycles; the systems exhibit some differences in function but neither has experienced a complete loss of function. In addition, the eutrophic state is more productive (generates biomass more quickly) than is the oligotrophic state of the pond. Lower productivity is connected to better health in many types of aquatic system. In contrast, higher productivity is connected to good functioning and health in many terrestrial systems (Chapter 3). This strongly suggests that value judgments are at work in these health judgments. Natural facts about functions alone do not track ecosystem health in a consistent manner.

Resilience: The oligotrophic and eutrophic states may both be resilient in the sense that they resist a regime shift. In fact, if the pond is in state 1b—an oligotrophic state that is experiencing increased nutrient input—it is likely to be *less* resilient than the eutrophic state. Nevertheless, we would judge a pond still in an oligotrophic state to be healthier than one which

has undergone eutrophication. As discussed previously, this puzzle has been overlooked by most authors who connect health to resilience. Resilience does not determine ecosystem health but amplifies judgments about health or degradation.

History: Researchers studying lake sediment cores have documented regime shifts in aquatic bodies, due to climate changes, going back many tens of thousands of years ago (Cvetkoska et al. 2016). In the hypothetical case, absent historical data, we cannot assume that the pond has been consistently oligotrophic or eutrophic at any point in the past. The eutrophic state (2) should be considered unhealthy regardless of when it first occurred. This shows that choice of a specific historical benchmark is unlikely to reliably track health judgments, as argued for previously.

Consider now how values can help to explain shared intuitive judgments about the health of these states. Judgments about pond health do align with aesthetic values and human wellbeing considerations. First, eutrophic bodies are widely considered ugly to look at compared to oligotrophic bodies. They also may support fewer charismatic organisms, and charismatic organisms are more difficult to observe due to water turbidity. Some environmental philosophers have previously noticed the connection between aesthetic properties and the health of ecosystems, although this connection has been subject to criticism.⁶⁷ In addition, humans seem to have a preference for ecosystems that contain greater numbers of trophic levels, particularly when they include large animals. Eutrophic bodies are trophically simplified, which plausibly contributes to the consensus that they are degraded. The valuation of trophic complexity may be explained in

⁶⁷ There are particularly concerns that thinking about nature's aesthetic value solely in terms of scenery does not capture the full scope of aesthetic value (Saito 1998). There is room for discussion about whether our assessments of ecological health ought to be sensitive to popular aesthetic judgments or whether they should be based on the more informed views of those who work in environmental aesthetics.

various ways, but it seems to partially amount to an aesthetic preference for ecosystems with large organisms. It is likely that maintaining trophic diversity is good for other reasons. Ecosystems which are extremely trophically simplified may be more susceptible to collapse or to runaway extinctions. However, this is not a danger in the pond case, since eutrophication results in a stable ecosystem state.

Next, concerns about human wellbeing are plausibly relevant to our judgment that certain ecosystem states are unhealthy, such as the eutrophic pond. Some types of algae bloom are toxic to humans and other animals, which can make eutrophication a direct risk to human health. Even in cases where toxicity is not a concern, eutrophic bodies are perceived to be less suited to many forms of outdoor recreation. To give a different example, assessing the health of forests may involve considering their effects on drinking water quality, air quality and other local climatic effects, considerations which pertain directly to human welfare. Human wellbeing considerations, as well as aesthetic considerations, may be particularly relevant to judgments about the health of urban ecosystems.

The wellbeing considerations which are relevant to ecosystem health are relatively diffuse effects of ecosystems on the environment, such as their contributions to nontoxic or toxic water and their suitability for certain forms of recreation or nature appreciation. Sometimes, a direct link has been suggested between ecosystem health and the provision of services. Ecosystem services are goods and benefits provided by ecosystem processes to humans, which include contributing to clean water but also more concrete goods such as timber and food. However, the ability to provide services does not always track health judgments. Agricultural monocultures are very efficient in the short run at providing specific services, but they are usually considered unhealthy qua ecosystems. Ordinarily, managing a forest for a specific kind of timber production is seen as

distinct from managing it for ecological health. Thus, while recognizing that certain kinds of diffuse service are relevant to judgments about ecosystem health, it is best not to define ecosystem health in terms of the ability to provide services, since this would conflate two distinct kinds of management goal.

Consider now why the farmed state (3) of the pond is not as ecologically healthy as the oligotrophic state. History has some intuitive relevance here, since farming has presumably caused a recent and major qualitative change from past states of the pond. Some might also think that biodiversity is relevant to this judgment, since the farmed pond has decreased animal diversity. However, species diversity alone is unlikely to reliably track judgments about ecological health. This is because, first, some types of ecosystem are naturally far more diverse than others, and second, changes in species diversity have very complex, uneven, and sometimes negligible effects on ecosystem properties (as discussed elsewhere in the dissertation). In our example, the eutrophic pond has less diverse macroscopic animals but may have a great diversity of algae and microbes. So, while some features of ecological organizational complexity are relevant to health judgments, the relationship between biodiversity and health will be complicated.

Some may also have the intuition that the farmed state is less healthy because it gets the most human intervention. However, it is also possible that intensive management has been used to maintain the oligotrophic state. For reasons discussed previously, we should be wary of the view that human intervention has some inherent bearing on ecological health. In addition, notice that the intuition about the farmed state can be restated in terms of resilience. Systems which need very intensive management typically need that management because they lack resilience and cannot self-sustain. So, we can recapture the intuition as having to do with compromised resilience rather than the occurrence of human management as such.

This leaves us back with accounting for why the farmed state is considered less ecologically healthy to begin with. It has partly to do with the aesthetic preference for trophic complexity as already mentioned, but there seems to be another factor related to history. The relevance of history to the farmed pond can be accounted for by considering the value of authenticity.⁶⁸ Some changes to ecosystems are viewed as compromising their health because they are perceived as artificial and as cutting the system off from ecological history. The ecosystem health concept permits ecosystems to change “naturally” over appropriate timescales. However, sudden and major qualitative changes to ecosystems, especially when they arise due to technological influences or poorly motivated intentional human actions, are perceived as disrupting the health of ecosystems because they render the systems artificial or inauthentic. This value helps to explain why we consider it unhealthy when a historically forested area converts to prairie due to increased browsing pressure. The perceived problem is not that we think prairies are defective ecosystems (since we want to restore prairies in other regions), but that having a prairie in a forested region seems inauthentic. Thus, intensive human management to restore the ecosystem to the forested state is thought to improve the health of the ecosystem.⁶⁹

Again, my claim is that authenticity does a good job of tracking and explaining some features of common judgments about ecological health. Moreover, if authenticity is relevant to these judgments, then it explains why historical fidelity and human intervention appear relevant to

⁶⁸ Newman et al. (2017) discuss authenticity, p. 363ff. Most recent discussions of authenticity in conservation ethics have been about de-extinction and genetic technology rather than ecosystems.

⁶⁹ This view is expressed frequently in scientific publications on overbrowsing and forest health. See Noonan et al. (2021) for an example. Interesting parallels might be drawn between attitudes towards forest restoration and art restoration.

some judgments about ecosystem health but not others. For instance, human actions which are meant to restore ecosystems are compatible with their health, but actions which cause the ecosystem to rapidly change to a novel state compromise health (particularly if the new state is trophically impoverished or has diminished resilience).

So, to summarize my proposal, a healthy ecosystem exhibits structures and functions which are judged to be an authentic reflection of its location and history; is in a state which is conducive to human wellbeing via its effects on the environment and is aesthetically pleasing; and is resilient and sustainable without intensive intervention.

My sense is that satisfying all of these factors is sufficient for health, but these conditions may not all be necessary in all cases. For instance, assessments of agricultural ecosystems may drop the authenticity requirement and place greater emphasis on specific human wellbeing requirements. For ecosystems which have not historically been considered scenic (like wetlands), greater emphasis may be placed on wellbeing and authenticity than on aesthetic value.⁷⁰

At this point it should be clear that some relevant values are culturally relative—meaning that either their recognition or interpretation will be sensitive to human historical and cultural context. This is an intended feature of my account. Views about ecosystems and their management changed dramatically throughout the twentieth century to the present and have varied across human cultures. It is no surprise if the way we assess the health of ecosystems now differs from how it was assessed in, say, Frederic Clements' time. An implication is that our conception of ecosystem health is not set in stone. For instance, if you share the concern that judgments about authenticity may be too backward-looking as we manage ecosystems in the context of a changing

⁷⁰ However, wetlands may also be subjects of “aesthetic rehabilitation” (Callicott 2003; Rolston 2000).

climate, then you might think we should somewhat alter how we evaluate health. For present purposes, my argument is that shared judgments about authenticity are currently central to assessments of ecological health in contemporary ecology, environmental studies, and ethics.

My proposal overlaps a fair bit with the definition of health developed by a team of marine ecologists in Tett et al. (2013). The ecologists' definition gives more detail about relevant components of the structure and function of ecosystems, while mine is more explicit about the values used to make judgments about "appropriateness." For comparison, their definition in full is stated below. I have added emphases where the ecologists use terms indicating that a non-empirical judgment has been made.

Ecosystem health depends on: the physiological health of the constituent organisms; the *characteristic* properties and interactions of the species present; and the emergent properties of the system comprising the biota and their environment. Healthy ecosystems can sustain services to humans. They are vigorous, resilient to externally imposed pressures, and able to maintain themselves without human management. They contain organisms and populations that are free of stress-induced pathologies, and biodiversity that includes (1) a functional diversity enabling all biogeochemical and trophic functions *appropriate* to the ecohydrodynamic conditions, and (2) a diversity of responses to external pressures [note: these properties would be expected to contribute to resilience]. All *expected* trophic levels are present and well interconnected, and there is good spatial connectivity amongst subsystems. (Tett et al. 2013, 12)

Notice that their account picks out: structural and functional properties "characteristic" of the ecosystem; the provision of services; and resilience, along with properties which are likely to contribute to resilience. This largely agrees with my own assessment of ecosystem health. I would add two things. First, the judgments about "characteristic" properties seem to be based on the value of authenticity to a large extent. Second, I would urge that "services" be interpreted in a broad

sense to include aesthetic, recreational and environmental effects, and that health is not contingent on the ecosystem providing specific goods in any particular quantities.⁷¹

A second related account of ecological health can be found in Roberta Millstein's reconstruction of Aldo Leopold's view on land health. According to Millstein,

(healthy) functioning for Leopold consists in 'characteristic' species in 'characteristic' numbers performing interconnected 'characteristic' functions, [...] in such a way that energy and matter continue to flow through the soil and through the other parts of the land community, with longer food chains tending to maintain cycling better by avoiding erosion, so that the land community can continue to support life over time. (Millstein 2020b, 1116)

Notice that this definition picks out, first, characteristic structures and functions, and second, features of the "land community" (a special term used by Leopold which has a meaning similar to "ecosystem") which will tend to make it sustainable or persistent. Millstein's definition does not specifically mention humans or human values, but it seems that Leopold did consider humans to be part of the land community (see Millstein 2018, fig. 13.1). So, it seems like a sophisticated reading of Leopold's views on land health is quite consistent with contemporary thought about ecological health. Although Leopold's views were specifically about the health of the land, similar considerations about composition and sustainability are relevant to assessing the health of any type of ecosystem.

In addition, my account has similarities to the account of ecosystem regeneration and degradation developed by Elis Jones (2021)—concepts which are related to that of health. Jones argues that regeneration efforts rely on choosing an appropriate benchmark state for the

⁷¹ For context, the quoted authors are thinking about marine systems, where fisheries collapse has been an ongoing problem. So, they are likely thinking that healthy marine systems are able to maintain game fish populations, which I agree with. My remarks are clarificatory rather than objections.

undegraded ecosystem, and that choosing this benchmark relies on judgments about the charisma, instrumental value, and freedom from human influence of certain reef states. The three values Jones identifies are similar to the three values I argued influence health judgements.⁷² This should be unsurprising given that restoring ecosystem health is usually an aim of ecosystem regeneration efforts, and that a degraded ecosystem state contrasts with a healthy one. A main difference between our accounts is that Jones argues that freedom from human influence is relevant to benchmark selection for at least some regeneration purposes. I agree with this as a descriptive matter, but have argued that the occurrence or lack of human influence per se is not a reliable indicator of ecosystem health. Jones also allows that coral reefs can occur in less desirable states in the absence of human influence.

To conclude this section, notice that our conception of ecosystem health has two components. One component pertains to the current state of the ecosystem, i.e. its structure and processes, aesthetic value, and so forth. The second component pertains to the sustainability of that state. At minimum, an account of ecosystem health needs to include both components. An account which omits the first component cannot distinguish between healthy and unhealthy alternative regimes in many ecosystem types. There are fewer accounts which completely omit the second component, given the centrality of resilience and sustainability in current thought about environmental management.

⁷² I had arrived at my account of health prior to reading Jones' paper, so the convergence may lend some extra plausibility to our accounts.

4.5 Arguments for Eliminativism

In this section I will consider two major objections to the ecosystem health concept. These objections aim to show that we should eliminate the concept from our discourse. The first objection is the “ecosystem antirealism objection”: the argument that ecosystems are not real entities, so it does not make sense to attribute states of health to them. The second objection is the “organicism objection”: the claim that the idea of ecosystem health relies on a faulty comparison of ecosystems with individual organisms. I summarize these objections and show that both of them fail. This does not by itself show that the concept of ecosystem health is a good one, but that some of the frequently cited reservations about it are unfounded.

4.5.1 The Ecosystem Antirealism Objection

Various philosophers have argued that ecosystems are not objectively real entities in the world. Recently, Newman, Varner and Lindquist (2017) have suggested that if ecosystems are not real entities, then it does not make sense to attribute health states to them:

[Given that] ecosystems are not real physical entities, how can they be healthy or not? ... And if ecosystems are *arbitrary constructs*, what reason is there to suppose that it matters whether or not they are “healthy,” assuming a non-arbitrary definition of health could be specified? (290, emphasis in original)

The authors do not consider these questions in great detail, because they arise in a chapter whose primary purpose is to evaluate ecologist ethics⁷³ rather than the ecosystem health concept.

⁷³ I.e., the view that ecological wholes are directly morally considerable.

However, it is clear from the surrounding discussion that the authors doubt the questions can be given satisfactory answers. I attempt to address their concerns here.

In order to evaluate this kind of objection, we should first reconstruct the argument for antirealism about ecosystems in order to clarify what this position amounts to. In the passage I quoted from, the authors appeal to the following works which provide arguments against realism about ecosystems: (Jax 2006; 2007; Garcia and Newman 2016).

Garcia and Newman (2016) make an argument against realism based on the observation that there are multiple non-equivalent definitions of “ecosystem.” They then make the following claims about these multiple definitions, which they claim provide a weak basis for realism about ecosystems.

- Several ecosystem definitions are intended to be antirealist or do not have a clear ontological interpretation (173).
- Choice among definitions tends to be based on ecologists’ research interests (174).
- Empirical evidence underdetermines the ontology of ecosystems (174).

They also make some further claims which they argue generate practical problems for considering ecosystems real entities, including:

- Different ecosystem definitions carve up ecosystems in conflicting ways, so there are very many distinct (and potentially overlapping) kinds of ecosystem (“disparity thesis,” 178).

I will call this the **conflicting definition argument** for ecosystem antirealism. The major concerns of this argument seem to be that ecosystem definitions are (a) somewhat arbitrary, unprincipled, or dependent on human research interests; (b) do not support a single clear ontology, or support

conflicting ontologies; and (c) carve ecosystems up in too many different ways, so we would need to accept the existence of an unreasonable number and variety of ecosystems.

The conflicting definition argument relies on some implicit premises about what our definitions of real objects should be like. Authors who raise definitional concerns about ecosystems seem to assume that we should be able to provide a single definition for a type of object, or a small number of consistent definitions. They may also assume that scientific definitions of objects should be based on principled grounds other than research interests. They additionally seem to assume that our ontology should be relatively sparse—i.e., we should be wary of accepting large numbers of cross-cutting objects.

There is a distinct kind of popular argument against ecosystem realism which I will call the **unclear boundary argument**. While the conflicting definition argument targets ecosystems at the type level, the unclear boundary argument targets ecosystems at the token level. This argument claims that ecosystem boundaries are vague or arbitrary, that ecosystem components are constantly changing over time, that ecosystems are not sufficiently closed, “integrated,” or internally structured systems, and so forth.⁷⁴ Again, this argument makes some implicit assumptions about what real objects are like, e.g. that they have relatively constant or clear boundaries or have the right kind of internal structure.

I will respond to both of these arguments together. First, note that the arguments both involve substantive assumptions about what counts as a real object. Most authors who have discussed ecosystem realism do not provide much argumentation for these assumptions, which

⁷⁴ See Millstein (2018) on boundary issues for ecosystems. There are related discussions about boundaries and ontology focused on communities (Sterelny 2001; Odenbaugh 2007; Lean 2018).

might be objected to. However, suppose for the sake of discussion that we accept the assumptions of either argument about the criteria for counting as a real object.

It should be apparent that if you take real objects to possess the properties assumed by either the definition or boundary argument, then many entities studied by biologists are not real objects. Notice that the definition and boundary arguments can both be applied with little modification to many organisms, genes, persons and species, among other items. So, if you accept either of the arguments for ecosystem antirealism, you should also be an antirealist about these other objects.⁷⁵ I have no objection to this conclusion (appropriately interpreted), but I suspect some authors may not intend to support such a sweeping position.

A more moderate position is to treat individuality as a matter of degree. If some precise conditions for being an individual are spelled out, reasonable arguments can be made that some entities should be treated as individuals of particular types and others as collectives of particular types (and perhaps, some entities can be treated as both) (e.g., Huneman 2011; Lean 2018). But in my view, this is a matter of model choice or representation rather than a matter of realism. To give a simple illustration, my coffee table does not fulfill the criteria for being an evolutionary individual, but that has no bearing on whether or not it is real. Similarly, the fact that ecosystems might fail some criteria for being a particular type of individual does not entail that they fail to be real. The basic problem is that positions about realism *tout court* cannot be read off of individual cases of model choice or representation.

⁷⁵ In my view, the common assumption that organisms are inherently more cohesive objects than ecosystems is probably a result of our perceptual bias towards things of approximately our own size, and the related tendency to think of large vertebrates as representative forms of life. I'm not aware of any convincing arguments that ecosystems in general are ontologically "worse" objects than organisms in general.

In any case, we need to consider whether it follows that ecosystems cannot have health properties if they are not real objects according to the given arguments, or according to any other antirealist view. A general appraisal of the implications of antirealism is outside the scope of this chapter. For the purpose of my argument, however, observe that being an antirealist about ecosystems or other objects is not the same as the view that we cannot have warranted scientific discussions about the things we call “ecosystems.” Most antirealists want to maintain a distinction (even if the distinction is one of degree rather than kind) between terms which are scientifically useful (“atom,” “species,” “the Greater Yellowstone Ecosystem”), versus terms which lack a current scientific application (“aether,” “Heffalump,” “Spock”⁷⁶). Even if both classes of entities include items which are not strictly real according to one’s metaphysics, it is clear that the former group can be scientifically studied and ascribed properties by scientists.⁷⁷

So, although ecosystems are messy entities with interchangeable parts, somewhat arbitrary boundaries and multiple conflicting definitions, there is nothing *prima facie* wrong with attributing properties to them. Organisms also suffer from notorious boundary and definition problems, but we have no problem with saying that humans, cats and even houseplants are healthy. Thus, there is nothing special about health states such that only “really real” objects can be healthy, supposing there is such a thing as a really real object. Positions about ecosystem realism therefore have no bearing on whether or not ecosystems can be healthy.⁷⁸

⁷⁶ Referring to the science fiction character rather than to any actual individual who may share this name.

⁷⁷ It seems we can also attribute properties to objects of the second class, but deciding what properties should be ascribed to them is outside the purview of science.

⁷⁸ Also see McShane (2004, 241 ff) for a related argument.

If this is not already plausible, consider parallel cases of ascribing properties to constructed entities. For example, societies and corporations are obviously human constructs, but we can still make both factual and evaluative claims about them. Hopefully, no one would argue that we should not be concerned about corporate corruption because corporations are not objectively real objects.

