

OTHER GAMES IN OTHER TOWNS: PLURALISM ABOUT BIOLOGICAL FUNCTION
AND REPRESENTATION

By

Robert Lazo

Copyright © by Robert Lazo 2024

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements

For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

UNIVERSITY OF ARIZONA

2024

Acknowledgements

Graduate school is very hard, and I could not have hoped to complete it without the help of many people. I have done my best to express my thanks for that help over time, but I'm very happy to have the opportunity to do so again now.

To start, thanks to my fellow graduate students. My cohort – Andrew Lichter, Will Schumacher, Jackie Sideris, and Shuai Liu – helped make the transition back into academia after a couple of years away easier than it could have been and more fun. Thanks also to the elder grads who quickly took me under their wings and provided invaluable advice and support: Robert Wallace, Brandon Ashby, Caroline King, Sarah Raskoff, Jacob Barrett, and Santi Borboa. As for those closer in cohort-year to myself, I am indebted to many friends for the same. In no particular order, thank you to Max Kramer, Aliya Dewey, Anna-Bella Sicilia, Lenin Vasquez-Toledo, Luke Goleman, Kaya Wisniewska, Nathan Oakes, Sean Whitton, Yili Zhou, Ke Zhang, Will Cailles, Jacob Blitz, Rhys Borchert, Danielle Guzman, Avital Hazony, Tim Kearn, Ritwik Agrawal, Travis Quigley, April Chan, Ella LaRose, Sabrina Gonzalez, Ding, Ben Keoseyan, Danilo Linhares, Dave Poplar, and many others. I would especially like to direct my gratitude towards the many members of the various TTRPG groups I was lucky enough to be a part of, as well as the writing and pedagogy support groups.

To my committee: thank you for your endless help. To my chair, Jonathan Weinberg, I could not have completed this dissertation without your teaching, kindness, and gentle prodding. Your efforts turned me – I hope! – into a half-decent philosopher, if still a teleosemanticist, and I promise to pay for your meal whenever we meet in the future. To Renee Duckworth, thank you for emailing the philosophy department inviting folks to join your evolutionary epigenetics seminar in my first year, and for inviting me to be a guest in your lab. You showed me how real live scientists think, an important antidote to years of mostly just imagining what you all were up to. To Terry Horgan, thank you for your continuous interest and support for a student who believes in a theory pretty close to the opposite of your own, and for insisting on catchy titles. At least we're both steadfast naturalists! To Beckett Sterner, thank you for being a wonderful guide through contemporary philosophy of biology, both as a discipline and a social circle. Your own work on pluralism about biological function and on conservation biology has been a deep inspiration.

There are many other professors to whom I owe gratitude for teaching me how to do philosophy and to think about biology: thanks to Jason Turner, Mark Timmons, Peg Weiser, Sara Aronowitz, Judie Bronstein, John Weins, Max Li, Mike Sanderson, Alex Badyaev, Houston Smit, Allen Buchanan, Chris Maloney, Joseph Tolliver, and Richard Healey. I also want to thank my undergraduate advisor, Joel Richeimer, whose classes on the philosophy of science and American pragmatism were the starting points that led me here, as well as Andrew and Clair Linzey of Oxford, who insisted I was cut out for this.

I would also like to thank the members of the department who keep the lights on. Sandra Kimball, Shauna Garland, Megan Gutierrez, Beth Zimmer, and Debbie Jackson, none of this would be possible without your tireless efforts.

Thanks also to some of my friends outside of the department: Elisa Pelgrift, Justine Orlovsky-Schnitzler, Meg Oakes, Andrew Stewart, and Drew Hogan. I'm so grateful to have been able to sometimes *not* talk or think about philosophy.

Thank you, of course, to my family. I could never have achieved any of this without all of you. Mom and Dad, thank you for helping me, believing in me, and loving me. Thank you for giving us such a wonderful home to grow up in that so loved books, art, food, and animals, and for showing me over and over again the joys of learning. Thank you also for showing me how we can all grow and change over time, and discover new things to do and love. Nicholas and James, thank you for being my brothers. I am so lucky to have two older brothers who are so kind – at least as adults! – and who are such excellent role models. You have pursued scary dreams, and have shown me it's worth it to do so. And I owe much of the above to my grandparents, Nancy and Ernest Stone and Rosemarie and Bob Lazo, who raised wonderful children who became my wonderful parents.

I would also like to thank my mother-in-law, Sandie Moffett, and her mother, Rose Stefanski, who have been nothing but welcoming and encouraging. And, of course, you raised my wife, whom I promise I'll thank shortly!

To all of the animals that I have loved and who have loved me, even when I've annoyed them, thank you for all of your playfulness, kindness, and comfort. Thank you to Yellow Dog, Gaffer, Seamus, Winston, Domino, Baggins, Pippin, Oreo et al., Hobo, Daisy, Scout, Audrey, and Mac. Thank you to Ganymede, Cesario, and Io. I hope that I have been able, in my own small way, to pay you all back for all that you have given me.

Finally, I want to thank my wife, Amanda. You have loved me and supported me through all of this, and much more besides. Thank you for being with me, for listening to me ramble about philosophy, and for moving to the desert so that I could do this in the first place. You are the inspiration for all that I do. I love you, and this dissertation is dedicated to you.

Table of Contents

Abstract.....	7
Introduction: Biological Functions, Ecology, and Evolutionary Biology.....	9
§1. Introduction.....	9
§2. Ecology and Evolutionary Biology.....	11
§3. Conclusion.....	20
Chapter 1: Realized Functions in Ecology.....	21
§1. Introduction.....	21
§2. Ecological Niche Modeling in Conservation Biology.....	22
§3. Realized Functions.....	29
§4. Objections and Replies.....	33
§4.1. The No-Real-Telicity Objection.....	34
§4.2. The Relativity Objection.....	43
§5. Conclusion.....	45
Chapter 2: What Fodor’s Frogs(’ Visual Systems) See.....	47
§1. Introduction.....	47
§2. Two Variants of the Problem.....	51
§3. Two Questions about the Frogs’ Visual Systems.....	53
§4. SE Functions and the Ecological Question.....	55
§5. The Determinate Realized Function of the Frogs’ Visual Systems.....	59
§6. Conclusion.....	63
Chapter 3: Swampman and the Problem of Novelty.....	64
§1. Introduction.....	64
§2. A Dilemma.....	66
§3. Papineau on Representations.....	69
§4. Swampman’s Realized Representations.....	73
§5. A Selected-Effects Solution to the Problems of Swampman and Novelty?.....	77
§6. Conclusion.....	82
Conclusion: Other Games in Other Towns.....	85
Works Cited.....	91

Abstract

At the heart of this dissertation is an argument about the differences in some of the explanatory practices of ecology and evolutionary biology. While evolutionary biology is often concerned with the causes of traits – their evolutionary history – ecology often studies traits *as* causes. I argue for this claim in the introduction, and the rest of the monograph concerns some of the consequences of this inversion of explanandum and explanans. Treating traits as causes rather than as effects has profound consequences for our understanding of functional explanation in biology, as well as in cognitive science, especially for teleosemantics. In general, treating traits only as effects leads to problems when we try to explain how organisms and populations respond to novelty in their environments. Treating traits as causes resolves this difficulty.

In the first chapter, I identify and develop a notion of function – realized function – implicitly used in some ecological niche models from conservation biology. A trait's realized function is its contribution to a population's ability to occupy its realized niche. I show why ecological niche models require this notion of function as well as how they make use of it via a case study: the use of an ecological niche model to predict the changes in the range of Sub-Saharan amphibian species due to anthropogenic climate change (Garcia et al. 2014).

In the second chapter, I re-examine the case of Fodor's frogs, showing how the notion of a realized function can help resolve the problem of indeterminacy for teleosemantics, according to which a representation's biological function is unable to provide a determinate content for that representation. I argue that a backward-looking function that *did* fix the content wouldn't be sufficient to resolve the problem this indeterminacy poses in ecology, but an appeal to the realized function of the frog's visual system can. As a result, teleosemanticists would be best served by adopting pluralism about the biological functions that can give contents to representations, becoming pluralists about representation as well.