As an aside, it is sometimes assumed that ecosystems must be objectively real entities in order for ecoholism (i.e., the view that ecological wholes are directly morally considerable) to be a viable position (Garcia and Newman 2016). My argument in this section suggests a *reductio* against this claimed link. The same reasons that have been given to support antirealism about ecosystems also support antirealism about individual organisms and persons. So, if there is a link between counting as an objectively real object and being morally considerable, then organisms and persons are not morally considerable. To avoid this conclusion, we should either accept that we can have moral obligations towards entities which are partly constructed, or we need to finesse our understanding of the criteria for counting as a real object.⁷⁹

4.5.2 The Organicism Objection

The second objection raises the concern that talking about ecosystem health commits us to a dated organicist view of ecosystems. Some have connected the concept of ecosystem health with

⁷⁹ Millstein (2018) also discusses ontology and moral considerability. Chapter 5 below contains some further discussion of ecoholism (a position which I reject for reasons unrelated to the present discussion).

(a popular conception of) the Clementsian superorganismic view of ecosystems.⁸⁰ Here are several examples of this objection:

Because health is a property of organisms, “ecosystem health” implies that ecosystems are super-organisms. This implies in turn that ecosystems have consistent structures, regular development, homeostasis, tight integration, and distinct identities. This vision of ecosystems as superorganisms was implicit in Frederick Clements’s theory of ecosystem development[....] (Suter 1993, 1533–34)

The phrase ecosystem health is based on an invalid analogy with human health requiring acceptance of an optimum condition and homeostatic processes maintaining the ecosystem at a definable optimum state. (Wicklum and Davies 1995, 997)

The idea that an ecosystem can be healthy or sick (in either a binary way or on some sliding scale of degrees of health) presupposes a biological structure of an ecosystem, analogous to that of an organism, which a professional medic might diagnose as sick or healthy based on metabolic indicators[....] (Lancaster 2000, 214)

Some authors, such as Lilly-Marlene Russow, seem to be concerned about both organicism and ecosystem realism together. Russow worries that ecosystems are not similar to organisms because they lack clear boundaries:

The danger here is that [human] patients are more-or-less self contained organisms: they have skin, a dividing line between them and the rest of the world. This often provides a foundation for identifying what is good for them as individuals, perhaps distinct from what is good for the family or community in which they are located, or good for the bacteria that have taken up residence in their gut. We cannot think of ecosystems in this way because they are both nested and open-ended systems. (Russow 1995, 364)

The simplest response to this objection is that “health” need not mean the same thing when applied to humans and ecosystems in order to be a legitimate concept. Concepts can have multiple uses and meanings, and applying a concept to different kinds of entity does not require that the

⁸⁰ Early Clements talked about “formations,” which are more like our concept of a community than an ecosystem. His superorganismic view was intended to be a metaphor for community succession rather than a model of the ecosystem. However, his view has been widely (mis)interpreted as a view about ecosystem ontology.

entities be otherwise comparable. Talk about ecological health is now well-established in ecology and environmentalism, so we should try to understand how the concept is actually being used in this context. As I sometimes alluded to in my above discussion, how we think about ecological health is in fact not very similar to how we think about human health.

It is also apparently false that working on ecological health requires thinking of ecosystems in terms of an organicist model. The organicist view of communities was a view primarily about succession, succession being one of the earliest described ecological phenomena.⁸¹ But various more recent developments in ecological theory are relevant to current thought about ecological health, including trophic dynamics,⁸² Island Biogeography (MacArthur and Wilson 1967), alternative state models, landscape and metapopulation-level insights, and the concepts of ecosystem functions and services. It is very implausible that most people working on ecological health have abandoned a century of developments in ecological theory and reverted to the organicist view of communities as end-directed developmental units. In short, this common concern has little basis in the actual contemporary work on ecosystem health.

Although I emphasize that “health” as applied to ecosystems is not identical in meaning to “health” as applied to organisms, I want to resist the view that ecosystem health is a metaphor. For example, Callicott (1995) takes the position that health is predicated literally of organisms and metaphorically of ecosystems. However, taking this position requires that “health” apply in a more objective sense to organisms, and that when we talk about ecosystems’ health we are thereby

⁸¹ The phenomenon of ecological succession was clearly described as early as 1808 by Alexander von Humboldt in *Ansichten der Natur*, although of course there was not yet a term for the process (Humboldt 2014). The theory of succession was a major focus of the new discipline of ecology around the turn of the twentieth century.

⁸² Trophic levels were originally described by Elton (1927).

comparing them with organisms. These are both problematic positions. First, as Callicott acknowledges, judgments about the health of organisms are themselves sensitive to human values. If judgments about both organism and ecosystem health are sensitive to human values, then there is no basis for thinking that organism health is more objective. Second, as I have argued, talk of ecosystem health does not need to presuppose that ecosystems are similar to organisms in any particular way. Conceiving ecosystem health as a metaphor only invites this comparison. Because of the many differences between ecosystems and organisms, it is better to conceive ecosystem health as a distinct concept which in the details has little to do with organism health.

4.6 Arguments for Ecosystem Health

Above I responded to two common arguments against the ecosystem health concept. In this section, I discuss some arguments in favor of the ecological health concept. I reject the first argument I consider before presenting some considerations of my own.

4.6.1 Argument from Intrinsic Value

Various environmental philosophers, including Callicott (1992), Sagoff (1992), and Newman et al. (2017, chap. 9), have connected ecosystem health with intrinsic value. The general idea is that if ecosystems have intrinsic value, then we have an obligation to consider their interests,

and that considering the interests of an ecosystem amounts to considering its health.⁸³ In order to evaluate the merits of this argument, I provide my own reconstruction of the argument here which makes explicit all of the implicit steps.

1. An entity has intrinsic value if and only if we are morally obligated to consider its interests.
2. If we are morally obligated to consider an entity's interests, then the entity must have (knowable) interests. (*ought implies can* principle)
3. So, if an entity has intrinsic value, then it must have (knowable) interests. (1, 2)
4. Ecosystems have intrinsic value. (assumption)
5. So, ecosystems have (knowable) interests. (3, 4)
6. All biological entities with (knowable) interests have a (knowable) interest in remaining healthy or avoiding dysfunction.
7. Ecosystems are biological entities.
8. So, ecosystems have a (knowable) interest in remaining healthy or avoiding dysfunction. (5, 6, 7)

The argument can be truncated at this point. Presumably, if ecosystems have an interest in remaining healthy, then we can conclude that it is appropriate to describe ecosystems in terms of their health. Thus, the ecological health concept is vindicated.

⁸³ Callicott and Sagoff run this argument in the forward direction, i.e. from the claim that ecosystems have intrinsic value to the claim that they have interests. Newman et al. consider and reject an argument in the opposite direction, i.e. they think you have to first show that ecosystems have interests to justify the claim that they have intrinsic value. I'm considering the argument in the first direction because the point here is to try and establish that ecosystems have states of health.

The explicit version of the argument makes it easy to identify its fragile points. The first fragile point is the first premise, which claims that intrinsic value is tied to fairly strong moral obligations. Along with this is the assumption that ecosystems have intrinsic value. It has been notoriously controversial whether ecosystems have intrinsic value, and what it amounts to if they do. Some interpretations of intrinsic value do not support premise (1).⁸⁴ For present purposes, I will simply note that these are controversial assumptions, and the argument is only as compelling as the assumptions.

Another fragile point is premise (6) which allows us to conclude that an ecosystem's interests coincide with states of health. Sagoff, for instance, seems to assume this as a conceptual point—i.e., that “health” just refers to whatever states are inherently good for living systems. Although I believe it is appropriate to ascribe states of health to ecosystems, we should not grant (6) as a premise. The problem is that living systems have many properties which they might have an interest in maintaining, such as resilience, sustainability, fitness, ability to persist, and so forth. It is not obvious as a conceptual point that for any given living system, it has an interest in maintaining health rather than one of these other properties. *Prima facie*, for example, a plant might have an interest in staying healthy but it alternatively might have an interest in reproducing or in simply staying alive (assuming for now that plants have interests). Similarly, an ecosystem could have an interest in staying healthy, or it could have an interest in expanding, persisting, or being able to adjust to climate change. Thus, even granting that all biological entities have interests does not entail that it is appropriate to describe those interests in terms of health.

⁸⁴ Chapter 5 includes a discussion of different views about the intrinsic value of ecological wholes. If you accept what I term a *weak* conception of intrinsic value in that chapter then you might reject premise (1).

For these reasons, this argument strategy is unconvincing as a defense of the ecological health concept.

4.6.2 Argument from Entrenchment

A major consideration in favor of retaining an ecosystem health concept is that this concept is already entrenched in the discourse about ecology and conservation. Trying to eliminate the concept would be impractical, so it should require a fairly strong argument that the concept is not just (perhaps) unclear or messy but positively harmful and misleading. The ecosystem health concept could be suboptimal in some ways, but the fact is that most sciences and human projects work with suboptimal, historically contingent concepts. It is more useful for philosophers to attempt to clarify existing concepts rather than arguing for the elimination of existing concepts without offering a good replacement. Exceptions occur when the use of a concept is harmful to people or unethical (e.g., by furthering harmful stereotypes), or when a concept is so confused it demonstrably impedes research goals, but no such case has been made convincingly for the ecosystem health concept.

4.6.3 Argument for a Plurality of Conservation Goals

A second consideration that favors retaining an ecosystem health concept is the need to distinguish health as a conservation target from other targets such as ecosystem services and biodiversity protection. Managing an ecosystem for its health roughly means managing it to exhibit structures and functions which are reflective of its location and history, to contribute to human wellbeing, and to be aesthetically pleasing and resilient. This is related to but distinct from other

projects such as managing an ecosystem to maintain as much native biodiversity as possible; to provide a specific service such as timber; or to approximate as closely as possible the prior composition of the system. Gillette et al. (2021) argue that the rise of the ecosystem health concept in the late 1980s represented a shift from a “compositional restoration” approach to ecosystem management. Assuming that it is sometimes reasonable to consider managing for health in contrast to these other goals, we need the health concept in order to talk about the different options. Importantly, the fact that there are no naturalistic answers to the status question means that we cannot appeal to an alternative natural property—e.g. functions, resilience or historical fidelity—as a replacement for *health*. Managing an ecosystem to promote just these properties is distinct from managing it to promote health. Elimination of the health concept would invite confusion among these different potential conservation goals.

4.6.4 Argument from the Microbiome

Although I will not develop this argument in detail here, I want to mention a third avenue in support of an ecosystem health concept which may be considered in future work. My dissertation focuses on large ecosystems including forests and lakes. However, microbial ecosystems will also provide insights about health concepts. In particular, various philosophers and biologists have thought that it is possible to treat the human microbiome as either part of the human individual (the holobiont view) or as composing an ecosystem (Inkpen 2019). In my preceding discussions I assumed that (macro) ecosystems cannot usefully be treated as organisms. However, if there are some entities that can be treated as either individual organisms or ecosystems, then this further bolsters the case for an ecosystem health concept. The argument, in outline, would say that individual organisms are healthy or unhealthy; that some individual organisms can also be treated

as ecosystems; so some ecosystems are healthy or unhealthy. This line of argument further erodes the frequently assumed ontological distinction between ecosystems and organisms. It remains to be considered whether a different conception of health would apply when considering an individual qua organism or qua ecosystem.

To fend off an immediate concern, much of the discussion of holobionts has focused on whether they are evolutionary individuals, i.e. whether they are units of selection (e.g., Skillings 2016). Even if holobionts are not evolutionary individuals, I would argue that this is irrelevant to whether they have states of health. In my view, ecosystems are individuals (in a generic sense—they are individual objects), but conventional macro ecosystem types are usually not usefully conceived as evolutionary individuals. As I have argued in this chapter, you do not need to appeal to selected functions in order to make sense of an ecological health concept, and therefore the health concept can apply to things which are not evolutionary individuals. Even for things which are evolutionary individuals, we should be wary of defining health primarily in terms of carrying out selected functions, since selection only cares about reproductive success in a given environment, not about wellbeing, efficiency, aesthetic value, or other qualities we associate with health.

4.7 Conclusion

To conclude, there is room for further consideration of the values underlying the ecological health concept—including aesthetic values, authenticity, and human wellbeing considerations. It is especially noteworthy that ecosystem health-focused conservation is often taken to be an alternative to more anthropocentric resource-focused conservation methods; yet ecosystem health

judgments are sensitive to anthropocentric values. Those who are less optimistic about current uses of the ecosystem health concept can contribute to change by criticizing the values at work and suggesting alternatives. In the meantime, ecosystem health researchers (in both science and philosophy) should be committed to transparency about what kinds of value judgment they are making (Gillette et al. 2021). Finally, conservation priorities based on properties like “intactness,” lack of human influence, or (mere) historical fidelity should be considered very carefully since I have argued that these properties are not closely connected to the health of ecosystems.

5.0 Values of Biodiversity at the Ecosystem and Global Scales

5.1 Introduction

Humans are responsible for an accelerating rate of species extinctions. According to the most recent IPBES report, on average, about a quarter of species are threatened with extinction within each of the best-studied taxa (Díaz et al. 2019). While there is general agreement that we should try to prevent species losses, philosophers have struggled to explain how this is justified. In particular, there has been a persistent divide between environmentalists who favor instrumental versus intrinsic value defenses (e.g., the varying perspectives of McCauley 2006; McShane 2007; Justus et al. 2009; Odenbaugh 2020; Newman 2020).

This chapter adds to the debate by arguing that a popular form of instrumental argument, the Ecosystem Services Argument, is unsuccessful and the core reason for its failure has been overlooked in the literature. The Ecosystem Services Argument (actually a cluster of related arguments) claims that biodiversity should be protected because it supports important ecosystem functions and services. Especially in popular media, this kind of argument is commonplace and has not always been subject to enough scrutiny. If this form of argument fails, then environmental philosophers must remain open to considering defenses which appeal to other kinds of value.

The general objection I develop to this argument form is a *scale problem*, which points to a scale mismatch between ecosystem services and the intended target of the argument. Most iterations of the Ecosystem Services Argument point to goods, functions or processes occurring within ecosystems which are putatively supported by local diversity levels. There have been relevant prior discussions about the complicated empirical support for a connection between local

diversity and ecosystem functioning or services (e.g., Newman et al. 2017). This chapter develops a distinct problem for the Ecosystem Services Argument, and will show that this problem is even worse. The problem arises because the argument can only support protecting global diversity if global species losses causally contribute to local species losses. But this step in the argument (which has rarely been discussed among environmental philosophers⁸⁵ or even made explicit) also lacks support. I present conceptual and empirical reasons that undermine the causal connection between global and local species losses presupposed by the argument. So, even if appealing to services can justify protecting local species diversity, this line of argument fails to scale up to justify protecting global biodiversity. Furthermore, the empirical evidence linking species richness to even local ecosystem functions is shaky. Some of the relevant empirical evidence behind these claims will be reviewed in this chapter.

A distinct way to put my conclusion is as follows. It is usually taken for granted that biodiversity has some instrumental value to humans. Environmentalists who object to the Ecosystem Services Argument usually claim that something is wrong with justifying the preservation of nature on instrumental grounds (e.g., McCauley 2006), but this is not my position. In contrast, I argue that global biodiversity—i.e. the total number of species on earth, distributed in roughly their native ranges⁸⁶—does not have any instrumental value to humans, or at least none that is currently plausible based on existing scientific research. As a result, arguments which appeal to standard instrumental values cannot possibly justify global biodiversity preservation.⁸⁷ To be

⁸⁵ Though see Desjardins et al. (2019).

⁸⁶ The “native ranges” clause is to prevent species maintained only in zoos or nonnative species grown annually in gardens from counting towards biodiversity totals in the present context.

⁸⁷ Section 5.6.4 below does some finessing of this claim.

clear, my contention is not that we should reject instrumental arguments in general, but that Ecosystem Services Arguments specifically are not successful as biodiversity defenses.

I must briefly address the nature of biodiversity before discussing its value. First, the aim of this chapter is primarily to consider arguments for species conservation. For this reason, “biodiversity” should be read as species richness unless specified otherwise. Second, for convenience, I will often write as if biodiversity were an entity, but I assume that biodiversity is instead a property (or cluster of properties) belonging to any given ecological system or area. Thus, strictly speaking, we are considering the value of the property of *being biodiverse*. Historically, environmental philosophers have been wary of the view that biodiversity has intrinsic value—preferring to ascribe intrinsic value only to objects—but have assumed that biodiversity has instrumental value to humans (e.g., Oksanen 1997). If my discussion shows that global biodiversity does not have instrumental value to humans—or at least that the usual arguments for its instrumental value are unsuccessful—this suggests that a major reorientation of the discourse is in order.

I also stress that my goal is not to undermine biodiversity conservation. Rather, my goal is to encourage further reflection on popular arguments made for saving species.

This chapter is organized as follows. First I overview the existing types of biodiversity defense. Then I present the standard Ecosystem Services Argument and review some empirical problems with its claims. Next I present my scale problem in more detail. After that I consider some alternative formulations of services arguments, which I argue are not much more successful than the original version. I conclude the chapter with a discussion of some alternative biodiversity defenses which may avoid the scale problem.

I suggest three upshots to my conclusion that ecosystem services defenses of biodiversity are unsuccessful. First, it is worth devoting more scientific and philosophical attention to potential roles of species diversity at the global scale. For instance, more knowledge of possible feedbacks or tipping points at the global scale resulting from interactions among climate change, ecosystem changes and species extinctions would help to inform conservation decision-making. Second, philosophers should remain open to considering biodiversity defenses which appeal to other kinds of value, including intrinsic and aesthetic values. Finally, a general moral of this discussion is that it is important for environmental philosophers to pay close attention to scale, since most ecological processes and patterns are scale-dependent, and neglecting this can lead to errors like the one committed by the Ecosystem Services Argument.

5.2 Preliminaries

There are notoriously many arguments in support of protecting biodiversity.⁸⁸ The main argument strategies I identify are:

- **Intrinsic value arguments:** claim that we ought to preserve biodiversity for its own sake; that we have moral duties towards nature; or that biodiversity or its existence is valuable as an end of its own.

⁸⁸ A survey of biodiversity defenses is provided by Howard et al. (2018); various arguments are reviewed critically by Maier (2012, chap. 6) and Newman et al. (2017).

- **Services arguments:** claim that biodiversity is instrumentally important because it contributes to ecosystem functioning, climate regulation, sustainable food resources, pharmaceutical discoveries, water quality, recreational industries, etc.
- **Risk avoidance arguments:** claim that biodiversity should be protected because biodiversity loss may compromise services in the future; that biodiversity loss will result in losses of future options; that biodiversity loss will compromise ecosystem resilience or adaptability; that biodiversity and genetic loss will compromise evolutionary adaptation and speciation; or that biodiversity loss should be stopped to prevent unknown future harms. These are usually instrumental arguments but they differ in form from standard services arguments by the introduction of one or more premises asserting a precautionary principle, the value of future options, etc.
- **Experience value arguments:** claim that biodiversity is aesthetically valuable; that biodiversity contributes to transformative human experiences; that biodiversity promotes human closeness to nature; that native biodiversity is valuable as a component of local cultural inheritance; or that scientific knowledge of existing biodiversity is valuable. Note that environmentalists disagree about whether aesthetic value is a species of intrinsic or instrumental value. Because this is a matter of disagreement, and because these arguments differ sufficiently from “standard” intrinsic and instrumental arguments, they are worth treating as a distinct group.
- **Arguments for instrumental value of the biodiversity concept:** finally, multiple commentators on my work have suggested that our having a *concept* of biodiversity may have instrumental value—e.g., it may promote concern about the composition of the environment, it may aid in communication about environmental problems, or it may

be integral to the design of ecological research. This claim should be clearly distinguished from the claim that biodiversity itself has instrumental value. This type of argument may represent a response to biodiversity eliminativism—i.e., the view that we should reject *biodiversity* as an appropriate concept guiding research or management discourse (Santana 2018). However, arguments from the value of the biodiversity concept cannot support efforts to protect biodiversity. To give a silly illustration, it is useful to have a concept of *train derailment*, but that does not mean we should promote train derailments. For this reason, I will not consider this particular argument further in this chapter. However, I agree with the insight that we should also think carefully about what concepts are used to communicate about environmental objectives.