In the third chapter, I show how this same pluralism about biological function can provide us with an answer to the challenge posed by Swampman. We do not need to deny that he would have any representations at all. We can instead say that he has realized representations but lacks selected-effects ones. I use this discussion as a springboard towards the more general problem of

novelty, according to which a backward-looking teleosemantics is unable to assign evolutionarily novel contents to representations. I consider the best-developed attempt to resolve this problem and argue that it fails for the same reasons a backward-looking teleosemantics generally fails to handle ecological questions. Again, pluralism about biological function and representation is our best bet to resolve the issue.

Finally, I conclude with some general remarks about the pluralism I've advocated for. There is no simple way to resolve the various notions of function I've used throughout into a single notion, and there is no "best game in town." Instead, there are many games in many towns. To do good work in biology and cognitive science, we must admit – at least for now – that there is more than one notion of function at play in these sciences. These various notions are more and less applicable in different explanatory and predictive contexts, but are also sometimes jointly required in order to better understand a system of interest. As a result, I advocate for integrative pluralism about biological function.

Introduction: Biological Functions, Ecology, and Evolutionary Biology

1. Introduction

I entered graduate school with the intent of pursuing research in teleosemantics, a theory of content according to which a representation's content is fixed by its function. I was and remain especially interested in how a proponent of the theory might respond to some of its most persistent criticisms: the problems of novelty, of content indeterminacy, and of Swampman. I quickly realized, though, that to do this work, I would need to know a lot more biology, and so I began graduate studies in ecology and evolutionary biology. This work led me to a second, seemingly independent research question. How do ecologists use the functions of organisms' traits to understand how those organisms fit into their environments, especially when it comes to conservation biology? Unlike the selected-effects functions studied by evolutionary biologists, which are historical, and the causal role functions studied by molecular biologists, which are often described independently of their impact on fitness, the functions ecologists are often interested in are what we might loosely term the present-utility functions of traits. They use trait-based models to describe organisms' niches – the environments the organisms could successfully inhabit. That is, they talk about what organisms' traits allow them to do in their environments and how these interactions help us to explain and predict their abilities to occupy those environments.

I am, of course, far from the first person to notice that different disciplines in biology make use of different notions of function. Pluralism about biological function has been the norm in the philosophy of biology since the 90s, if not a little earlier (see Preston 1998 and Godfrey-Smith 1993). It does seem to me, though, that the literature on teleosemantics has largely ignored this lesson. While admitting pluralism is likely true in the context of biology, most act as if there is only one notion of function to which the teleosemanticist can appeal, typically the notion of a selected-effects (SE) function. This is, I think, an odd assumption to make. If we can admit that traits have their functions in virtue of different underlying states of affairs – their history and their causal role, for instance – then we would expect the same to be true of representations, or that it at least *could* be. It could be that more than one notion of function underwrites claims about the content of representations.

I show the backward-looking teleosemanticist that, if we can get the right notions of function in place, then embracing this sort of pluralism could do them substantial good. It could

allow them to finally respond to some of their most stubborn critics in a way that not only satisfies themselves, but also satisfies those critics. Consider, for example, the trouble with Swampman, the exact physical duplicate of Davidson without any of his history. If the classical, backward-looking (i.e. endorsing a backward-looking notion of function) teleosemanticists are correct, then Swampman has no mental states, since he has no history and so no functions, but this result is, according to the theory's opponents, clearly wrong. Since Swampman is a physical duplicate, we can use exactly the same psychology, neuroscience, and cognitive science to predict and explain all of his behaviors as we would have used for Davidson. We know well ahead of time that Swampman will tell "his" family members that he loves them, because he has the right set of mental states. He represents "his" family members as, among other things, those he ought to love. More than just this, though, it seems likely that we *couldn't* explain or predict Swampman's behaviors without appealing to his mental states. Appeals just to physics and chemistry wouldn't be enough.¹

This argument from the indispensability of Swampman's mental states seems to me a powerful one for the same reasons I find it a powerful argument for the indispensability of mental states more generally. To borrow from Fodor, to explain why I'll be at the airport a week from now at the appointed time to pick up my friend, we'll need more than just physics (1987). We just can't get around in the world without believing in mental states. Having said that, it is surely also the case that Swampman's cognitive system is importantly different from our own. It wasn't shaped by natural selection, but instead by chance. Many of our explanations for why and how we represent will therefore differ significantly from those we deploy to explain the same for Swampman. My cognitive system represents snakes and snake-shaped objects as dangerous because that representation has been selected for over time. Swampman's cognitive system does the same merely because it happened to come into being with that feature.