The following discussion will focus on arguments that fall under the services and risk avoidance categories, though I will more briefly discuss intrinsic value and experience value arguments at the end of the chapter.

There are at least three types of consideration when evaluating arguments in environmental ethics. One consideration is the scientific support for empirical claims or factual predictions of an argument. Another is the moral acceptability of premises appealing to ethical principles or claims. A third is rhetorical effectiveness—whether the argument is likely to be compelling or to sway public opinion. My discussion of the Ecosystem Services Argument in this chapter focuses on the first issue. I identify an implicit premise which is needed for the argument to go through and argue that, contrary to what might be expected, scientific evidence supports the negation of this premise. Given reasons to think that a key premise is false, I assume the other two considerations are moot. However, I do not think philosophical arguments always need to rest on empirical premises which

have clear scientific evidence in their favor. Interesting arguments often utilize uncertain premises. That is not the issue with the Ecosystem Services Argument, where the premise at issue is not uncertain or questionable but apparently just false.

In a few places I will discuss questions of rhetorical effectiveness when commenting on other arguments. However, whether an argument will sway public opinion is an empirical question, so this question should be addressed through experimental research. Because my dissertation does not utilize experimental methodology, I will not attempt to draw any final conclusions about rhetorical effectiveness. However, I will in some cases engage with philosophers' concerns about the clarity or accessibility of concepts, since clarifying concepts is within the purview of non-experimental philosophical work.

I will also clarify that I am only considering arguments for positions about our moral or prudential justifications for conserving biodiversity. This is distinct from positions about policy or about effective interventions. My conclusion that the Ecosystem Services Argument fails does not suggest that biodiversity conservation efforts should be abandoned. However, it does suggest that policymaking bodies and environmentalists should be careful about publicizing the Ecosystem Services Argument as the primary reason for biodiversity conservation efforts.

Before proceeding, I will discuss why services arguments have been popular. Arguments appealing to ecosystem services have been common in the past several decades in popular media, government documents, scientific and philosophical publications, and more. Services argument have some advantageous features compared to the other major argument forms presented above. They are particularly thought to be strong with respect to rhetorical effectiveness.

First, other types of biodiversity defense rely on more controversial philosophical concepts. For example, it has been a matter of dispute whether biodiversity (or any other ecological unit) has

intrinsic value; whether the aesthetic value of nature is objective enough to generate moral duties; and whether precautionary principles form a reasonable basis for decision-making. Most services arguments do not rely on any controversial principles or kinds of value, making them more widely acceptable. Second, everyone can recognize that various ecosystem services are very important for human wellbeing. This makes services arguments seem accessible and compelling to nonspecialists. In contrast, there have been concerns that arguments based on intrinsic value may not be credible or understandable by the public (Odenbaugh 2003).⁸⁹ To sum up, *prima facie*, services arguments seem to be among the best arguments for protecting biodiversity because they are widely acceptable and mostly do not rely on concepts that are controversial in philosophy.

We should acknowledge the environmentalists who believe that services arguments are inappropriate because they believe that instrumental arguments are too anthropocentric or do not fully capture the value of nature. Some find it morally objectionable to discuss the value of nature in terms of uses or economic quantities (McCauley 2006; Morelli and Møller 2015). I set aside these types of concern for the present discussion, since these authors already reject the Ecosystem Services Argument, though for different reasons from the ones I will give.

I do not think that instrumental argumentation should be rejected in general. I hold a pluralist view about values in environmental ethics, on the assumption that it is impossible to reduce all reasonable ways of valuing nature to a single ethical theory of type of value (Norton and Noonan 2007; Arias-Arévalo et al. 2018). Although debates in environmental ethics have often been framed as competitions between different kinds of value, the argument types described above

⁸⁹ Ironically, though, there have also been claims that the ecosystem services concept may be inaccessible to the public (Bekessy et al. 2018).

are not mutually exclusive—in particular, I do not think we need to choose between either instrumental or intrinsic arguments.

Before I move on, there are two more clarificatory comments. The first is about “biodiversity.” In this context, the goal is to generate arguments to support preserving the variety of species on earth. Thus, “biodiversity” means roughly “species richness,” i.e. the number of extant species, unless stated otherwise. Stipulating that I will use “biodiversity” this way here is a matter of convenience and follows the convention in the relevant literature, and I do not assume that biodiversity is best understood as species richness for other purposes.⁹⁰ I also am not assuming that saving species is the only or primary goal of conservation work, although I assume that most environmental thinkers agree that it is an important goal.

Second, the intent of the arguments under consideration is to support protecting species globally, or equivalently, to protect the number of species on earth. Throughout this chapter, I will utilize a distinction between “local” and “global” processes and properties. The meaning of these terms is often context-dependent. In the present context, I will stipulate that I consider properties and processes to be local when they are at the scale of a token ecosystem or part of an ecosystem, and to be global when they are at scales substantially larger than that of an ecosystem. The reason for drawing the distinction here has to do with the way services arguments are conventionally formulated; this will become clearer in the discussion below.

⁹⁰ There have been suggestions that *biodiversity* is a poor concept for conservation work because it is unclear, superfluous (i.e., can be replaced by more directly measurable quantities), or does not always align with what we value in nature (Santana 2014, 2017, 2018). I hope to circumvent these issues by stipulating that in this chapter, I’m only discussing the success of arguments for protecting global species richness.

On examination, the Ecosystem Services Argument does not provide support for protecting global biodiversity, so it fails with respect to its intended conclusion. The reason for this is a scale mismatch between cited ecosystem services and the global scale of biodiversity loss. As a result of this mismatch, the argument fails to show that global biodiversity has any instrumental value to humans. Some arguments of the other forms do not suffer from this scale mismatch and so they remain *prima facie* tenable, as I will show in the later sections of the chapter.

The conclusion about services arguments is unfortunate, since I support biodiversity conservation. However, I believe it is a matter of concern that both popular outlets and philosophers have frequently cited services arguments as major justifications for worrying about biodiversity losses. Many formulations of this argument are basically speculative and even conflict with reasonable interpretations of scientific evidence. I am also concerned by the lack of attention to how global biodiversity losses may interact with global-scale processes, including the climate and large-scale ecological processes. My goal is to encourage further investigation of the global roles of biodiversity and the ways in which global biodiversity is valuable to humans.

5.3 The Ecosystem Services Argument

I will next summarize the Ecosystem Services Argument and review some initial problems. Standardly, ecosystem services are defined as ecosystem processes, properties, or products which benefit humans. It is also possible to conceive of ecosystem services as benefits provided to any or all species. However, biocentric arguments based on the broader conception of services are

thought to suffer from serious flaws, so I will focus on the more conventional anthropocentric version.⁹¹

A basic version of the argument takes the following form.

1. There is a positive causal relationship between biodiversity and some ecosystem services. Biodiversity is needed to support these services; or these services would be lost or diminished if biodiversity were lost.⁹²
2. These services are very instrumentally valuable; we have good reason to want to retain the services due to the human wellbeing or financial costs of their loss.
3. So, we have good reason to protect biodiversity as a means to maintain ecosystem services.⁹³

Past discussion of this argument has focused on the empirical justification for the first premise. The main empirical evidence for this premise comes from the biodiversity-ecosystem functioning (BEF) research program, which includes long-term studies of the effects of different biodiversity levels on various indicators of ecosystem functioning. The initial BEF studies were manipulative field experiments in which ecologists altered the number of species in plots and then

⁹¹ One reason is that species have different and conflicting ecological needs, so a service for one species will be a disservice to another. This makes it challenging to specify what should count as a service, or to prioritize among species' needs. In addition, we risk approaching circularity if we argue that biodiversity should be saved to protect services, and services are defined as conditions that benefit all species.

⁹² These claims posit a causal relationship between variables. The causal claim can be helpfully read in terms of Woodward's interventionist theory (Woodward 2003), but other theories of causation may be equally apt as long as they allow for causal relations between variables.

⁹³ A similar argument is reviewed critically in *Defending Biodiversity* (Newman et al. 2017) on p. 47 and following.

tracked properties like the plot's biomass over time (e.g., Tilman and Downing 1994; Tilman 1996; Tilman et al. 1996; Tilman et al. 2001; Spehn et al. 2005; Tilman et al. 2006; Hector et al. 2007; Hector and Bagchi 2007; Hautier et al. 2015). The larger context for this research was the inherited diversity-stability debate about whether more diverse communities are more or less stable, which is one aspect of ecological functioning (McCann 2000; Jax 2010, chap. 3; deLaplante and Picasso 2011).

BEF studies have generally found a positive (sometimes asymptotic or saturating) correlation between diversity and measures of ecosystem functioning. However, there are questions about whether this provides evidence for a causal connection between diversity and the functioning of ecosystems, especially at larger scales. I will briefly review some main issues for the Ecosystem Services Argument, issues which are discussed at length by other sources (McCann 2000; Jax 2010; deLaplante and Picasso 2011; Newman, Varner, and Lindquist 2017). Although the problems sketched below represent important challenges to services arguments, I will argue that my scale objection is even more serious.

First, the causal interpretation of BEF findings remains challenging. One reason is that there are many distinct proposed mechanisms to explain diversity-functioning relationships, and only some of the mechanisms are causal, while others explain correlations as (merely) probabilistic outcomes or even as artifacts of study design (Doak et al. 1998; Giller et al. 2004; Loreau and Mazancourt 2013; Downing et al. 2014). Initial BEF studies were not designed in a way to distinguish among possible mechanisms. More recent BEF research has moved towards designing studies to distinguish among possible mechanisms for biodiversity-functioning relationships (Hector et al. 2010; Hallett et al. 2014). However, due to the ongoing nature of this research, there seems to be no consensus about the relative importance of the proposed mechanisms.

It is of course likely that different mechanisms are at work in different communities. The initial BEF research focused on a small number of ecosystem types which are amenable to this kind of study. Most of the studies were done on grasslands (often measuring productivity as the main indicator of functioning), with additional research on decomposition rates or community stability in aquatic systems (e.g., Little and Altermatt 2018). Although ongoing research is investigating a greater variety of ecosystem types, we still have uneven knowledge about the important mechanisms in different ecosystems.

A second problem, as discussed previously in Chapter 3, is with the external validity of classic BEF studies. To review, many of these studies utilize artificially simplified, small-scale experimental plots, often including only a single trophic level. In addition, many of the studies have focused on a limited number of ecosystem types, especially temperate grasslands, yet we expect that different ecosystems could respond in qualitatively different ways to diversity changes.

If the results of these studies are to inform the management of ecosystems, we need reason to believe that the same causal principle(s) describing the behavior of the experimental plots also describe the behavior of real ecosystems (Cartwright and Hardie 2012, chap. I.B; Marcellesi 2015). There are reasons to think this is not the case, since BEF experiments omit many factors which are known to be causally important in ecosystems (e.g., higher trophic levels; realistic community assembly processes; disturbance regimes). It is a well-known feature of ecology that patterns and processes at one scale often do not scale up nicely. A pattern observed in a single trophic level at a small spatial scale might be qualitatively different from the pattern observed in a larger system. So, results from classic BEF studies do not specify how realistic ecosystems will respond to diversity changes (Hooper et al. 2005; Brose and Hillebrand 2016). This has been widely acknowledged among BEF researchers but sometimes has been overlooked in popular reporting.

As one might expect, the BEF research program has moved in the direction of studies that focus on more realistic or multi-trophic systems (Little and Altermatt 2018; Moi et al. 2021). This includes observational in addition to manipulative studies (Hector et al. 2007) and studies of BEF relationships at larger spatial scales (e.g., Patrick et al. 2021). Very large-scale studies are increasingly possible due to remote sensing technology. Although these studies face their own statistical and interpretive challenges, we can expect that they will continue to improve our understanding of real ecosystem behavior.

A few observations about the Ecosystem Services Argument are in order. First, we might not want our obligation to conserve biodiversity to hang on currently open-ended empirical research. Second, it remains the case the BEF research focuses on simple ecosystem functions—such as productivity—which are amenable to various forms of estimation. Ecosystem functions, however, are conceptually distinct from services (*pace* Newman et al. 2017, 51). Functions are the processes that occur in ecosystems as studied by core ecological science. Services are goods and opportunities provided to humans by ecosystems, such as water purification, food, and recreational opportunity. Ecosystem functions are not necessarily services; some, including periodic flooding or fires, can be disservices. This adds complexity to extrapolating from ecological studies of ecosystem functions to conclusions about how to promote ecosystem services.

It should also be noted that managing an ecosystem for services can undermine native diversity levels. Monocultural agriculture is the most obvious example of this. Reyers et al. (2012) consider some examples of “win-lose” cases where managing for either services or biodiversity results in negative effects on the other item—for instance, damming rivers may be good for services but harmful to native diversity. This observation casts further doubt on the usefulness of arguments which connect local diversity levels to ecosystem services. Ecosystem services

arguments presuppose that managing ecosystems for services is beneficial, yet managing for services in some cases is detrimental to biodiversity. This means we may sometimes need to choose between prioritizing services and prioritizing biodiversity. For this purpose, we need independent reasons for thinking promoting biodiversity is a good goal.

A final relevant issue has to do with the interpretation of “biodiversity” in the hypothesis that higher biodiversity contributes to better functioning or stability.⁹⁴ The initial BEF studies measured species richness, assuming that taxonomic diversity is a good predictor of community-level functions. However, there is evidence that functional trait diversity better explains community-level outcomes than does the diversity of species (Mori et al. 2013; Gagic et al. 2015; Sakschewski et al. 2016; Mason et al. 2005 on functional diversity indices). This makes intuitive sense given that it is species’ role functions rather than their evolutionary history which proximately explain their contributions to ecosystem behavior (Chapter 2; Odenbaugh 2010). Therefore, to the extent that diversity is causally relevant to ecosystem functioning, it is likely the diversity of role functions rather than the number of species present that is most explanatory.

It is also known that some individual species or functional types of species are disproportionately important to ecosystem functions (on functional types, Díaz and Cabido 2001). In some cases there is evidence that one or a few dominant species, rather than diversity of any kind, drives community stability (Sasaki and Lauenroth 2011). So, setting aside for now potential external validity problems, BEF findings provide reason to protect functionally important species

⁹⁴ The meaning and utility of the biodiversity concept has been controversial (Maclaurin and Sterelny 2008; Burch-Brown and Archer 2017; Santana 2018). If a given biodiversity index is the best predictor of functioning, this does not entail that “biodiversity” ought to be interpreted in the same way for all purposes.

but do not univocally support the assumption that modest species richness losses will cause consequential breakdowns in ecosystem functioning.

The hedge term “modest” is important here, since it is clear that a major loss of species will cause ecosystem function breakdowns. For management purposes, we should focus on a degree of species loss that can be reasonably expected in the near term. In this respect, again, classic BEF experiments may not be representative of likely scenarios, since they often use many plots with just a few species. Realistic projections about biodiversity changes will help to inform our expectations and should be considered when designing BEF research.

I have challenged the assumption that modest species richness declines will be a proximate cause of lost services. However, there is also a problem in the opposite direction. Population declines among species which directly support services (e.g., pollinators and fished ocean species) or which are functionally important (e.g. keystone species, top trophic levels) can and do cause disruptions to ecosystem service provision. Population declines fall short of extinction and therefore do not affect species richness. This further shows that, if we are primarily concerned with services, species richness is an inadequate variable to focus on.

As serious as they might be, the above problems for the Ecosystem Services Argument are not the primary reason the argument fails. BEF research is ongoing and we can expect future research to provide further evidence about how biodiversity declines will affect ecosystem processes. Even as future research brings more clarity, the standard Ecosystem Services Argument will not be successful because of the scale problem, which I will present next.

5.4 The Scale Problem

The scale problem shows that any standard-form services argument cannot be successful, regardless of the details discussed above. If standard services arguments are sound, they give us reason to protect the species diversity of individual (token) ecosystems in order to protect the services provided to humans by those ecosystems. But biodiversity loss is a global phenomenon. The argument under consideration is meant to support preventing global species losses. This goal is only justified if we add a further premise to the argument, stating that global biodiversity losses contribute to biodiversity losses within individual ecosystems.

Here is a reformulated Ecosystem Services Argument, which makes the implicit premise and the intended conclusion explicit (additions are italicized).

1. There is a positive causal relationship between *local* biodiversity and some ecosystem services. *Local* biodiversity is needed to support these services; or these services would be lost or diminished if *local* biodiversity were lost.
2. These services are very instrumentally valuable; we have good reason to want to retain the services due to the human wellbeing or financial costs of their loss.
3. *Global biodiversity losses will cause local biodiversity losses.*
4. So, we have good reason to protect *global* biodiversity as a means to maintain ecosystem services.

For the purpose of this discussion, ‘local’ scale entities, processes, and properties occur at the scale of an individual ecosystem. Throughout, I have opted to use phrases like ‘local biodiversity’ to avoid equivalent but cumbersome phrases like ‘the species richness within individual ecosystems.’ Local-scale phenomena here contrast with much larger-scale phenomena, i.e. those characteristic of landscapes, continents, or the whole globe. Although I recognize that precisely individuating

ecosystems is a problem, I assume that we can distinguish between ecosystem-scale and global-scale phenomena clearly enough to assess the argument at hand.

First, I will state why the above argument is in tension with some relevant empirical findings. Second, I will explain on a conceptual level why we should not expect the above argument to be sound.

The link between global and local species losses is a matter of empirical investigation. Surprisingly, some large meta-analyses have concluded that although global diversity is declining, local species richness has not on average been declining (Vellend et al. 2013; Dornelas et al. 2014). Intuitively, this is because the ranges of some species are expanding while the ranges of others are shrinking, such that in some areas the total diversity is increasing or not changing.

It is certainly the case that local biodiversity is declining in some places as ecosystems collapse or are removed for development. Various urban locations, monocultural agriculture sites, or sites of recent deforestation will be characterized by low biodiversity levels. So, the conclusion cannot be that there are few sites which have experienced biodiversity loss. Instead, the claim is that the increasing rate of global extinction has not so far contributed in a consistent manner to local biodiversity losses within some studied ecosystems.⁹⁵ Critical responses to the cited meta-analyses point to the prevalence of human-influenced sites and question whether the meta-analyses sample a representative set of studies, among other methodological features (Gonzalez et al. 2016; Cardinale et al. 2018).

The previously cited meta-analyses report changes in biodiversity composition even in areas where net loss was not detected. As previously discussed, there are reasons to believe that

⁹⁵ See the exchange between Cardinale (2014) and authors of the Dornelas et al. study.

species richness is less explanatorily relevant than more sophisticated measures of biodiversity change (Hillebrand et al. 2018). So, the moral is not that there is nothing to be concerned about. Rather, arguments which focus on species richness may not capture more salient features of biodiversity change.

It is an empirical question to what extent local systems are experiencing species richness declines—and, apparently, a question which is not fully settled to the mutual agreement of ecologists. But even if species richness were declining at a majority of sites, it is a further question what caused the declines. Human land use, climate change, and sometimes biological invasion seem to be the main causes of major local biodiversity reductions. Unsurprisingly, studies which track land use detect a much greater degree of local species loss (Newbold et al. 2015). This further supports my claim that the global extinction rate is not what is primarily driving local biodiversity changes.

I have outlined some empirical questions surrounding the extent to which local species richness is declining. In addition, I suggested that the global extinction rate is not the main driver of local changes, even where species richness is declining. The reason this is possible may be further clarified by applying the distinction among within-system, between-system, and total diversity, termed α -, β - and γ -diversity respectively. Think of a landscape consisting of a patchwork of different communities. In this example, α -diversity refers to the diversity of species within a local community, while β -diversity quantifies the spatial heterogeneity—the differences in composition between communities. And γ -diversity refers to the total number of species within the landscape. BEF research has mostly investigated α -diversity, and standard services arguments appear to have α -diversity in mind, since they claim that services will be lost as species are lost from ecosystems.

The conceptual problem is that constant or increasing levels of α -diversity are consistent with decreasing levels of β - and γ -diversity. This can occur as a result of declines in between-system heterogeneity—i.e., when the composition of different local communities becomes more similar. If all that matters causally is α -diversity, then there is no reason not to maintain local diversity via homogenization. This is at odds with the conservation goals of preserving distinctive local communities and maintaining the total number of species at larger scales (Lean 2021). Therefore, again, there is a scale mismatch between the services argument, which cites a putative causal role of α -diversity, and the intended conclusion which is to protect global diversity.

It is also possible for α -diversity to decline much more quickly than γ -diversity, which is what happens when an area is deforested. Extreme local declines in α -diversity are a concern from the perspective of service provision. But again the proximate cause is human activities, not the global extinction rate.

Recently there have been calls for more study of how β -diversity affects BEF relationships. There is evidence that β -diversity declines can affect ecosystem functioning, so this is an important area for future research (van der Plas et al. 2016; Patrick et al. 2021). However, most existing discussions have been from a theoretical or modeling perspective rather than experimental (Burley et al. 2016; Wang and Loreau 2016; Mori et al. 2018; Gonzalez et al. 2020). Existing predictions should be considered preliminary, and we cannot form any very precise conclusions about how changing β -diversity will affect ecosystem services.⁹⁶

⁹⁶ I'm assuming that predictions of ecological models without experimental or observational confirmation need to be taken with a grain of salt. There's a history of early ecological models making predictions which are later disconfirmed or substantially reinterpreted when more empirical data come in. The history of BEF research provides one example of this; see Justus and Wakil (2021) for another.

An additional consideration is that at very large scales, ecosystem characteristics are driven primarily by climate and by *neutral* processes like speciation/extinction and colonization rates. β -diversity is likely to be most causally important at intermediate scales. Those effects of β -diversity on local ecosystem functioning will be mediated by the degree of synchrony or connectivity among communities. In short, the effects of β -diversity on ecosystem functions are likely to be complex and scale-dependent, much like the effects of α -diversity. However, if an argument can be made that declines in β -diversity, of a magnitude that is likely to occur over human-relevant timescales, will significantly compromise ecosystem resilience or service-supporting ecosystem processes, then this would go a long way towards justifying some intuitions about conservation. So, although I think arguments about β -diversity are currently far from conclusive, this is an area that warrants further study.

At best, the standard version of the Ecosystem Services Argument, which focuses on α -diversity, would show that the number of species at a given location ought to be maintained, regardless of the identity of those species. Yet environmentalists agree that rare or unique species should be given special priority as a part of biodiversity conservation (Deliège and Neuteleers 2015; Lean 2021). We also usually think that introduced invasive species do not count towards biodiversity conservation goals.⁹⁷ Thus, we do seem to value compositional heterogeneity (β -diversity), not just the total count of species at each location. As I have now argued, the standard services argument cannot justify this preference.

Complaints that instrumental arguments cannot distinguish between native and nonnative species have been made previously. For example, concerns have been raised that services

⁹⁷ Lean (2021) discusses a similar set of issues as applied to invasive species.

arguments would justify introducing species or even replacing ecosystems with artificial technology if this results in better performance of functions or services (Deliège and Neuteleers 2015; Newman et al. 2017, sec. 2.7 ff; Desjardins et al. 2019). Concerns have also been raised that if biodiversity or species richness is the main locus of value, then we might have a counterintuitive obligation to increase the number of species at any location rather than to protect the native species (Santana 2017).

My presentation of the scale problem explains why these are persistent concerns. Some authors have implied that these problems affect instrumental arguments as such. However, on my assessment the root of the issue is that the Ecosystem Services Argument operates at the wrong scale. Protecting native diversity and rare species is warranted assuming that our goal is to protect compositional heterogeneity at larger scales and to protect the total number of species globally. But because services arguments have focused on processes and properties at the scale of individual ecosystems, they have no bearing on these potential larger-scale goals. The unintuitive consequences of the argument arise because of this misalignment.

5.5 Alternative Services Arguments

The argument I have discussed so far is what I characterized as the standard version of the Ecosystem Services Argument, which connects α -diversity to ecosystem services or functions. However, there are other forms of service argument and you might wonder if they avoid the preceding objections. In this section I present some alternative versions and argue that they do not fare much better than the original argument.

5.5.1 Conjunctive Argument

The first reformulation is a conjunctive argument which appeals to the roles of individual species in conjunction rather than to the roles of diversity per se. This argument claims that since many individual species support important services, considering all of these services collectively gives us reason to preserve biodiversity in general. This results in an argument with fewer empirical limitations, since the services provided by individual species are often easier to detect than the services supported by biodiversity as such.

An argument of this form has recently been endorsed by Jay Odenbaugh (2020), though what I provide here is my own summary.

1. Many individual species support important services such as pollination, food provision, drug development, attracting tourism, etc.
2. These services are very instrumentally valuable and will be lost as the relevant species are lost.
3. The collective instrumental values of individual species give us good reason to protect global biodiversity.
4. So, we have good reason to protect global biodiversity.

This kind of argument is subject to the ‘patchiness’ objection (Newman 2020). Briefly, according to this criticism, although many species have important uses, there are also very many species which have no apparent uses to humans. The existence of useful species does not provide an instrumental reason to protect non-useful species. If we already know which species support services, then from the standpoint of protecting those services, we should focus on trying to protect just those species with known use value.

A common reaction is that instrumentally important species may rely on mutualist or food web relationships with other species. This gives us reason to protect those other species. But keep in mind that the total number of species on earth swamps the number of species which directly support services. Moreover, some service-providing species are generalists, so they do not rely on the presence of other particular species. So it remains a problem to account for the instrumental value of the vast majority of species.

Although the argument is not sound as stated, there is a straightforward way to improve it. The strategy is to add a precautionary principle which can support protecting species without currently known uses.

5.5.2 Precautionary Conjunctive Argument

A precautionary version of the above argument is as follows.

1. Many individual species support important services such as pollination, food provision, drug development, attracting tourism, etc.
2. These services are very instrumentally valuable and will be lost as the relevant species are lost.
3. We are not in a position to know which species may turn out to provide important services.
4. So, by a precautionary principle, we have good reason to protect all species.

There have been many different formulations of precautionary principles.⁹⁸ The principle relevant here is one which states something like, “we should take steps to prevent serious negative consequences, even when the likelihood and exact nature of those consequences are unknown.” Thus, the idea behind this argument is that we should try to stop species losses in order to prevent possible future harms, even ones we cannot now foresee.

There are several reasons this precautionary conjunctive argument may not be compelling. First, one might think that we *are* in a position to know which organisms are likely to be useful; that we have good reason to believe the majority of species are not useful; or that we have reason to believe that many of the most useful organisms, e.g., ones with agricultural applications, are in the least danger of extinction (Maier 2012, chap. 6). For example, we have good knowledge about which plant species are edible by humans. Where we lack such knowledge, it is likely because the species are uncatalogued, rare, remote from humans, or difficult to cultivate, making them unlikely candidates to become important food sources. So, you might think this argument relies on a false generalization about our state of knowledge.

Precautionary arguments in general have been subject to criticism. Some philosophers argue that precautionary principles are misguided because they seem to advise a strategy of complete risk avoidance, a strategy which is impossible to implement and contrary to ordinary reasoning (e.g., Newman et al. 2017, chap. 3). Ordinarily, people are willing to accept a small chance of major harm in exchange for a likely moderate gain, suggesting that most people do not accept a principle that states we should take action to avoid all potential future harms (Sober 1986).

⁹⁸ There is a large literature on the precautionary principle. Some starting points are Manson (2002) for an overview; Steele (2006) for a favorable discussion; Newman et al. (2017, chap. 3) for a discussion critical of the principle.

If you share these kinds of reservation about precautionary principles, then you will not find the above argument effective.

This argument avoids the scale problem since it concludes directly that we should try to preserve all species in order to avoid risks of future harm. However, if we have evidence that some species do not have instrumental value to humans, the argument seems to allow that we no longer have reason to protect them. Problematically, the argument only justifies protecting all species to the extent that we have incomplete knowledge about the consequences of their loss. We would expect that our impetus to protect all species will decrease as we gain better scientific knowledge. So, this argument does not capture the intuition that global biodiversity decline as such is undesirable, irrespective of our state of knowledge about the uses of particular species.

Further, in practice, conservation work faces a large number of competing goals. If we want to focus on instrumental justifications, the precautionary argument for protecting all species will be pitted against more pressing instrumental justifications for protecting currently very valuable species such as pollinators and crop species. Appeals to unknown, unspecified future risks cannot reasonably outweigh appeals to currently known use and economic values. Therefore, in practice, it is hard to see how the precautionary conjunctive argument can give priority to efforts to protect all biodiversity.

5.5.3 Option Value Argument

Finally, consider an argument which appeals to the option value of species. Erik Persson defines option value as “[t]he value something has because it provides an alternative way of

promoting something else that has instrumental and/or end value” (Persson 2016, 167).⁹⁹ When multiple species can potentially sustain the same service, those species have option value. We value the fact that these species are intersubstitutable because they provide some insurance against losing services in the event that one species is lost. In general, we value the existence of alternative options because they leave room for more choices or new preferences in the future. Here is my summary of Persson’s option value argument for protecting biodiversity:

1. Many individual species support important services such as pollination, food provision, drug development, attracting tourism, etc.
2. These services are very instrumentally valuable and will be lost as the relevant species are lost.
3. In many cases, multiple species support the same service.
4. When multiple species support the same service, we have good reason to protect all of them because of their option value.
5. So, we have good reason to protect all species which support services.
6. So, we have good reason to protect species in general.

Option value considerations can provide good reason to protect a diversity of species that provide specific important services which would be difficult to replicate without those species. For example, I agree with Persson that option value gives us very good reason to protect a diversity of pollinating insects, given that pollination is a vital service which would be difficult to replicate without the pollinators. However, option value is only important in cases when a limited number of species relatively directly support an important ecosystem service or function. Where there is a

⁹⁹ This is distinct from the technical sense of “option value” in economics.

large number of species that support the same function (e.g., the number of primary producers in many systems), these species have minimal option value, because there is no reasonable chance of losing the function if some species are lost.¹⁰⁰ You might think that species are *less* instrumentally valuable to the extent that they are extremely functionally redundant. Thus, (5) follows only in cases where the functional role of the species is relatively uncommon. In addition, (6) does not follow from (5), given that many species do not support irreplaceable services. So, service-based option value arguments will be subject to the same patchiness objection discussed previously. They can justify protecting species that support important and relatively irreplaceable services, but they cannot justify protecting arbitrary species (such as, presumably, many non-pollinating invertebrates, wild herbaceous plants, and small mammals, reptiles and amphibians).

You might think that there are various species which do not directly support a service, but which help to sustain services indirectly by virtue of their role in a community. The problem with option value and these “indirect” roles is that many species are highly intersubstitutable, especially from the perspective of service provision. For example, having wetlands and forests is good for water quality, but there are huge numbers of species that could make up a wetland or forest. Thus, relevant species will have low option value with respect to water quality. So, the option value services argument does not give us a reason to save most species.

I do think option value might justify protecting a diversity of ecosystems, and this could go some way to supporting biodiversity protection at larger scales. The argument would state, briefly, that we should protect a variety of ecosystems which provide vital services as insurance

¹⁰⁰ Donald Maier makes this point more colorfully: “there seems to be a gross superfluidity of species—a true *embarras de richesses*—associated with the majority of ecosystem services” (Maier 2012, 178).

against the loss of those services. However, there will still be problems surrounding, e.g., low abundance species which do not contribute appreciably to ecosystem functions. Put differently, even if we have good reason to protect ecosystems, this will only provide us with justification for protecting the functionally important species within those ecosystems. The reason is that ecosystems can maintain their identity and functions over time even as their composition changes to some extent.

5.6 What is Global Biodiversity Good For?

In the preceding sections I showed that the major conventional services arguments and precautionary-services arguments for protecting biodiversity fail to justify their intended conclusion. Some versions suffer from the scale problem, such that they fail to demonstrate any instrumental value of global biodiversity. Other versions can be formulated to avoid this problem in principle, but they do so at the cost of having comparatively weak conclusions which rely on people agreeing to risk-averse approaches to management. These also result in making decisions based on our lack of knowledge rather than based on what evidence we do have. These arguments, which appeal to unknown harms, cannot reasonably be taken to give strong priority to protecting global biodiversity given the existence of many other ongoing environmental problems with known, immediate costs and harms to human life.

In this section, I consider possible arguments for the value of global biodiversity as such, which I take to consist of the global richness of species, with a preference for maintaining species in roughly their native ranges and habitats. My major aim here is to show that these arguments do not suffer from the scale problem, although I will show that the first argument I consider is

inconclusive for other reasons. I will not attempt here to fully justify any of the remaining arguments; my goal instead is to show which arguments are *prima facie* viable because they operate at the correct scale. I recommend that philosophers with environmental concerns continue to advance work on these arguments, even though they have all been controversial.

5.6.1 The Tipping Point Argument

This argument is a precautionary, global-scale version of the ecosystem services argument. Of the arguments I consider in this section, it is closest in form to the “standard” instrumental arguments.

The services arguments considered above have a common feature: they try to link the α -diversity within ecosystems to services provided by those ecosystems. So, the services under consideration are those generated by local processes within ecosystems (although the services might be consumed by humans either locally or globally). However, it is also possible to think of services produced at the global scale by global-scale processes. Climate stability is a plausible example of a global service—it is a service that cannot be explained at the scale of individual ecosystems.

If global biodiversity can be causally linked to some global services, the resulting argument will avoid the scale problem. Yet we currently seem to have poor knowledge about (a) what are the candidates for global services and (b) how biodiversity changes at the global scale will affect other global-scale processes. Recall, for example, that the majority of BEF research pertains to (very-)sub-ecosystem-scale processes. A major takeaway is that more attention should be devoted to these global-scale questions.

Despite the current lack of empirical information, one may construct a global-scale services argument utilizing a precautionary principle, as in the following.

1. If global biodiversity losses become great, we may reach a tipping point beyond which there will be a widespread collapse of major ecosystems, breakdowns in biogeochemical processes, runaway climate feedbacks, or accelerating further biodiversity losses.
2. We are not in a position to know when such a tipping point would be reached.
3. So, by a precautionary principle, we have good reason to protect global biodiversity.

This argument is my updated version of the famous airplane rivet analogy (Ehrlich and Ehrlich 1981; Ehrlich and Walker 1998). Anne and Paul Ehrlich and their coauthors compared species to the rivets in an airplane. While an airplane may continue flying as rivets are removed (due to redundancies), at some point, a threshold will be crossed and the airplane will experience a catastrophic breakdown. Similarly, the Ehrlichs suggest that the removal of species would eventually result in catastrophic breakdowns of ecosystem functioning. They think that preventing species losses is urgent because of our ignorance about when the breakdown will occur. As summarized, their argument is subject to the scale problem. My version is formulated to operate at the correct scale by pointing to potential global consequences of global species losses. Although this argument avoids the scale problem and deserves further consideration, I will argue that it is currently inconclusive.

My version of this argument is inspired by the recent debates about global-scale tipping points in climate science. There have been debates about whether the drivers of climate change will cause the climate system to pass a tipping point, after which accelerating positive feedbacks will result in rapid species losses and global functional changes (Brook et al. 2013; Lenton and

Williams 2013). It should be noted that relevant scholars (including the cited authors) disagree about the nature of tipping points and whether global tipping points are likely to occur.

An important difference between aircraft and ecosystems is that ecosystems can undergo major reorganizations without losing all functionality.¹⁰¹ The notion of “ecosystem collapse” at work in this argument should be defined carefully (see Keith et al. 2013). “Collapse” means something like a major, undesirable, and relatively irreversible change to the state of an ecosystem (or the biosphere), particularly one which involves major native biodiversity loss or a major decrease in populations of functionally characteristic species. Climate change and biodiversity changes are resulting in ecosystems undergoing regime shifts, but this does not automatically mean that local or global-scale functions will be lost. Sometimes, new regimes are considered less healthy or desirable (Chapter 4). However, even a less healthy ecosystem regime can support a diverse biota and characteristic ecosystem functions. Thus, a collapse is an extreme kind of regime change which involves major loss of biota. Most standard ecological regime shifts should not be considered a collapse of the ecosystem.

My characterization of ecosystem collapse raises an issue with the causal interpretation of the first premise of the argument. Because biodiversity losses are partly constitutive of ecosystem collapse, biodiversity losses should not be treated as separate cause of this phenomenon as perhaps suggested by premise (1). Biodiversity losses are more naturally considered components, indicators, or effects of ecosystem collapse.

¹⁰¹ For convenience, I will continue talking about ecosystems, but the global-scale argument predicts that there will be widespread ecosystem collapses which will affect global processes.

So, it is more natural to run arguments like this in the opposite direction. The fact that climate change and human activities are in the process of wiping out whole ecosystem types and threatening others with irreversible regime shifts is a matter of serious immediate concern. Ecosystem changes are known to contribute to species extinctions. Thus, we should try to prevent ongoing ecosystem collapses (among other reasons) as a means to prevent accelerating biodiversity losses. However, this argument presupposes that we already have reasons to think we should try to prevent biodiversity losses.

Similarly, we have knowledge about how climate change contributes to species extinctions, and we are working on understanding how climate change and ecosystem changes might influence each other, but it is not well understood how changes in biological diversity might contribute to climate change. Therefore, we have good reason to try to mitigate climate change as a means to protect species, but no clear reason to try to protect biodiversity as a means to mitigate climate change. (The reader is reminded that even though some individual species contribute to ecosystem functions which contribute to climate change mitigation, this does not justify protecting global species richness, for the reasons given in sections 5.3 and 5.4.)

Although I am sympathetic to precautionary arguments that we should try to prevent runaway global-scale processes like the ongoing climate crisis,¹⁰² I do not think these types of argument can function primarily as justifications for protecting arbitrary species. The current climate situation is known to be a result of human atmospheric pollution, and although biodiversity

¹⁰² Precautionary arguments related to climate change are more defensible than those previously considered, since climate change carries known serious risks to human life. There are still interpretive problems about the formulation and implications of the precautionary principle, but I think precautionary arguments which point to known risks must be taken more seriously than ones which point to the mere possibility of harms of unknown character.

may contribute to climate dynamics in ways that are not currently understood, the primary causal direction at the moment seems to be from climate change to biodiversity loss. So, we still need an independent reason for believing that biodiversity loss is undesirable and that we should make a special effort to prevent biodiversity losses in addition to climate mitigation and ecosystem management efforts.

This concludes my discussion of potential instrumental arguments for conserving biodiversity. I will next consider intrinsic value, followed by some discussion of values which sit somewhere between the standard use values and intrinsic value.

5.6.2 Future Options

Above I presented an option value argument that appealed to species' actual or potential contributions to ecosystem services. This is not the only possible formulation of an argument from option value. James Maclaurin and Kim Sterelny (2008, chap. 8) favorably discuss option value arguments for the protection of species richness or taxonomic diversity which appeal to (a) future human preferences and (b) evolutionary potential. First, they argue that human attitudes towards nature have changed over time, so a major reason to conserve biodiversity is for “insuring against changes in what we ourselves want and value” (156).

This is not a convincing argument because the possibility that we could change our preferences applies to too many things. At worst, the idea that we should save everything that we might want in the future advocates a hoarding mentality. In addition, this “possible future preference value” leaves us with no way to prioritize among the very many things humans might conceivably value in the future. Finally, it is overly burdensome to expect all the potential values of future humans to guide our current actions. This is not to say we should discount the wellbeing

of future humans, but that we cannot possibly anticipate all possible desires future humans might have.

Charitably, the idea behind this argument is that, because we cannot anticipate future values, we should try to protect a variety of potentially valuable entities. While this seems like an innocuous idea, the problem is that we have limited resources to protect the huge variety of potentially valuable entities on earth. It is inevitable that some species will go extinct in the near future. In practice, we must make decisions about which entities to prioritize. These decisions are better guided by our current values—or by values which seem to be justified by existing arguments or preferences—or by the values we can reasonably anticipate future humans having. We have a better chance of specifying these values clearly enough to aid prioritization decisions, for instance.