Pluralism about biological function – assuming we have the right notions of biological function in mind – could allow us to take up the middle ground between these extreme responses to Swampman. One could admit that Swampman lacks the backward-looking functions of the classical teleosemanticist, but has some other sort of biological function – shared with the original Davidson and with us – that relies on an appeal to something other than the past. In this

¹ Notice that, if we *could* explain everything about Swampman without appealing to his cognitive states, the same would have to be true of *ourselves*, since Swampman is a physical duplicate of us (Sebastián 2017). I'll expand on this point in chapter 3.

way, one could say both that Swampman has mental states, thereby explaining why psychology does such a good job explaining and predicting his behavior, and that he lacks the sort of representation so central to our historical explanations for so many of Davidson's and our own cognitive features.

Of course, there's a great deal of work one needs to do to get to that point. That's the work I mean to accomplish in this dissertation. In the first chapter, I'll identify and develop a notion of function that I argue is found implicitly in some niche distribution models used in conservation biology and in ecology more generally. This notion of function relies on claims about how an organism's traits allow it to successfully occupy its realized niche by enabling certain interactions with its environment broadly construed. In the second chapter, I argue that we can better address the problem of indeterminacy posed by Fodor's frogs by appealing to the novel notion of function I discuss in chapter 1. In chapter three, I show how pluralism about notions of biological function can do the work described in the last few paragraphs. We can say both that Swampman's cognitive system lacks features important for understanding our own and that it is still capable of representing. I conclude with some general remarks on pluralism about representation and the future research directions it suggests.

For the moment, though, it will be helpful to lay out a pair of distinctions: first, between ecological and evolutionary research programs, and second, between the study of traits as effects and of traits as causes. Doing so will make clearer why appeals to the traits of organisms can serve distinct explanatory and predictive purposes. We sometimes mean to explain the present existence or distribution of a trait, but in other circumstances we mean to use a trait in order to explain and predict some future state of affairs. Research on one side of each of these distinctions will not always do well answering questions posed by the other, and some of the arguments surrounding Fodor's frogs and Swampman fail to take this into account.

2. Ecology and Evolutionary Biology

Not every biological science has the same explanatory aims when discussing what a trait allows an organism to do. Consider the differences between the explanations and explanatory structures found in ecology and its subdisciplines and the more philosophically familiar ones found in evolutionary biology. Most importantly for our purposes, where evolutionary biology concerns itself with the history and causes of evolution (Futuyma 2005), ecology studies "the interactions that determine the distribution and abundance of organisms" (Krebs 2001, p. 15).

That is, while evolutionary biology means to explain the presence of traits in a population, ecology is often more interested in what those traits allow a population to do and the ways in which these interactions with the environment explain a population's size and range. Ecology is by no means ahistorical, but it is often interested not in the history of a species *per se*, but in its present and future.

This distinction in subject was there from the beginning. Haeckel introduced the term “ecology” – or, rather, “œcology” – in 1866 and offered a more precise definition in 1870:

By ecology we mean the body of knowledge concerning the economy of nature – the investigation of the total relations of the [organism]² both to its inorganic and to its organic environment; including above all, its friendly and inimical relations with those [other organisms] with which it comes directly or indirectly into contact – in a word, ecology is the study of all those complex interrelations referred to by Darwin as the conditions of the struggle for existence (Allee et al. 1949).³

Here we can see both the tight connection between ecology and evolutionary biology and some of the ways in which they pull apart. Modern ecological science began to establish itself in the latter half of the 19th century as part of the Darwinian research programme. In doing so, though, it set out a distinct subject of study. Haeckel identifies ecological knowledge as that which underpins many⁴ of our claims about evolution. To know which evolutionary changes a lineage

² Haeckel uses the word “animal,” but I’ve substituted “organism” to better match onto the more expansive definition in use today.