Consider now a somewhat different future option argument, which states that species should be protected because they support future evolutionary options. Maclaurin and Sterelny write, “We should deplore the extinction of any species because every species represents a new and potentially important trajectory in a space of evolutionary possibility” (157). While this represents a compelling way to think about extinction, it does not explain why species diversity should be valued in the first place. There are a lot of species of beetles, and presumably many of these species have the potential to give rise to new species in the future. But the only reason to think these new species would be valuable is if we already think it is good for there to be a lot of beetles. The argument that species should be protected because they can give rise to more species threatens to beg the question.¹⁰³

¹⁰³ Carlos Santana (2014, 775) provides some further objections to arguments from evolutionary potential. One of his points is that speciation has often followed extinction events. If it’s correct that mass extinctions promote speciation,

5.6.3 Intrinsic Value

Environmentalists have been divided over intrinsic value (Appendix A), which is roughly whatever value something has independently of its uses or its contribution to other valuable ends. It is worth considering further because intrinsic value defenses may not suffer from the scale problem. If biodiversity as such has intrinsic value, then that value belongs to it “directly” and provides direct reason to protect it.

Various authors have disputed the claim that biodiversity (Oksanen 1997) or other ecological wholes (Justus et al. 2009; Newman et al. 2017) are the sorts of things that can have intrinsic value. To make sense of the debate, I propose a distinction between two conceptions of intrinsic value held by various environmentalists. “Intrinsic value” has been used in a lot of different ways by philosophers, and I will not attempt a full survey here (see McShane 2017). Instead, I focus on one distinction that is relevant to this context.

Some philosophers treat intrinsic value as coextensive with direct moral considerability (e.g., Vucetich et al. 2015). On this view, something has intrinsic value just when it has interests (in a suitably robust sense) and when those interests sometimes create direct obligations on the part of agents. I will refer to this conception as *strong* intrinsic value. This is the conception of intrinsic value considered by Oksanen (1997) and Newman et al. (2017, chap. 9) who both argue that intrinsic value fails to apply to ecological wholes.

then it’s a bit strained to argue that extinction should be prevented in order to promote speciation. In addition, when evaluating these arguments further, attention should be paid to the timescale over which events are supposed to occur.

To see whether biodiversity might have strong intrinsic value, consider an influential ecologist position—the position that ecological wholes are directly morally considerable. A recent argument for considering ecosystems directly morally considerable is provided by Roberta Millstein in her reconstruction of Aldo Leopold’s land ethic (Millstein 2020a). Briefly, she argues that if we intrinsically value human communities, then consistency demands that we also intrinsically value land communities, since (according to Leopold’s views) we share similar relationships of interdependency with members of our land community¹⁰⁴ that we do with members of a human community. Millstein cites the fact that soldiers are willing to sacrifice their lives for their community as evidence that some people do intrinsically value human communities (i.e., they view the community itself as generating duties on the part of human agents).

Suppose for the sake of discussion that we are willing to grant this position, and consider whether Leopoldian ecophilism could directly support biodiversity conservation. A problem arises here due to the ontological difference between communities and biodiversity. Communities are collections of concrete objects—you can often draw approximate boundaries around their parts in the world. In contrast, “biodiversity” refers to the variety of life in an area, so it is better understood as a property (or multiple related properties). As a result, it is not the right sort of entity to generate moral duties. Asserting that we have moral duties *to* diversity or another property looks like a plausible example of a category error.¹⁰⁵

¹⁰⁴ “Land community” is Leopold’s term for something in between the contemporary concept of a community and a token ecosystem. See Millstein (2018).

¹⁰⁵ My argument has nothing to do with claims about the reality of collectives or biodiversity. There is a background debate about whether biodiversity is a real property or kind (e.g., Burch-Brown and Archer 2017), in addition to a

If you are a Leopoldian ecologist who thinks ecosystems are directly morally considerable, then you more likely think that we have a defeasible duty to protect the diversity within the ecosystem for the sake of the ecosystem. That is, a reason to protect biodiversity comes from our duties towards ecosystems. However, this position is subject to the scale problem. It gives us a reason to protect the diversity within individual ecosystems, but no reason to protect global biodiversity.

In order to get a duty to protect global biodiversity, you would have to adopt the view that the earth as a whole is directly morally considerable, and that protecting its biological diversity is among our duties towards the whole earth system. I mention this as a possible viewpoint, but it has been extremely unpopular among both philosophers and scientists.¹⁰⁶

Since I have brought up ecologist positions, a further issue is whether these views get off the ground at all. If you share the common intuition that entities must usually be sentient, self-aware or rational (or capable of becoming one of these) in order to be directly morally considerable, then you will agree that ecosystems are not directly morally considerable and thus not intrinsically valuable in the strong sense. For these reasons, ecologism does not provide a clear avenue to defending global biodiversity.

debate about whether ecosystems are real entities (Chapter 4). I think these debates about realism are irrelevant to our moral obligations for reasons discussed in Chapter 4 (and to be developed in future work).

¹⁰⁶ Probably because the idea is associated with mystical views. I see no immediate reason why “geologism” is less credible as a possible position than ecologism. If it is worth taking seriously that we might have duties to ecosystems, it is also worth taking seriously that we might have duties to the whole earth system. I think the latter view need not involve any controversial scientific, metaphysical, or spiritual views about the nature of the planet. However, I personally would not endorse either ecologism or geologism for the reason stated in the following paragraph.

Some environmental authors end the discussion at this point and conclude that intrinsic value defenses are no good, but this overlooks a distinct interpretation of intrinsic value. Some philosophers have treated intrinsic value in a weaker sense as the value something has when it is valuable or valued for its own sake (McShane 2007; Callicott 2017). On this *weak* conception of intrinsic value, it makes sense to attribute intrinsic value not just to sentient individuals but also to certain non-sentient individuals, works of art, ecological wholes, and even properties such as diversity and beauty. Sometimes intrinsic value of this sort is referred to as “end value,” which emphasizes the contrast with use or instrumental value (Persson 2016; McShane 2017). I believe part of the reason for the persistent disagreements about intrinsic value is a failure to recognize the difference between stronger and weaker conceptions of intrinsic value. Objections to intrinsic value often target the stronger sense, while some of those who favor intrinsic value only have the weaker sense in mind.

For clarity, this distinction between strong and weak intrinsic value is orthogonal to the distinction between objective and subjective value. Objective value is value an entity has “mind-independently” or in virtue of its own properties, while subjective value is had in virtue of being valued by someone. The stronger and weaker senses of intrinsic value I have distinguished here are both in principle consistent with any of these views about the nature of value. However, it seems more common for weak intrinsic value to be associated with a subjective conception of value and for strong intrinsic value to be treated as objective.¹⁰⁷ (The reader is referred to Appendix A for more detail on my understanding of intrinsic value.)

¹⁰⁷ If it’s not clear how strong intrinsic value could be subjective, one such idea would be that entities are directly morally considerable only if there are agents which can recognize duties towards those entities. Conversely, weak

Some environmental literature has also talked about the related concept of existence value, which is the value something has when humans would prefer that it existed, regardless of its uses (Stabell 2019). Existence value is closely related to subjective weak intrinsic value—if something has subjective intrinsic value, then presumably (all else being equal) we would prefer that it exist, and vice versa. Both intrinsic valuation and existence valuation give us a non-instrumental reason to protect the thing which we value in that way.

Consider the following argument for protecting biodiversity:

1. Biodiversity is a property or feature of the biosphere which is valuable for its own sake (i.e., it has some form of weak intrinsic value/end value/existence value).
2. *Ceteris paribus*, we should try to protect things that are valuable for their own sake.
3. So, *ceteris paribus*, we should try to protect biodiversity.

Two questions should be asked about this argument. First, what is the reason to accept premise (1)? And second, how strong is the conclusion, given the *ceteris paribus* clause?

A major reason to accept premise (1) is the observation that many people do apparently hold a preference for biodiverse nature, independently of biodiversity's further actual or perceived uses. If enough people have a strong, use-independent preference for maintaining biodiversity, then we should treat it as potentially worth protection on that basis (Stabell 2019 provides an extended argument for this conditional). Similar justifications can be given for the maintenance of art museums, classic works of literature, or catalogues of rare insects.

intrinsic value would be considered objective if you believe that values do not depend on the existence of valuers, and if you believe that weak intrinsic value derives from something's non-subjective properties. Someone who believes in objective aesthetic properties might endorse objective weak intrinsic value.

Suppose you accept premise (1). The sticking point is how far this argument gets us to a justification for drastic and costly measures to protect biodiversity rather than other potential uses of resources. That something is intrinsically valuable in the weaker sense or has existence value to humans does not justify expending arbitrarily large amounts of resources to save it, nor does it mean that the thing's value is arbitrarily or infinitely high (*pace* McCauley 2006). If an art museum catches fire, we should save the humans inside before trying to save the art, even supposing that the art has intrinsic value (in the weak sense). We can also distinguish between more and less valuable pieces of art based on additional considerations including their rarity, popularity, expert evaluations of the skill of the artist, etc. So, just because something is weakly intrinsically valuable does not mean that we must try to save it at any cost, or that it should be valued equally to everything else with intrinsic value.

So, this intrinsic value argument does not provide a justification to save biodiversity at any cost. In my view this is a feature rather than a bug, since most people do not take saving biodiversity to be an overriding goal. It is instead one goal among many—protecting the climate, protecting human quality of life, protecting particular ecosystems, protecting local economies, not placing undue burdens on developing countries—and hard decisions have to be made about how to set priorities where these goals conflict or resources are limited. A similar point is commonly made about aesthetic value. For example, Elliott Sober writes: “it is the material comforts of civilization that make possible a serious concern for both aesthetic and environmental values.... Such values are not frivolous, but they can become important to us only after certain fundamental human needs are satisfied” (Sober 1986, 191). The weak conception of intrinsic value does not attribute infinite value to its bearers so it allows that important instrumental needs can take priority over saving intrinsically valued entities.

A related position has been taken by Kristen Shrader-Frechette (1996) on her “hierarchical holism” view, which I interpret as ascribing weak intrinsic value to various environmental properties. She argues against both “extreme individualistic” ethics, which fail to consider collective or environmental health, and against strong forms of ecophilism such as Callicott’s view which may give precedence to ecological wholes over human individuals. On her view, certain “strong” human rights must take precedence over protecting the environment, but we should still ascribe values to the environment which might take precedence over less important human rights. I am not sure that a “hierarchy of rights” is the best way to formulate our ethical thinking, but I agree with her insight that we can ascribe some form of morally weighty value to features of the environment without falling into a strong and problematic form of ecophilism.

Ironically, intrinsic value defenses have been criticized for being both too demanding and too undemanding. Some have argued that intrinsic value does not effectively motivate conservation work or cannot be prioritized over instrumental value (Maguire and Justus 2008). At the same time, there have been suggestions that intrinsic value is “infinite” and overrides other kinds of value (McCauley 2006). I have argued that the second claim is incorrect, at least for weak intrinsic value. As for the first claim, it is worth maintaining a plurality of values in conservation decision-making. Intrinsic valuation may well help to motivate conservation efforts for some people, and especially if global biodiversity lacks instrumental value, we should be more hesitant to discount other forms of value. At the same time, intrinsic valuation does not need to be the primary basis for setting specific management priorities or for designing specific interventions. For instance, intrinsic valuation might serve to motivate concern about biodiversity loss. But decisions about reserve design should also consider instrumental values of local ecosystems. There is certainly a problem in environmental ethics about weighing very different forms of value against

each other. However, responding to this problem by urging that one form of value be rejected (whether intrinsic or instrumental) is counterproductive, since it reduces the number of available tools for motivating conservation work.

Some of the philosophers who are critical of intrinsic value arguments have claimed that intrinsic value is a metaphysically objectionable or confusing concept (Odenbaugh 2003; Justus et al. 2009). I am sympathetic to these concerns as applied to certain interpretations of intrinsic value. However, I have indicated that weaker forms of intrinsic value need not be metaphysically onerous. In fact, if you adopt a subjectivist conception of weak intrinsic or existence value, then it may become an empirical matter whether people ascribe this value to the relevant entities, though of course this position raises some problems of its own (e.g., the kinds of problem associated with moral relativism).

There are various remaining questions about how to decide which entities are intrinsically valuable; how to deal with disagreements about the preceding; how to incorporate intrinsic valuation into decision-making procedures; and so forth. However, given the failure of instrumental justifications for protecting global biodiversity, it is worth taking intrinsic value defenses seriously and continuing to address these questions.

5.6.4 Experience Values

For the sake of length, I will not attempt to cover aesthetic and other experience-value arguments in great detail. But I will briefly show that if global biodiversity or biological heterogeneity has experience value of some kind, then these arguments also operate at the correct scale so they are *prima facie* acceptable as biodiversity defenses. By “experience value” I mean to include potential sources of value such as:

- The aesthetic properties of biodiversity,
- The potential of (experiences of) biodiversity to be productive of human wellbeing or partly constitutive of human flourishing,
- The potential of (experiences of) biodiversity to transform humans or human pro-environmental preferences,
- The potential of engaging in or supporting biodiversity conservation to contribute to pro-environmental attitudes in the public,
- The connection of native biodiversity with local cultural heritage,
- The contributions of biodiversity to science education and scientific knowledge generation.

As I mentioned previously, philosophers differ about whether aesthetic value and similar forms of value should be considered intrinsic or instrumental (e.g., whether aesthetically valuable entities are valuable for their own sake because of their aesthetic properties, or whether they are valued for their ability to provide satisfying experiences to humans). I will not take a position about this issue. If you are committed to instrumental defenses, then appealing to some form of experience or scientific value, construed as a type of instrumental value, may be a way to rescue instrumental defenses in view of the apparent failure of standard ecosystem services defenses.

If experience value arguments are formulated to avoid the scale problem, then they will also avoid patchiness objections. A common thought is that unless you are willing to argue that all components of nature are equally aesthetically valuable, such arguments will ascribe more value to some species or communities than to others. However, if you argue that biological diversity *per se* has aesthetic or experience value, then patchiness concerns are avoided, because the value attaches directly to diversity rather than to individual species.

Surprisingly, there is not much philosophical work attributing aesthetic value to biological diversity. Most of the work in environmental aesthetics is focused on the aesthetic value of ecosystems or more nebulous items such as “nature,” “wilderness,” or “scenery” (Thompson 1995; Saito 1998; Varandas 2015). Discussions of aesthetic value which are supposed to be about biodiversity tend to slip into talking about the aesthetic value of species rather than of diversity per se (Welchman 2020; Linquist 2020). So, the aesthetic value of biodiversity per se appears to be an empty niche for an environmental philosopher to fill.

5.7 Conclusion

When thinking about biodiversity defenses, it has been tempting to form the opinion that there is safety in numbers. There are a lot of proposed arguments for protecting biodiversity, and one might hope that at least a few of them will be successful, or that the arguments will somehow generate a cumulative case for protecting biodiversity, even if each individual argument is subject to objections. A similar convergence strategy is sometimes pursued elsewhere in applied ethics. For example, one might show that many different ethical frameworks all justify respecting competent patients’ decisions about their medical care. Thus, one can conclude that respecting patients’ decisions is well-justified in medical ethics, even if any individual ethical framework is subject to objections.

Unfortunately, I have demonstrated in this chapter that the arguments for the value of biodiversity are not as abundant as one might think. The scale problem shows that popular services arguments cannot be successful. They are unsuccessful in a more fundamental way than has been appreciated. The problem is not only that they suffer from some empirical objections or that they

have some unintuitive consequences. The problem is that they rest on a faulty implicit premise linking global species losses to local species losses.

In view of this, I have suggested that environmental philosophers should continue devoting attention to other forms of value that might be attributed to global biodiversity, such as intrinsic value, aesthetic value, existence value, and so forth.

In addition, it is important to be more attentive to scale when discussing the values of species diversity. Ecological processes and properties are notoriously scale-dependent, and phenomena across multiple scales should be considered when discussing environmental management problems (Desjardins et al. 2019). Finally, we should devote more scientific and philosophical attention to potential causal roles of species richness at the global scale. For instance, more knowledge of possible feedbacks or tipping points at the global scale resulting from interactions among climate change, ecosystem changes and species extinctions would help to inform conservation decision-making.

6.0 Conclusion: Major Themes and Future Directions

Functional concepts in ecology are interestingly different from those in other parts of biology. One key reason is that ecosystems are far more flexible than are many familiar organisms. Ecosystems can undergo major regime shifts while maintaining their identity. As a result, the same role functions of components can contribute to very different types of functioning within different ecosystem regimes.

In Chapter 2, I showed that values are indirectly relevant to role function assignments, but not in the ways some people assume. For instance, invasive or destructive species do have role functions (*pace* Garson 2016, 88). We can identify role functions which are, from our perspective, either beneficial or detrimental to a given ecosystem. Instead, value judgments may apply when deciding which ecosystem states should be subject to inquiry and explanation by ecological researchers. This is partly a choice of research focus, but the choice will influence role function assignments according to CCI because role functions are those activities which are stable across relevant selected ecosystem states.

It is not a given that the ecological states selected for focus within ecology are ones that we ought to focus on. Overwhelmingly, ecologists study the relatively “natural” states of ecosystems, though there is growing research on urban ecology, agricultural ecology, and so forth. The focus of ecologists on relatively unaltered ecosystems seems in tension with the fact that a majority of earth’s land surface is heavily altered by humans (Plumptre et al. 2021). This focus also seems in tension with the perspective on which humans are components of ecosystems rather than outside interveners. It remains to be seen whether this bias towards less-human-altered ecosystems influences role function assignments or functional explanations in important ways.

One potential concern arises from the fact that our attempts to manage the most delicate or novel parts of the biosphere are currently guided by theoretical frameworks which focus on systems which are least affected by humans. This fact merits further consideration among ecologists and philosophers of science.

In Chapters 3 and 4, I showed that the concepts of *ecosystem functioning* and *ecosystem health* are:

- a. normative—they represent ecosystems as systems with certain ways of functioning that are appropriate or characteristic;
- b. evaluative—they treat some ecosystem states as better or preferable to others; and
- c. anthropocentric—they make the above judgments from the perspective of human-centric needs and preferences about the environment.

It has been widely recognized that *ecosystem health* is normative, and there has been some previous discussion about whether it is appropriate for ecologists to study a value-laden concept like this and about how exactly to interpret the concept (Lackey 2001). In contrast, to my knowledge, the value-laden nature of *ecosystem functioning* has not previously been discussed among philosophers and has less frequently been commented on by ecological researchers. This will be a fruitful case for further study. It is especially important to examine the values implicit in assessments of ecosystem functioning given that restoring ecosystem functioning is often the end goal of management projects (E. Jones 2021).

Because ecological research is often explicitly motivated by conservation and management concerns, it may not be possible to disentangle “basic” ecological research from the more obviously applied and value-laden enterprise of conservation biology (Soulé 1985). Conservation goals will influence ecological research at the level of experimental design and data interpretation.

I showed why this will occur in the case of research on ecosystem functioning: experimenters must make judgments about what ecosystems are like when they function *well* (a judgment sensitive to human interests) in order to decide what variables to measure in a study and how to interpret their data. This case should be of interest to other philosophers working on values in science, and is especially important because of the potential evidentiary relevance of BEF research to biodiversity conservation.

The fact that our understanding of ecosystem health is anthropocentric also merits further consideration among environmental ethicists. If my position is correct, then it may be impossible to design ecosystem interventions which are non-anthropocentric in nature. This represents a novel and important challenge to non-anthropocentrism as a general position in environmental philosophy.

In contrast with the above position about ecosystems, I argued in Chapter 5 that certain anthropocentric biodiversity defenses are unsuccessful. I showed that biodiversity defenses which appeal to intrinsic value are more plausibly successful than are those which appeal to instrumental value, at least given our current state of knowledge about causal relations among global-scale processes. This is opposite of the conventional wisdom among environmental philosophers, who have usually thought that biodiversity has instrumental value but might not have intrinsic value (Oksanen 1997). Clearly, further research is needed on the values of biodiversity, research which should attend more closely to the scale-relativity of ecological processes. Since the biodiversity crisis (and the climate crisis) are global-scale phenomena, it is vital to keep studying ecological phenomena at very large scales. We need detailed projections about how biodiversity and ecosystem functions are changing at the global scale. In addition, this case shows that detailed

engagement with empirical research is sometimes needed to generate sound moral arguments, given the causal complexity of ecosystems.

The above two points taken together further support the importance of value pluralism in environmental ethics. Both anthropocentric and non-anthropocentric values may be needed to justify and formulate management objectives at different scales. Regardless, given the speed at which the climate and ecosystems are changing, it is not possible or prudent for most of our environmental management work to be preservationist—i.e., to try and “dial back the clock” and restore ecosystems to a real or imagined pristine state (see Callicott et al. 1999). Instead, we need to think carefully about how we want earth’s ecosystems to function in a future that will be characterized by ongoing climate and ecological change.

Appendix A Intrinsic Value

Unfortunately, “intrinsic value” is used inconsistently throughout the environmental ethics literature. Here I will clarify how I am distinguishing intrinsic from instrumental value in this dissertation (particularly in Chapter 5).

In Chapter 5 I introduce a distinction between weak and strong intrinsic value, and assert that the distinction is orthogonal to that between objective and subjective value. So, I must acknowledge that some authors explicitly take intrinsic value to *mean* objective intrinsic value (e.g., Rolston 1982; MacLaurin and Sterelny 2008, 150; Justus et al. 2009). Some authors then consider existence value to be a type of instrumental value (e.g., Reyers et al. 2012). (That’s not all: Santana (2017) takes existence value to be a type of value distinct from both instrumental and intrinsic value—he takes instrumental value to mean economic value and intrinsic value to mean objective value.) These terminological differences mean that, for instance, an argument which I classify as intrinsic might be classified as instrumental by another author. Although one might read this as merely a verbal difference, I believe there are good reasons to prefer some uses of terminology over others.

For authors who treat subjective existence value as instrumental, they seem to take preference satisfaction to be an end, and treat the existence of an entity to be a means to satisfy that preference. While not *prima facie* an unreasonable position, I think it is misleading to classify subjective end or existence value as instrumental value (in this context) for the following reasons.

First, as mentioned in the main text, there is a metaphysical debate in the background about the nature of value. A viable position is that all value is subjective—i.e., all value depends on valuers. If this position is correct, then nothing has objective intrinsic value, so if intrinsic value is

defined as objective value, then nothing has intrinsic value at all. But this position is problematic—even if all value is subjective, we can make sense of a distinction between valuing something for its own sake and valuing something because it contributes to a distinct goal. It has been argued that failing to recognize this distinction, especially with respect to humans, amounts to a moral error (McShane 2007). So, objectivists who simply define intrinsic value as objective unfairly stack the deck against subjectivism, and subjectivists who define intrinsic value as objective undermine their own ability to make an important moral distinction.

Second, preference satisfaction is different in kind from other ends or uses which are usually appealed to in instrumental defenses. For example, promoting clean water and regulating the climate are ends to which ecosystems contribute; valuing ecosystems because they contribute to these ends are core examples of instrumental valuation. There is a further question why clean water and climate regulation are valued. Presumably, we have a preference to live in a world with clean water and a regulated climate. If so, then these instrumental values also ultimately reduce to preference satisfaction.

Notice the additional step in the instrumental chain:

<i>Intrinsic valuation</i>	ecosystem exists → satisfies a preference (for ecosystems)
<i>Instrumental valuation</i>	ecosystem exists → contributes to clean water → satisfies a preference (for clean water)

(For the sake of exposition, I am leaving aside endpoints alternative to preference satisfaction. An example of an alternative is supplied by an environmental virtue ethic which treats contact with the environment as partly constitutive of human flourishing (e.g., Cafaro 2001). On this view, the last item in the chains would read “promotes or partly constitutes flourishing” rather than “satisfies a preference.”)

In environmental ethics, the purpose of calling some arguments “instrumental” is to point out the external step in the value chain. Canonically instrumental arguments point to an effect of an ecosystem (or biodiversity, species, etc.), where that effect provides the main justification for caring about the entity in question.

For other philosophical purposes it might make sense to draw distinctions in different places, but I think this is the most perspicuous way to divide up arguments in our context. Even if some environmentalists reject the concept or existence of subjective intrinsic value, my way of distinguishing intrinsic from instrumental arguments is far from novel. For example, Newman et al. (2017, chap. 1) distinguish between objective and non-objective forms of intrinsic value in a similar way (they draw more distinctions than I do, but they do hold that there are subjective forms of intrinsic value). There are also other environmental philosophers who attribute subjective forms of intrinsic value to nature (e.g., Callicott 2017). And there are plenty of other philosophers who have been subjectivists about value in general. So, the possibility of subjective intrinsic value should not be excluded as a matter of definition.

Environmental philosophers have occasionally attacked the distinction between intrinsic and instrumental value. The main allegations are that (1) no clear positive definition of intrinsic value has been produced, and (2) for some instances of valuation, it is unclear whether they should be considered intrinsic or instrumental (Norton 2008). These are both basically vagueness complaints about the concept of intrinsic value. If my discussion has made it plausible that there is a sensible and morally relevant difference between “direct” and instrumental modes of valuation, then it is reasonable to retain the distinction even if it involves some vagueness.

Bibliography

- Adam, Thomas C., Russell J. Schmitt, Sally J. Holbrook, Andrew J. Brooks, Peter J. Edmunds, Robert C. Carpenter, and Giacomo Bernardi. 2011. “Herbivory, Connectivity, and Ecosystem Resilience: Response of a Coral Reef to a Large-Scale Perturbation.” *PLoS ONE* 6 (8): e23717. <https://doi.org/10.1371/journal.pone.0023717>.
- Ahn, Soohyun. 2020. “How Non-Epistemic Values Can Be Epistemically Beneficial in Scientific Classification.” *Studies in History and Philosophy of Science Part A* 84: 57–65. <https://doi.org/10.1016/j.shpsa.2020.08.002>.
- Alexandrova, Anna. 2018. “Can the Science of Well-Being Be Objective?” *The British Journal for the Philosophy of Science* 69: 421–45. <https://doi.org/10.1093/bjps/axw027>.
- Allen, Colin, and Jacob Neal. 2020. “Teleological Notions in Biology.” In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/spr2020/entries/teleology-biology/>.
- Amundson, Ron, and George V. Lauder. 1994. “Function Without Purpose: The Uses of Causal Role Function in Evolutionary Biology.” *Biology & Philosophy* 9 (4): 443–69. <https://doi.org/10.1007/BF00850375>.
- Arias-Arévalo, Paola, Erik Gómez-Baggethun, Berta Martín-López, and Mario Pérez-Rincón. 2018. “Widening the Evaluative Space for Ecosystem Services: A Taxonomy of Plural Values and Valuation Methods.” *Environmental Values* 27 (1): 29–53. <https://doi.org/10.3197/096327118X15144698637513>.
- Beatty, John. 1997. “Why Do Biologists Argue like They Do?” *Philosophy of Science* 64: S432–43. <https://doi.org/10.1086/392620>.

- Beisner, B. E., D. T. Haydon, and K. Cuddington. 2003. "Alternative Stable States in Ecology." *Frontiers in Ecology and the Environment* 1 (7): 376–82. [https://doi.org/10.1890/1540-9295\(2003\)001\[0376:ASSIE\]2.0.CO;2](https://doi.org/10.1890/1540-9295(2003)001[0376:ASSIE]2.0.CO;2).
- Bekessy, S. A., M. C. Runge, A. M. Kusmanoff, D. A. Keith, and B. A. Wintle. 2018. "Ask Not What Nature Can Do for You: A Critique of Ecosystem Services as a Communication Strategy." *Biological Conservation* 224: 71–74. <https://doi.org/10.1016/j.biocon.2018.05.017>.
- Bigelow, John, and Robert Pargetter. 1987. "Functions." *The Journal of Philosophy* 84 (4): 181–96. <https://doi.org/10.2307/2027157>.
- Boorse, Christopher. 1977. "Health as a Theoretical Concept." *Philosophy of Science* 44 (4): 542–73.
- Brook, Barry W., Erle C. Ellis, Michael P. Perring, Anson W. Mackay, and Linus Blomqvist. 2013. "Does the Terrestrial Biosphere Have Planetary Tipping Points?" *Trends in Ecology & Evolution* 28 (7): 396–401. <https://doi.org/10.1016/j.tree.2013.01.016>.
- Brose, Ulrich, and Helmut Hillebrand. 2016. "Biodiversity and Ecosystem Functioning in Dynamic Landscapes." *Philosophical Transactions of the Royal Society B: Biological Sciences* 371 (1694): 20150267. <https://doi.org/10.1098/rstb.2015.0267>.
- Broughton, Jack M, and Elic M Weitzel. 2018. "Population Reconstructions for Humans and Megafauna Suggest Mixed Causes for North American Pleistocene Extinctions." *Nature Communications* 9 (5441). <https://doi.org/10.1038/s41467-018-07897-1>.
- Burch-Brown, Joanna, and Alfred Archer. 2017. "In Defence of Biodiversity." *Biology & Philosophy* 32 (6): 969–97. <https://doi.org/10.1007/s10539-017-9587-x>.

- Burley, Hugh M., Karel Mokany, Simon Ferrier, Shawn W. Laffan, Kristen J. Williams, and Tom D. Harwood. 2016. "Macroecological Scale Effects of Biodiversity on Ecosystem Functions under Environmental Change." *Ecology and Evolution* 6 (8): 2579–93. <https://doi.org/10.1002/ece3.2036>.
- Callicott, J. Baird. 1992. "Aldo Leopold's Metaphor." In *Ecosystem Health: New Goals for Environmental Management*, edited by Robert Costanza, Bryan G. Norton, and Benjamin D. Haskell, 42–56. Washington, D.C.: Island Press.
- . 1995. "The Value of Ecosystem Health." *Environmental Values* 4 (4): 345–61. <https://doi.org/10.3197/096327195776679448>.
- . 2003. "Wetland Gloom and Wetland Glory." *Philosophy & Geography* 6 (1): 33–45. <https://doi.org/10.1080/1090377032000063306>.
- . 2017. "What Good Is It, Anyway?" In *The Routledge Handbook of Philosophy of Biodiversity*, edited by Justin Garson, Anya Plutynski, and Sahotra Sarkar, 168–82. London: Routledge.
- Callicott, J. Baird, Larry B. Crowder, and Karen Mumford. 1999. "Current Normative Concepts in Conservation." *Conservation Biology* 13 (1): 22–35. <https://doi.org/10.1046/j.1523-1739.1999.97333.x>.
- Cardinale, Bradley. 2014. "Overlooked Local Biodiversity Loss." *Science* 344 (6188): 1098–1098. <https://doi.org/10.1126/science.344.6188.1098-a>.
- Cardinale, Bradley J., Andrew Gonzalez, Ginger R.H. Allington, and Michel Loreau. 2018. "Is Local Biodiversity Declining or Not? A Summary of the Debate over Analysis of Species Richness Time Trends." *Biological Conservation* 219: 175–83. <https://doi.org/10.1016/j.biocon.2017.12.021>.

- Carpenter, Steve, Brian Walker, J. Marty Anderies, and Nick Abel. 2001. "From Metaphor to Measurement: Resilience of What to What?" *Ecosystems* 4 (8): 765–81. <https://doi.org/10.1007/s10021-001-0045-9>.
- Cartwright, Nancy, and Jeremy Hardie. 2012. *Evidence-Based Policy: A Practical Guide to Doing It Better*. Oxford: Oxford University Press.
- Clements, Frederic E. 1936. "Nature and Structure of the Climax." *The Journal of Ecology* 24 (1): 252. <https://doi.org/10.2307/2256278>.
- Coggan, Nicole V., Matthew W. Hayward, and Heloise Gibb. 2018. "A Global Database and 'State of the Field' Review of Research into Ecosystem Engineering by Land Animals." *Journal of Animal Ecology* 87 (4): 974–94. <https://doi.org/10.1111/1365-2656.12819>.
- Costanza, Robert, and Michael Mageau. 1999. "What Is a Healthy Ecosystem?" *Aquatic Ecology* 33 (1): 105–15.
- Cronon, William. 1996. "The Trouble with Wilderness: Or, Getting Back to the Wrong Nature." *Environmental History* 1 (1): 7–28. <https://doi.org/10.2307/3985059>.
- Cummins, Robert. 1975. "Functional Analysis." *The Journal of Philosophy* 72 (20): 741–65.
- Cvetkoska, Aleksandra, Elena Jovanovska, Alexander Francke, Slavica Tofilovska, Hendrik Vogel, Zlatko Levkov, Timme H. Donders, Bernd Wagner, and Friederike Wagner-Cremer. 2016. "Ecosystem Regimes and Responses in a Coupled Ancient Lake System from MIS 5b to Present: The Diatom Record of Lakes Ohrid and Prespa." *Biogeosciences* 13 (10): 3147–62. <https://doi.org/10.5194/bg-13-3147-2016>.
- deLaplante, Kevin, and Valentin Picasso. 2011. "The Biodiversity-Ecosystem Function Debate in Ecology." In *Philosophy of Ecology*, edited by Kevin deLaplante, Bryson Brown, and Kent

- A. Peacock, 169–200. Amsterdam: Elsevier. <https://doi.org/10.1016/B978-0-444-51673-2.50007-8>.
- Deliège, Glenn, and Stijn Neuteleers. 2015. “Should Biodiversity Be Useful? Scope and Limits of Ecosystem Services as an Argument for Biodiversity Conservation.” *Environmental Values* 24 (2): 165–82. <https://doi.org/10.3197/096327114X13947900181275>.
- Desjardins, Eric. 2015. “Historicity and Ecological Restoration.” *Biology & Philosophy* 30 (1): 77–98. <https://doi.org/10.1007/s10539-014-9467-6>.
- Desjardins, Eric, Gillian Barker, Zoë Lindo, Catherine Dieleman, and Antoine C. Dussault. 2015. “Promoting Resilience.” *The Quarterly Review of Biology* 90: 147–65.
- Desjardins, Eric, Justin Donhauser, and Gillian Barker. 2019. “Ecological Historicity, Novelty and Functionality in the Anthropocene.” *Environmental Values* 28 (3): 275–303. <https://doi.org/10.3197/096327119X15519764179791>.
- Díaz, Sandra, Josef Settele, Eduardo Brondízio, Hien T. Ngo, Maximilien Guèze, John Agard, Almut Arneth, et al. 2019. *Summary for Policymakers of the Global Assessment Report on Biodiversity and Ecosystem Services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*. Bonn, Germany: IPBES Secretariat. <https://doi.org/10.5281/zenodo.3553579>.
- Dietrich, Michael R. 2020. “What Is the Nature of Scientific Controversies in the Biological Sciences?” In *Philosophy of Science for Biologists*, edited by Kostas Kampourakis and Tobias Uller, 1st ed., 235–54. Cambridge University Press. <https://doi.org/10.1017/9781108648981.013>.

- Díaz, Sandra, and Marcelo Cabido. 2001. “Vive La Différence: Plant Functional Diversity Matters to Ecosystem Processes.” *Trends in Ecology & Evolution* 16 (11): 646–55. [https://doi.org/10.1016/S0169-5347\(01\)02283-2](https://doi.org/10.1016/S0169-5347(01)02283-2).
- Doak, D. F., D. Bigger, E. K. Harding, M. A. Marvier, R. E. O’Malley, and D. Thomson. 1998. “The Statistical Inevitability of Stability-Diversity Relationships in Community Ecology.” *The American Naturalist* 151 (3): 264–76. <https://doi.org/10.1086/286117>.
- Döring, Thomas F., Anja Vieweger, Marco Pautasso, Mette Vaarst, Maria R. Finckh, and Martin S. Wolfe. 2015. “Resilience as a Universal Criterion of Health.” *Journal of the Science of Food and Agriculture* 95 (3): 455–65. <https://doi.org/10.1002/jsfa.6539>.
- Dornelas, M., N. J. Gotelli, B. McGill, H. Shimadzu, F. Moyes, C. Sievers, and A. E. Magurran. 2014. “Assemblage Time Series Reveal Biodiversity Change but Not Systematic Loss.” *Science* 344 (6181): 296–99. <https://doi.org/10.1126/science.1248484>.
- Douglas, Heather. 2000. “Inductive Risk and Values in Science.” *Philosophy of Science* 67 (4): 559–79. <https://doi.org/10.1086/392855>.
- . 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
- Downing, Amy L., Bryan L. Brown, and Mathew A. Leibold. 2014. “Multiple Diversity–Stability Mechanisms Enhance Population and Community Stability in Aquatic Food Webs.” *Ecology* 95 (1): 173–84. <https://doi.org/10.1890/12-1406.1>.
- Dussault, Antoine C. 2018. “Functional Ecology’s Non-Selectionist Understanding of Function.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 70: 1–9. <https://doi.org/10.1016/j.shpsc.2018.05.001>.

- Dussault, Antoine C., and Frédéric Bouchard. 2017. “A Persistence Enhancing Propensity Account of Ecological Function to Explain Ecosystem Evolution.” *Synthese* 194: 1115–45.
- Ehrenfeld, David. 1992. “Ecosystem Health and Ecological Theories.” In *Ecosystem Health: New Goals for Environmental Management*, edited by Robert Costanza, Bryan G. Norton, and Benjamin D. Haskell, 135–43. Washington, D.C.: Island Press.
- Ehrlich, Paul R., and Anne H. Ehrlich. 1981. *Extinction: The Causes and Consequences of the Disappearance of Species*. New York: Random House.
- Ehrlich, Paul, and Brian Walker. 1998. “Rivets and Redundancy.” *BioScience* 48 (5): 387–387. <https://doi.org/10.2307/1313377>.
- Elton, Charles. 1927. *Animal Ecology*. New York: The Macmillan Company.
- Finegan, Bryan, Marielos Peña-Claros, Alexandre de Oliveira, Nataly Ascarrunz, M. Sydonia Bret-Harte, Geovana Carreño-Rocabado, Fernando Casanoves, et al. 2015. “Does Functional Trait Diversity Predict Above-Ground Biomass and Productivity of Tropical Forests? Testing Three Alternative Hypotheses.” *Journal of Ecology* 103 (1): 191–201. <https://doi.org/10.1111/1365-2745.12346>.
- Folke, Carl, Steve Carpenter, Brian Walker, Marten Scheffer, Thomas Elmqvist, Lance Gunderson, and C.S. Holling. 2004. “Regime Shifts, Resilience, and Biodiversity in Ecosystem Management.” *Annual Review of Ecology, Evolution, and Systematics* 35 (1): 557–81. <https://doi.org/10.1146/annurev.ecolsys.35.021103.105711>.
- Frank, David M. 2022. “Science and Values in the Biodiversity-Ecosystem Function Debate.” *Biology & Philosophy* 37: 7. <https://doi.org/10.1007/s10539-022-09835-4>.
- Frelich, Lee E, Bernd Blossey, Erin K Cameron, Andrea Dávalos, Nico Eisenhauer, Timothy Fahey, Olga Ferlian, et al. 2019. “Side-Swiped: Ecological Cascades Emanating from

- Earthworm Invasions.” *Frontiers in Ecology and the Environment* 17 (9): 502–10.
<https://doi.org/10.1002/fee.2099>.
- Gagic, V., I. Bartomeus, T. Jonsson, A. Taylor, C. Winqvist, C. Fischer, E. M. Slade, et al. 2015. “Functional Identity and Diversity of Animals Predict Ecosystem Functioning Better than Species-Based Indices.” *Proceedings of the Royal Society B: Biological Sciences* 282 (1801): 20142620–20142620. <https://doi.org/10.1098/rspb.2014.2620>.
- Garcia, Robert K., and Jonathan A. Newman. 2016. “Is It Possible to Care for Ecosystems? Policy Paralysis and Ecosystem Management.” *Ethics, Policy & Environment* 19 (2): 170–82.
<https://doi.org/10.1080/21550085.2016.1204054>.
- Garson, Justin. 2014. “What Is the Value of Historical Fidelity in Restoration?” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 45: 97–100. <https://doi.org/10.1016/j.shpsc.2013.10.003>.
- . 2016. *A Critical Overview of Biological Functions*. SpringerBriefs in Philosophy. Springer International Publishing. <https://doi.org/10.1007/978-3-319-32020-5>.
- Giller, Paul S., Helmut Hillebrand, Ulrike-G. Berninger, Mark O. Gessner, Stephen Hawkins, Pablo Inchausti, Cheryl Inglis, et al. 2004. “Biodiversity Effects on Ecosystem Functioning: Emerging Issues and Their Experimental Test in Aquatic Environments.” *Oikos* 104 (3): 423–36. <https://doi.org/10.1111/j.0030-1299.2004.13253.x>.
- Gillette, Kinley, S. Andrew Inkpen, and C. Tyler DesRoches. 2021. “Does Environmental Science Crowd Out Non-Epistemic Values?” *Studies in History and Philosophy of Science Part A* 87: 81–92. <https://doi.org/10.1016/j.shpsa.2021.01.008>.