³ In choosing this definition, I mean only to gesture towards many of the key aspects of ecology however one might define it. I don’t mean to suggest Haeckel’s is an exhaustive definition of the discipline or that Haeckel founded it (or, indeed, even intended to; see Kingsland 2005). As McIntosh has succinctly put it, ecology may well have no singular founder or foundation; it is “more a bush with multiple stems and a diffuse rootstock than a tree with a single, well-defined trunk and roots” (1985, p. 7). Aspects of ecological thought were developed throughout history, perhaps as far back in the Western tradition as Herodotus and at least to Theophrastus, and many key aspects of the science appear in the 18th century with Linnaeus, Buffon, and Humboldt.

⁴ This isn’t to say that *all* knowledge of evolutionary history depends on ecological knowledge. For instance, molecular biology has developed many statistical methods for detecting selection on parts of the genome that don’t require any idea as to the source of that selection or the role function of the relevant gene.

Still, it’s worth noting that this kind of knowledge still wouldn’t give us enough to know the gene’s evolutionary history. We would know that selection had occurred; we still wouldn’t know why it had. To know that, we would need to know why the gene was favored over alternatives in the first place.

has undergone and why they've undergone them, we often need to know, or at least have a rough idea of, the selection pressures they've faced. Ecology, being the study of the circumstances that generate selection pressures, is therefore often not only a key source of knowledge about evolutionary history, but must be epistemically accessible *prior* to that evolutionary knowledge.⁵ We see this in practice: the discovery of evolutionary history very often *begins* with an observation of present interactions and their expected fitness effects (Davies, Krebs, and West 2012).⁶ Notice that this implies that we can study those conditions and the selection pressures organisms face before we know anything about their evolution. We can, and indeed do, ask ecological questions without asking any evolutionary ones.

Let's consider a brief, toy case to illustrate this distinction. Imagine a population of finches with unusually shaped beaks endemic to the island of Fringuello. We might ask one of two questions about their odd beaks: on the evolutionary side of things, we might ask why they evolved the beaks in the first place; on the ecological side, we might ask what interactions with their environment their unusual beaks allow them to take part in that help to explain their present abundance and distribution (i.e. we might ask why there are as many of them as there are and why we find them where we find them). To answer this latter question, we need to observe the finches in their environment. Suppose that, upon observation, we discover that the finches use their beaks to open the seed pods of a plant that is also endemic to Fringuello, and that this is their primary source of food. Given this, their beaks play into our explanations for their present abundance and distribution via the role they play in their ability to gather food on Fringuello. They can, in part, successfully occupy the island because they can access food there thanks to their beak shape. Of course, this isn't a full explanation for the finches' present abundance and range, but that wasn't the question we set out to answer in the first place. Rather, we hoped to know what role the finches' unusual beaks play in *helping* to answer those questions.

⁵ My thanks to Renee Duckworth for directing my attention toward this point about the structure of many evolutionary explanations.

⁶ Now, this isn't a perfect procedure; the problems with adaptationist thinking, whether metaphysical or methodological, are well-documented (Gould and Lewontin 1979; Lloyd 2015). That being said, it is nevertheless often a valuable heuristic and is employed throughout evolutionary biology. Indeed, it is often all that a working biologist has to go on when a trait's relation to any genetic underpinning is unknown (e.g., when studying the evolution of a complex behavior).

To answer the former, evolutionary question – why they have evolved their beaks in the first place – we would need a very different sort of information. Rather than a set of facts concerning the present, we would need to investigate the past: how has the finches’ unusual beak shape historically contributed to the fitness of the finches when compared with alternative beak shapes? Supposing the past environment was similar to the present one, our answer to the ecological question can help us here. Finches with the unusual beak shape outperformed conspecifics with alternative beak shapes because it allowed them to access an abundant food source that was otherwise either inaccessible or less accessible. Notice that one would need to make a similar inference even if one knew both the genes controlling the finches’ beak shape and that selection had favored those genes over alternatives. In the absence of information about how the finches’ beaks contributed to their survival and reproduction, we would know only that they had done so, not the means by which they did.

We see here both the distinctions and the overlap between the sorts of information relevant to answering the ecological and evolutionary questions. Our answers to both sorts of questions rely on claims about the role a trait plays in enabling certain interactions with an organism’s environment, but they split apart with respect to their temporal focuses. This isn’t to say that evolutionary questions concern exclusively history or that ecological ones concern exclusively the present. One might be concerned with past ecological interactions or present evolutionary forces. Instead, I argue that ecological questions typically require information about what we might roughly call *narrow* temporal windows, whereas evolutionary questions typically require information about a *wide* temporal window, which includes many of the narrow ones studied by ecology and their various interrelations. This, of course, requires clarification.