- Glackin, Shane N. 2019. "Grounded Disease: Constructing the Social from the Biological in Medicine." *The Philosophical Quarterly* 69 (275): 258–76. <https://doi.org/10.1093/pq/pqy063>.
- Godfrey-Smith, Peter. 1994. "A Modern History Theory of Functions." *Noûs* 28 (3): 344–62. <https://doi.org/10.2307/2216063>.
- Gonzalez, Andrew, Bradley J. Cardinale, Ginger R. H. Allington, Jarrett Byrnes, K. Arthur Endsley, Daniel G. Brown, David U. Hooper, Forest Isbell, Mary I. O'Connor, and Michel Loreau. 2016. "Estimating Local Biodiversity Change: A Critique of Papers Claiming No Net Loss of Local Diversity." *Ecology* 97 (8): 1949–60. <https://doi.org/10.1890/15-1759.1>.
- Gonzalez, Andrew, Rachel M. Germain, Diane S. Srivastava, Elise Filotas, Laura E. Dee, Dominique Gravel, Patrick L. Thompson, et al. 2020. "Scaling-up Biodiversity-Ecosystem Functioning Research." *Ecology Letters* 23 (4): 757–76. <https://doi.org/10.1111/ele.13456>.
- Gross, Kevin, Bradley J. Cardinale, Jeremy W. Fox, Andrew Gonzalez, Michel Loreau, H. Wayne Polley, Peter B. Reich, and Jasper van Ruijven. 2014. "Species Richness and the Temporal Stability of Biomass Production: A New Analysis of Recent Biodiversity Experiments." *The American Naturalist* 183 (1): 1–12. <https://doi.org/10.1086/673915>.
- Hallett, Lauren M., Joanna S. Hsu, Elsa E. Cleland, Scott L. Collins, Timothy L. Dickson, Emily C. Farrer, Laureano A. Gherardi, et al. 2014. "Biotic Mechanisms of Community Stability Shift along a Precipitation Gradient." *Ecology* 95 (6): 1693–1700. <https://doi.org/10.1890/13-0895.1>.
- Harwell, Mark A., John H. Gentile, Larry D. McKinney, John W. Tunnell, William C. Dennison, R. Heath Kelsey, Kiersten M. Stanzel, Gregory W. Stunz, Kim Withers, and Jace Tunnell. 2019. "Conceptual Framework for Assessing Ecosystem Health." *Integrated*

<https://doi.org/10.1002/ieam.4152>.

Hausman, Daniel M. 2012. “Health, Naturalism, and Functional Efficiency.” *Philosophy of Science* 79 (4): 519–41. <https://doi.org/10.1086/668005>.

Hautier, Y., D. Tilman, F. Isbell, E. W. Seabloom, E. T. Borer, and P. B. Reich. 2015. “Anthropogenic Environmental Changes Affect Ecosystem Stability via Biodiversity.” *Science* 348 (6232): 336–40. <https://doi.org/10.1126/science.aaa1788>.

Hawkins, Bradford A., Miguel Á Rodríguez, and Stephen G. Weller. 2011. “Global Angiosperm Family Richness Revisited: Linking Ecology and Evolution to Climate.” *Journal of Biogeography* 38 (7): 1253–66. <https://doi.org/10.1111/j.1365-2699.2011.02490.x>.

Hector, A., Y. Hautier, P. Saner, L. Wacker, R. Bagchi, J. Joshi, M. Scherer-Lorenzen, et al. 2010. “General Stabilizing Effects of Plant Diversity on Grassland Productivity through Population Asynchrony and Overyielding.” *Ecology* 91 (8): 2213–20. <https://doi.org/10.1890/09-1162.1>.

Hector, A., J. Joshi, M. Scherer-Lorenzen, B. Schmid, E. M. Spehn, L. Wacker, M. Weilenmann, et al. 2007. “Biodiversity and Ecosystem Functioning: Reconciling the Results of Experimental and Observational Studies.” *Functional Ecology* 21 (5): 998–1002. <https://doi.org/10.1111/j.1365-2435.2007.01308.x>.

Hector, Andy, and Robert Bagchi. 2007. “Biodiversity and Ecosystem Multifunctionality.” *Nature* 448 (7150): 188–90. <https://doi.org/10.1038/nature05947>.

Hillebrand, Helmut. 2004. “On the Generality of the Latitudinal Diversity Gradient.” *The American Naturalist* 163 (2): 192–211. <https://doi.org/10.1086/381004>.

- Hillebrand, Helmut, Bernd Blasius, Elizabeth T. Borer, Jonathan M. Chase, John A. Downing, Britas Klemens Eriksson, Christopher T. Filstrup, et al. 2018. “Biodiversity Change Is Uncoupled from Species Richness Trends: Consequences for Conservation and Monitoring.” *Journal of Applied Ecology* 55 (1): 169–84. <https://doi.org/10.1111/1365-2664.12959>.
- Holling, C. S. 1973. “Resilience and Stability of Ecological Systems.” *Annual Review of Ecology and Systematics* 4 (1): 1–23. <https://doi.org/10.1146/annurev.es.04.110173.000245>.
- Hooper, D. U., F. S. Chapin, J. J. Ewel, A. Hector, P. Inchausti, S. Lavorel, J. H. Lawton, et al. 2005. “Effects of Biodiversity on Ecosystem Functioning: A Consensus of Current Knowledge.” *Ecological Monographs* 75 (1): 3–35. <https://doi.org/10.1890/04-0922>.
- Howard, Bruce, Leon C. Braat, Rob J. F. Bugter, Esther Carmen, Rosemary S. Hails, Allan D. Watt, and Juliette C. Young. 2018. “Taking Stock of the Spectrum of Arguments for Biodiversity.” *Biodiversity and Conservation* 27 (7): 1561–74. <https://doi.org/10.1007/s10531-016-1082-1>.
- Humboldt, Alexander von. 2014. *Views of Nature*. Edited by Stephen T. Jackson and Laura Dassow Walls. Translated by Mark W. Person. Chicago: University of Chicago Press.
- Huneman, Philippe. 2011. “About the Conceptual Foundations of Ecological Engineering: Stability, Individuality and Values.” *Procedia Environmental Sciences* 9: 72–82. <https://doi.org/10.1016/j.proenv.2011.11.013>.
- Huston, M. A., L. W. Aarssen, M. P. Austen, B. S. Cade, J. D. Fridley, E. Garnier, J. P. Grime, et al. 2000. “No Consistent Effect of Plant Diversity on Productivity.” *Science* 289: 1255a. <https://doi.org/10.1126/science.289.5483.1255a>.

- Huston, M. A., and A. C. McBride. 2002. "Evaluating the Relative Strengths of Biotic versus Abiotic Controls on Ecosystem Processes." In *Biodiversity and Ecosystem Functioning: Synthesis and Perspectives*, edited by Michel Loreau, Shahid Naeem, and Pablo Inchausti, 47–60. Oxford: Oxford University Press.
- Inkpen, S Andrew. 2019. "Health, Ecology and the Microbiome." *ELife* 8: e47626. <https://doi.org/10.7554/eLife.47626>.
- IUCN. 2016. *A Global Standard for the Identification of Key Biodiversity Areas, Version 1.0*. 1st ed. Gland, Switzerland: IUCN.
- Jamieson, Dale. 1995. "Ecosystem Health: Some Preventative Medicine." *Environmental Values* 4: 333–44.
- Jax, Kurt. 2005. "Function and 'Functioning' in Ecology: What Does It Mean?" *Oikos* 111 (3): 641–48. <https://doi.org/10.1111/j.1600-0706.2005.13851.x>.
- . 2006. "Ecological Units: Definitions and Application." *The Quarterly Review of Biology* 81: 237–58.
- . 2007. "Can We Define Ecosystems? On the Confusion Between Definition and Description of Ecological Concepts." *Acta Biotheoretica* 55 (4): 341–55. <https://doi.org/10.1007/s10441-007-9024-7>.
- . 2010. *Ecosystem Functioning*. Cambridge: Cambridge University Press.
- Jeffrey, Richard C. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science* 23 (3): 237–46.
- Jones, Clive G., John H. Lawton, and Moshe Shachak. 1994. "Organisms as Ecosystem Engineers." *Oikos* 69 (3): 373–86. <https://doi.org/10.2307/3545850>.

- Jones, Elis. 2021. “Distinguishing Regeneration from Degradation in Coral Ecosystems: The Role of Value.” *Synthese*. <https://doi.org/10.1007/s11229-021-03023-9>.
- Justus, James. 2008. “Complexity, Diversity, and Stability.” In *A Companion to the Philosophy of Biology*, edited by Sahotra Sarkar and Anya Plutynski, 321–50. Malden, MA: Blackwell Publishing.
- . 2011. “A Case Study in Concept Determination: Ecological Diversity.” In *Philosophy of Ecology*, edited by Kevin deLaplanche, Bryson Brown, and Kent A. Peacock, 147–68. Amsterdam: Elsevier.
- Justus, James, Mark Colyvan, Helen Regan, and Lynn Maguire. 2009. “Buying into Conservation: Intrinsic versus Instrumental Value.” *Trends in Ecology & Evolution* 24 (4): 187–91. <https://doi.org/10.1016/j.tree.2008.11.011>.
- Justus, James, and Samantha Wakil. 2021. “The Algorithmic Turn in Conservation Biology: Characterizing Progress in Ethically-Driven Sciences.” *Studies in History and Philosophy of Science Part A* 88: 181–92. <https://doi.org/10.1016/j.shpsa.2021.05.013>.
- Keith, David A., Jon Paul Rodríguez, Kathryn M. Rodríguez-Clark, Emily Nicholson, Kaisu Aapala, Alfonso Alonso, Marianne Asmussen, et al. 2013. “Scientific Foundations for an IUCN Red List of Ecosystems.” *PLoS ONE* 8 (5): e62111. <https://doi.org/10.1371/journal.pone.0062111>.
- Lackey, Robert T. 2001. “Values, Policy, and Ecosystem Health.” *BioScience* 51 (6): 437. [https://doi.org/10.1641/0006-3568\(2001\)051\[0437:VPAEH\]2.0.CO;2](https://doi.org/10.1641/0006-3568(2001)051[0437:VPAEH]2.0.CO;2).
- Lancaster, Jill. 2000. “The Ridiculous Notion of Assessing Ecological Health and Identifying the Useful Concepts Underneath.” *Human and Ecological Risk Assessment: An International Journal* 6 (2): 213–22. <https://doi.org/10.1080/10807030009380056>.

- Lean, Christopher Hunter. 2018. "Indexically Structured Ecological Communities." *Philosophy of Science* 85 (3): 501–22. <https://doi.org/10.1086/697746>.
- . 2020. "Invasive Species and Natural Function in Ecology." *Synthese*. <https://doi.org/10.1007/s11229-020-02635-x>.
- . 2021. "Invasive Species Increase Biodiversity and, Therefore, Services: An Argument of Equivocations." *Conservation Science and Practice* 3 (12): e553. <https://doi.org/10.1111/csp2.553>.
- Lenton, Timothy M., and Hywel T.P. Williams. 2013. "On the Origin of Planetary-Scale Tipping Points." *Trends in Ecology & Evolution* 28 (7): 380–82. <https://doi.org/10.1016/j.tree.2013.06.001>.
- Leopold, Aldo. 1991. "Wilderness as a Land Laboratory." In *The River of the Mother of God and Other Essays*, edited by Susan L. Flader and J. Baird Callicott, 287–89. Madison, Wisconsin: University of Wisconsin Press.
- Linquist, Stefan. 2020. "Two (and a Half) Arguments for Conserving Biodiversity on Aesthetic Grounds." *Biology & Philosophy* 35: 6. <https://doi.org/10.1007/s10539-019-9726-7>.
- Little, Chelsea J., and Florian Altermatt. 2018. "Species Turnover and Invasion of Dominant Freshwater Invertebrates Alter Biodiversity–Ecosystem-Function Relationship." *Ecological Monographs* 88 (3): 461–80. <https://doi.org/10.1002/ecm.1299>.
- Loreau, Michel, and Claire de Mazancourt. 2013. "Biodiversity and Ecosystem Stability: A Synthesis of Underlying Mechanisms." *Ecology Letters* 16: 106–15. <https://doi.org/10.1111/ele.12073>.
- Ludwig, David. 2016. "Ontological Choices and the Value-Free Ideal." *Erkenntnis* 81 (6): 1253–72. <https://doi.org/10.1007/s10670-015-9793-3>.

- Ma, Xuanlong, Alfredo Huete, Caitlin E. Moore, James Cleverly, Lindsay B. Hutley, Jason Beringer, Song Leng, Zunyi Xie, Qiang Yu, and Derek Eamus. 2020. "Spatiotemporal Partitioning of Savanna Plant Functional Type Productivity along NATT." *Remote Sensing of Environment* 246: 111855. <https://doi.org/10.1016/j.rse.2020.111855>.
- MacArthur, Robert H., and Edward O. Wilson. 1967. *The Theory of Island Biogeography*. Princeton University Press.
- Maclaurin, James, and Kim Sterelny. 2008. *What Is Biodiversity?* Chicago: The University of Chicago Press.
- Maguire, Lynn A., and James Justus. 2008. "Why Intrinsic Value Is a Poor Basis for Conservation Decisions." *BioScience* 58 (10): 910–11. <https://doi.org/10.1641/B581002>.
- Maier, Donald S. 2012. *What's So Good About Biodiversity? A Call for Better Reasoning About Nature's Value*. Dordrecht: Springer.
- Malaterre, Christophe, Antoine C Dussault, Sophia Rousseau-Mermans, Gillian Barker, Beatrix E Beisner, Frédéric Bouchard, Eric Desjardins, et al. 2019. "Functional Diversity: An Epistemic Roadmap." *BioScience* 69 (10): 800–811. <https://doi.org/10.1093/biosci/biz089>.
- Manson, Neil A. 2002. "Formulating the Precautionary Principle." *Environmental Ethics* 24 (3): 263–74. <https://doi.org/10.5840/enviroethics200224315>.
- Marcellesi, Alexandre. 2015. "External Validity: Is There Still a Problem?" *Philosophy of Science* 82 (5): 1308–17. <https://doi.org/10.1086/684084>.
- Mason, Norman W. H., David Mouillot, William G. Lee, J. Bastow Wilson, and Heikki Setälä. 2005. "Functional Richness, Functional Evenness and Functional Divergence: The Primary Components of Functional Diversity." *Oikos* 111 (1): 112–18.
- May, Robert M. 1972. "Will a Large Complex System Be Stable?" *Nature* 238 (5364): 413–14.

- Mazel, Florent, Matthew W. Pennell, Marc W. Cadotte, Sandra Diaz, Giulio Valentino Dalla Riva, Richard Grenyer, Fabien Leprieur, et al. 2018. “Prioritizing Phylogenetic Diversity Captures Functional Diversity Unreliably.” *Nature Communications* 9 (1): 2888. <https://doi.org/10.1038/s41467-018-05126-3>.
- McCann, Kevin Shear. 2000. “The Diversity–Stability Debate.” *Nature* 405 (6783): 228–33. <https://doi.org/10.1038/35012234>.
- McCauley, Douglas J. 2006. “Selling Out on Nature.” *Nature* 443 (7107): 27–28. <https://doi.org/10.1038/443027a>.
- McShane, Katie. 2004. “Ecosystem Health.” *Environmental Ethics* 26: 227–45.
- . 2007. “Why Environmental Ethics Shouldn’t Give Up on Intrinsic Value.” *Environmental Ethics* 29 (1): 43–61.
- . 2017. “Is Biodiversity Intrinsically Valuable? (And What Might That Mean?).” In *The Routledge Handbook of Philosophy of Biodiversity*, edited by Justin Garson, Anya Plutynski, and Sahotra Sarkar, 155–67. London: Routledge.
- Merchant, Carolyn. 2007. *American Environmental History: An Introduction*. New York: Columbia University Press.
- Millikan, Ruth Garrett. 1989. “In Defense of Proper Functions.” *Philosophy of Science* 56 (2): 288–302. <https://doi.org/10.1086/289488>.
- Millstein, Roberta L. 2018. “Is Aldo Leopold’s ‘Land Community’ an Individual?” In *Individuation, Process, and Scientific Practices*, edited by Otávio Bueno, Ruey-Lin Chen, and Melinda Bonnie Fagan, 279–302. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780190636814.003.0013>.

- . 2020a. “Defending a Leopoldian Basis for Biodiversity: A Response to Newman, Varner, and Linquist.” *Biology & Philosophy* 35 (12). <https://doi.org/10.1007/s10539-019-9724-9>.
- . 2020b. “Functions and Functioning in Aldo Leopold’s Land Ethic and in Ecology.” *Philosophy of Science* 87 (5): 1107–18. <https://doi.org/10.1086/710619>.
- Mitchell, Sandra D. 1995. “Function, Fitness and Disposition.” *Biology & Philosophy* 10 (1): 39–54. <https://doi.org/10.1007/BF00851986>.
- Moi, Dieison A., Gustavo Q. Romero, Pablo A. P. Antiqueira, Roger P. Mormul, Franco Teixeira de Mello, and Claudia C. Bonecker. 2021. “Multitrophic Richness Enhances Ecosystem Multifunctionality of Tropical Shallow Lakes.” *Functional Ecology* 35 (4): 942–54. <https://doi.org/10.1111/1365-2435.13758>.
- Morelli, Federico, and Anders Pape Møller. 2015. “Concerns about the Use of Ecosystem Services as a Tool for Nature Conservation: From Misleading Concepts to Providing a ‘Price’ for Nature, but Not a ‘Value.’” *European Journal of Ecology* 1 (1): 68–70. <https://doi.org/10.1515/eje-2015-0009>.
- Moreno, Alvaro, and Matteo Mossio. 2015. *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Dordrecht: Springer.
- Mori, Akira S., Takuya Furukawa, and Takehiro Sasaki. 2013. “Response Diversity Determines the Resilience of Ecosystems to Environmental Change.” *Biological Reviews* 88 (2): 349–64. <https://doi.org/10.1111/brv.12004>.
- Mori, Akira S., Forest Isbell, and Rupert Seidl. 2018. “ β -Diversity, Community Assembly, and Ecosystem Functioning.” *Trends in Ecology & Evolution* 33 (7): 549–64. <https://doi.org/10.1016/j.tree.2018.04.012>.

- Mossio, M., C. Saborido, and A. Moreno. 2009. "An Organizational Account of Biological Functions." *The British Journal for the Philosophy of Science* 60 (4): 813–41. <https://doi.org/10.1093/bjps/axp036>.
- Neander, Karen. 1991. "Functions as Selected Effects: The Conceptual Analyst's Defense." *Philosophy of Science* 58 (2): 168–84. <https://doi.org/10.1086/289610>.
- Neuenkamp, Lena, Suzanne M. Prober, Jodi N. Price, Martin Zobel, and Rachel J. Standish. 2019. "Benefits of Mycorrhizal Inoculation to Ecological Restoration Depend on Plant Functional Type, Restoration Context and Time." *Fungal Ecology* 40: 140–49. <https://doi.org/10.1016/j.funeco.2018.05.004>.
- Newbold, Tim, Lawrence N. Hudson, Samantha L. L. Hill, Sara Contu, Igor Lysenko, Rebecca A. Senior, Luca Börger, et al. 2015. "Global Effects of Land Use on Local Terrestrial Biodiversity." *Nature* 520 (7545): 45–50. <https://doi.org/10.1038/nature14324>.
- Newman, Jonathan A. 2020. "Biodiversity, Ecosystem Functioning, and the Environmentalist Agenda: A Reply to Odenbaugh." *Biology & Philosophy* 35 (1): 17. <https://doi.org/10.1007/s10539-019-9721-z>.
- Newman, Jonathan A., Gary Varner, and Stefan Lindquist. 2017. *Defending Biodiversity: Environmental Science and Ethics*. Cambridge: Cambridge University Press.
- Noonan, Meghan, Shawn J. Leroux, and Luise Hermanutz. 2021. "Evaluating Forest Restoration Strategies After Herbivore Overbrowsing." *Forest Ecology and Management* 482: 118827. <https://doi.org/10.1016/j.foreco.2020.118827>.
- Norton, Bryan G., and Douglas Noonan. 2007. "Ecology and Valuation: Big Changes Needed." *Ecological Economics* 63 (4): 664–75. <https://doi.org/10.1016/j.ecolecon.2007.02.013>.

- Nunes-Neto, Nei de Freitas, Ricardo Santos do Carmo, and Charbel N. El-Hani. 2016. "Biodiversity and Ecosystem Functioning: An Analysis of the Functional Discourse in Contemporary Ecology." *Filosofia e História Da Biologia* 11 (2): 289–321.
- Nunes-Neto, Nei Freitas, and Charbel Niño El-Hani. 2011. "Functional Explanations in Biology, Ecology, and Earth System Science: Contributions from Philosophy of Biology." In *Brazilian Studies in Philosophy and History of Science*, edited by Décio Krause and Antonio Videira, 290:185–99. Dordrecht: Springer Netherlands.
https://doi.org/10.1007/978-90-481-9422-3_13.
- Nunes-Neto, Nei Freitas, Alvaro Moreno, and Charbel N. El-Hani. 2014. "Function in Ecology: An Organizational Approach." *Biology & Philosophy* 29 (1): 123–41.
<https://doi.org/10.1007/s10539-013-9398-7>.
- Nunes-Neto, Nei Freitas, Cristian Saborido, Charbel N. El-Hani, Blandina Felipe Viana, and Alvaro Moreno. 2016. "Function and Normativity in Social-Ecological Systems." *Filosofia e História Da Biologia* 11 (2): 259–87.
- Odenbaugh, Jay. 2003. "Values, Advocacy and Conservation Biology." *Environmental Values* 12 (1): 55–69.
- . 2007. "Seeing the Forest *and* the Trees: Realism about Communities and Ecosystems." *Philosophy of Science* 74 (5): 628–41. <https://doi.org/10.1086/525609>.
- . 2010. "On the Very Idea of an Ecosystem." In *New Waves in Metaphysics*, edited by Allan Hazlett, 240–58. London: Palgrave Macmillan.
- . 2019. "Functions in Ecosystem Ecology: A Defense of the Systemic Capacity Account." *Philosophical Topics* 47 (1): 167–80.

- . 2020. “Biodiversity, Ecosystem Functioning, and the Environmentalist Agenda.” *Biology & Philosophy* 35 (1): 16. <https://doi.org/10.1007/s10539-019-9723-x>.
- Oksanen, Markku. 1997. “The Moral Value of Biodiversity.” *Ambio* 26 (8): 541–45.
- Olds, Andrew D., Kylie A. Pitt, Paul S. Maxwell, and Rod M. Connolly. 2012. “Synergistic Effects of Reserves and Connectivity on Ecological Resilience.” *Journal of Applied Ecology* 49 (6): 1195–1203. [https://doi.org/10.1111/jpe.12002@10.1111/\(ISSN\)1365-2745.BES_WIN_2012](https://doi.org/10.1111/jpe.12002@10.1111/(ISSN)1365-2745.BES_WIN_2012).
- Ollerton, Jeff. 2017. “Pollinator Diversity: Distribution, Ecological Function, and Conservation.” *Annual Review of Ecology, Evolution, and Systematics* 48 (1): 353–76. <https://doi.org/10.1146/annurev-ecolsys-110316-022919>.
- Ostertag, Rebecca, Faith Inman-Narahari, Susan Cordell, Christian P. Giardina, and Lawren Sack. 2014. “Forest Structure in Low-Diversity Tropical Forests: A Study of Hawaiian Wet and Dry Forests.” Edited by Bruno Hérault. *PLoS ONE* 9 (8): e103268. <https://doi.org/10.1371/journal.pone.0103268>.
- Parker, Wendy. 2014. “Values and Uncertainties in Climate Prediction, Revisited.” *Studies in History and Philosophy of Science Part A* 46: 24–30. <https://doi.org/10.1016/j.shpsa.2013.11.003>.
- Pâslaru, Viorel. 2009. “Ecological Explanation between Manipulation and Mechanism Description.” *Philosophy of Science* 76 (5): 821–37. <https://doi.org/10.1086/605812>.
- Patrick, Christopher J., Kevin E. McCluney, Albert Ruhi, Andrew Gregory, John Sabo, and James H. Thorp. 2021. “Multi-Scale Biodiversity Drives Temporal Variability in Macrosystems.” *Frontiers in Ecology and the Environment* 19 (1): 47–56. <https://doi.org/10.1002/fee.2297>.

- Peralta, Ariane L., Jeffrey W. Matthews, and Angela D. Kent. 2010. "Microbial Community Structure and Denitrification in a Wetland Mitigation Bank." *Applied and Environmental Microbiology* 76 (13): 4207–15. <https://doi.org/10.1128/AEM.02977-09>.
- Persson, Erik. 2016. "Option Value, Substitutable Species, and Ecosystem Services." *Environmental Ethics* 38 (2): 165–81. <https://doi.org/10.5840/enviroethics201638214>.
- Peters, Uwe. 2020. "Values in Science: Assessing the Case for Mixed Claims." *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1712235>.
- Plas, Fons van der. 2019. "Biodiversity and Ecosystem Functioning in Naturally Assembled Communities." *Biological Reviews*, 1220–45. <https://doi.org/10.1111/brv.12499>.
- Plas, Fons van der, Pete Manning, Santiago Soliveres, Eric Allan, Michael Scherer-Lorenzen, Kris Verheyen, Christian Wirth, et al. 2016. "Biotic Homogenization Can Decrease Landscape-Scale Forest Multifunctionality." *Proceedings of the National Academy of Sciences* 113 (13): 3557–62. <https://doi.org/10.1073/pnas.1517903113>.
- Plumptre, Andrew J, Daniele Baisero, R. Travis Belote, Ella Vázquez-Domínguez, Soren Faurby, Włodzimierz Jędrzejewski, Henry Kiara, et al. 2021. "Where Might We Find Ecologically Intact Communities?" *Frontiers in Forests and Global Change* 4: 626635. <https://doi.org/10.3389/ffgc.2021.626635>.
- Raerinne, Jani. 2011. "Causal and Mechanistic Explanations in Ecology." *Acta Biotheoretica* 59: 251–71. <https://doi.org/10.1007/s10441-010-9122-9>.
- Raffard, Allan, Frédéric Santoul, Julien Cucherousset, and Simon Blanchet. 2019. "The Community and Ecosystem Consequences of Intraspecific Diversity: A Meta-Analysis." *Biological Reviews* 94 (2): 648–61. <https://doi.org/10.1111/brv.12472>.

- Rahbek, Carsten. 1995. "The Elevational Gradient of Species Richness: A Uniform Pattern?" *Ecography* 18 (2): 200–205.
- Rapport, David. 1995. "Ecosystem Health: Exploring the Territory." *Ecosystem Health* 1 (1): 5–13.
- . 1998. "Answering the Critics." In *Ecosystem Health*, edited by David Rapport, Robert Costanza, Paul R. Epstein, Connie Gaudet, and Richard Levins, 41–50. Malden, MA: Blackwell Science.
- Rapport, David, R. Costanza, and A. J. McMichael. 1998. "Assessing Ecosystem Health." *Trends in Ecology & Evolution* 13 (10): 6.
- Reyers, Belinda, Stephen Polasky, Heather Tallis, Harold A. Mooney, and Anne Larigauderie. 2012. "Finding Common Ground for Biodiversity and Ecosystem Services." *BioScience* 62 (5): 503–7. <https://doi.org/10.1525/bio.2012.62.5.12>.
- Ricklefs, Robert E., and Roger Earl Latham. 1993. "Global Patterns of Diversity in Mangrove Floras." In *Species Diversity in Ecological Communities*, edited by Robert E. Ricklefs and Dolph Schluter, 215–29. Chicago: The University of Chicago Press.
- Rolston, Holmes III. 2000. "Aesthetics in the Swamps." *Perspectives in Biology and Medicine* 43 (4): 584–97. <https://doi.org/10.1353/pbm.2000.0052>.
- Rudner, Richard. 1953. "The Scientist Qua Scientist Makes Value Judgments." *Philosophy of Science* 20 (1): 1–6. <https://doi.org/10.1086/287231>.
- Russow, Lilly-Marlene. 1995. "Ecosystem Health: An Objective Evaluation?" *Environmental Values* 4 (4): 363–69. <https://doi.org/10.3197/096327195776679466>.

- Saborido, C., M. Mossio, and A. Moreno. 2011. “Biological Organization and Cross-Generation Functions.” *The British Journal for the Philosophy of Science* 62 (3): 583–606. <https://doi.org/10.1093/bjps/axq034>.
- Saccone, Patrick, Kristiina Hoikka, and Risto Virtanen. 2017. “What If Plant Functional Types Conceal Species-Specific Responses to Environment? Study on Arctic Shrub Communities.” *Ecology* 98 (6): 1600–1612. <https://doi.org/10.1002/ecy.1817>.
- Sagoff, Mark. 1992. “Has Nature a Good of Its Own?” In *Ecosystem Health: New Goals for Environmental Management*, edited by Robert Costanza, Bryan G. Norton, and Benjamin D. Haskell, 57–71. Washington, D.C.: Island Press.
- Saito, Yuriko. 1998. “The Aesthetics of Unscenic Nature.” *The Journal of Aesthetics and Art Criticism* 56 (2): 101–11. <https://doi.org/10.2307/432249>.
- Sakschewski, Boris, Werner von Bloh, Alice Boit, Lourens Poorter, Marielos Peña-Claros, Jens Heinke, Jasmin Joshi, and Kirsten Thonicke. 2016. “Resilience of Amazon Forests Emerges from Plant Trait Diversity.” *Nature Climate Change* 6 (11): 1032–36. <https://doi.org/10.1038/nclimate3109>.
- Santana, Carlos. 2014. “Save the Planet: Eliminate Biodiversity.” *Biology & Philosophy* 29 (6): 761–80. <https://doi.org/10.1007/s10539-014-9426-2>.
- . 2017. “Biodiversity Eliminativism.” In *The Routledge Handbook of Philosophy of Biodiversity*, edited by Justin Garson, Anya Plutynski, and Sahotra Sarkar, 86–97. London: Routledge.
- . 2018. “Biodiversity Is a Chimera, and Chimeras Aren’t Real.” *Biology & Philosophy* 33: 15. <https://doi.org/10.1007/s10539-018-9626-2>.

- Sasaki, Takehiro, and William K. Lauenroth. 2011. "Dominant Species, Rather than Diversity, Regulates Temporal Stability of Plant Communities." *Oecologia* 166 (3): 761–68. <https://doi.org/10.1007/s00442-011-1916-1>.
- Shrader-Frechette, Kristin. 1996. "Individualism, Holism, and Environmental Ethics." *Ethics and the Environment* 1 (1): 55–69.
- Skillings, Derek. 2016. "Holobionts and the Ecology of Organisms: Multi-Species Communities or Integrated Individuals?" *Biology & Philosophy* 31 (6): 875–92. <https://doi.org/10.1007/s10539-016-9544-0>.
- Sober, Elliott. 1986. "Philosophical Problems for Environmentalism." In *The Preservation of Species: The Value of Biological Diversity*, edited by Bryan G. Norton, 173–94. Princeton: Princeton University Press.
- Soulé, Michael E. 1985. "What Is Conservation Biology?" *BioScience* 35 (11): 727–34. <https://doi.org/10.2307/1310054>.
- Spehn, E. M., A. Hector, J. Joshi, M. Scherer-Lorenzen, B. Schmid, E. Bazeley-White, C. Beierkuhnlein, et al. 2005. "Ecosystem Effects of Biodiversity Manipulations in European Grasslands." *Ecological Monographs* 75 (1): 37–63. <https://doi.org/10.1890/03-4101>.
- Stabell, Espen Dyrnes. 2019. "Existence Value, Preference Satisfaction, and the Ethics of Species Extinction." *Environmental Ethics* 41 (2): 165–80. <https://doi.org/10.5840/enviroethics201941215>.
- Steele, K. 2006. "The Precautionary Principle: A New Approach to Public Decision-Making?" *Law, Probability and Risk* 5 (1): 19–31. <https://doi.org/10.1093/lpr/mgl010>.
- Sterelny, Kim. 2001. "The Reality of Ecological Assemblages: A Palaeo-Ecological Puzzle." *Biology & Philosophy* 16 (4): 437–61. <https://doi.org/10.1023/A:1011972529013>.

- Steudel, Bastian, Christine Hallmann, Maike Lorenz, Stefan Abrahamczyk, Kathleen Prinz, Cornelia Herrfurth, Ivo Feussner, Johannes W. R. Martini, and Michael Kessler. 2016. “Contrasting Biodiversity–Ecosystem Functioning Relationships in Phylogenetic and Functional Diversity.” *New Phytologist* 212 (2): 409–20. <https://doi.org/10.1111/nph.14054>.
- Suding, Katharine N., Katherine L. Gross, and Gregory R. Houseman. 2004. “Alternative States and Positive Feedbacks in Restoration Ecology.” *Trends in Ecology & Evolution* 19 (1): 46–53. <https://doi.org/10.1016/j.tree.2003.10.005>.
- Suter, Glenn W. 1993. “A Critique of Ecosystem Health Concepts and Indexes.” *Environmental Toxicology and Chemistry* 12 (9): 1533–39. <https://doi.org/10.1002/etc.5620120903>.
- Szasz, Thomas S. 1960. “The Myth of Mental Illness.” *American Psychologist* 15 (2): 113–18.
- Tett, P., R. J. Gowen, S. J. Painting, M. Elliott, R. Forster, D. K. Mills, E. Bresnan, et al. 2013. “Framework for Understanding Marine Ecosystem Health.” *Marine Ecology Progress Series* 494: 1–27. <https://doi.org/10.3354/meps10539>.
- Thompson, Janna. 1995. “Aesthetics and the Value of Nature.” *Environmental Ethics* 17 (3): 291–305.
- Tilman, David. 1996. “Biodiversity: Population Versus Ecosystem Stability.” *Ecology* 77 (2): 350–63. <https://doi.org/10.2307/2265614>.
- Tilman, David, and John A. Downing. 1994. “Biodiversity and Stability in Grasslands.” *Nature* 367: 363–65.
- Tilman, David, Peter B. Reich, and Johannes M. H. Knops. 2006. “Biodiversity and Ecosystem Stability in a Decade-Long Grassland Experiment.” *Nature* 441 (7093): 629–32. <https://doi.org/10.1038/nature04742>.

- Tilman, David, Peter B. Reich, Johannes Knops, David Wedin, Troy Mielke, and Clarence Lehman. 2001. "Diversity and Productivity in a Long-Term Grassland Experiment." *Science* 294 (5543): 843–45. <https://doi.org/10.1126/science.1060391>.
- Tilman, David, David Wedlin, and Johannes M. H. Knops. 1996. "Productivity and Sustainability Influenced by Biodiversity in Grassland Ecosystems." *Nature* 379: 718–20.
- Varandas, Maria José. 2015. "The Land Aesthetic, Holmes Rolston's Insight." *Environmental Values* 24 (2): 209–26. <https://doi.org/10.3197/096327115X14247121957607>.
- Väyrynen, Pekka. 2008. "Slim Epistemology with a Thick Skin." *Philosophical Papers* 37 (3): 389–412. <https://doi.org/10.1080/05568640809485228>.
- . 2011. "Thick Concepts and Variability." *Philosophers' Imprint* 11 (1): 1–17. <https://doi.org/10.1093/acprof:oso/9780199314751.001.0001>.
- . 2013. *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics*. Oxford: Oxford University Press.
- Vellend, M., L. Baeten, I. H. Myers-Smith, S. C. Elmendorf, R. Beausejour, C. D. Brown, P. De Frenne, K. Verheyen, and S. Wipf. 2013. "Global Meta-Analysis Reveals No Net Change in Local-Scale Plant Biodiversity Over Time." *Proceedings of the National Academy of Sciences* 110 (48): 19456–59. <https://doi.org/10.1073/pnas.1312779110>.
- Vitousek, Peter M., and Lawrence R. Walker. 1989. "Biological Invasion by *Myrica faya* in Hawai'i: Plant Demography, Nitrogen Fixation, Ecosystem Effects." *Ecological Monographs* 59 (3): 247–65. <https://doi.org/10.2307/1942601>.
- Vucetich, John A., Jeremy T. Bruskotter, and Michael Paul Nelson. 2015. "Evaluating Whether Nature's Intrinsic Value Is an Axiom of or Anathema to Conservation." *Conservation Biology* 29 (2): 321–32. <https://doi.org/10.1111/cobi.12464>.

- Wang, Shaopeng, and Michel Loreau. 2016. “Biodiversity and Ecosystem Stability across Scales in Metacommunities.” *Ecology Letters* 19 (5): 510–18. <https://doi.org/10.1111/ele.12582>.
- Welchman, Jennifer. 2020. “Commentary on Jonathan A. Newman, Gary Varner, and Stefan Linquist: Defending Biodiversity: Environmental Science and Ethics, Chapter 11: Should Biodiversity Be Conserved for Its Aesthetic Value?” *Biology & Philosophy* 35: 13. <https://doi.org/10.1007/s10539-019-9720-0>.
- Westman, Walter E. 1978. “Measuring the Inertia and Resilience of Ecosystems.” *BioScience* 28 (11): 705–10. <https://doi.org/10.2307/1307321>.
- Wicklum, D., and Ronald W. Davies. 1995. “Ecosystem Health and Integrity?” *Canadian Journal of Botany* 73 (7): 997–1000. <https://doi.org/10.1139/b95-108>.
- Wimsatt, William C. 1972. “Teleology and the Logical Structure of Function Statements.” *Studies in History and Philosophy of Science Part A* 3 (1): 1–80. [https://doi.org/10.1016/0039-3681\(72\)90014-3](https://doi.org/10.1016/0039-3681(72)90014-3).
- Winsberg, Eric. 2012. “Values and Uncertainties in the Predictions of Global Climate Models.” *Kennedy Institute of Ethics Journal* 22 (2): 111–37. <https://doi.org/10.1353/ken.2012.0008>.
- Wolf, Evan C., David J. Cooper, and N. Thompson Hobbs. 2007. “Hydrologic Regime and Herbivory Stabilize an Alternative State in Yellowstone National Park.” *Ecological Applications* 17 (6): 1572–87. <https://doi.org/10.1890/06-2042.1>.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Wright, Larry. 1973. “Functions.” *The Philosophical Review* 82 (2): 139–68. <https://doi.org/10.2307/2183766>.

Zeder, Melinda A. 2011. "The Origins of Agriculture in the Near East." *Current Anthropology* 52 (S4): S221–35. <https://doi.org/10.1086/659307>.