To know that the finches’ beaks enable them to open seed pods and that this is better for them than a beak shape that did not allow them to do so, we need not observe them for very long. It would be enough to see a handful of them use their beaks to open the seed pods and to study the particular mechanics of the opening process to know enough to confidently assert that the beaks help explain their abundance and distribution in virtue of their enabling the finches to gather the food in their environment.⁷ Contrast this with the answer to the evolutionary question.

⁷ To know that the ability to do so would be better than the alternative would be obvious enough that it would not require experimental confirmation. To quote Tinbergen, “the survival value of many attributes, behaviour and structure alike, is so obvious as to make experimental confirmation ludicrous. One need not starve an animal to death to show that its feeding

To know why the finches evolved their beak shape, we need to know not just the beak's contribution to fitness at some point in time, but also its average contribution over *many* points in time. We've seen this point already in our discussion of Haeckel's definition of ecology as the study of the interactions that, when summed together, so to speak, make up "the struggle for existence." Ecology studies those particular processes and mechanisms, whereas evolutionary biology considers their total effects.

It's helpful here to point towards similar remarks on the structure of scientific explanations made by Wimsatt. Think about a piece of granite (Wimsatt 1972). You're interested in investigating its various properties, and so you run a variety of experiments on it to determine its chemical composition and the patterns of thermal conductivity, tensile strength, density, and so on of its various spatial parts. You notice that these different properties strongly correlate with one another; in fact, by and large, the parts you pick out when looking for each of these properties prove to be spatially coincident with one another. Now imagine that you want to investigate a housefly. Seeing as chemical composition proved so explanatory with the granite, you start with that again, but then things get messy. You notice, for one, that chemical composition does not correlate as well with cell types, or organ systems, or functional parts; instead, the chemicals are mixed up throughout. Even worse, you can't get these other ways of breaking the fly down into its parts to neatly line up with one another, either. The fly's parts, when viewed from different sciences, aren't spatially coincident, but instead exhibit many different mappings between one another (Wimsatt 1972).

Following the work of Simon (1962) and Kauffman (1971), Wimsatt (1972; 2007) argues that objects can be organized in more than one way at the same time. House flies are made of cell types, organs and organ systems, functional parts, and a variety of chemicals all at once. To commit myself to a less metaphysically heavy claim, we at least have to pretend they are in order for most of our scientific projects to work. Our sciences assume that we can sensibly talk about these different organizations of objects such that some questions are best addressed using different spatial organizational schemes, and oftentimes interactions between the objects in those schemes. To explain the functioning of a fly's eye, we need to talk about not just chemistry, or cellular biology, or physiology, but all of those things at the same time. One result of this is that

behaviour has survival value, nor need one cut off a Blackbird's bill to show that this organ is necessary for successful feeding" (1963, p. 419).

different sciences sometimes pick out different particular objects even as they talk about the same system. Molecular biologists and behavioral ecologists might both study western honey bees, but they likely pay attention to very different parts of them.

I mean to make the same point here about evolutionary biology and ecology, but with reference to temporal rather than spatial organizational schemes. A population interacting with its environment at different times can involve very different sets of objects. Most obviously, individuals will have been added and removed if we are considering different generations, whether due to birth, death, immigration, or emigration. The environment and the population's interactions with it may have also changed: a season with plentiful rainfall might give way to a period of drought; an invasive species might move into the space; or a migratory species might move from its summer habitat to its winter one. As we expand the temporal window under consideration, the set of objects and their relations to one another change. In general, they become more complicated and diffuse. Moreover, different generalizations emerge (Levin 1992; McGill 2010; Potochnik and McGill 2012). Variation in some trait might be mysterious to us when considered from the perspective of a narrow temporal window but become explicable when considered through wider ones. For instance, we might see a distribution of wider and narrower beaks in a population of finches and find that a narrower beak leads to lower rates of infant mortality (Gibbs and Grant 1987). We wonder, then, what the point of the wider beak is. Only once we've observed the finches for longer will we discover that wider beaks are advantageous during periods of drought, since they allow finches to eat harder, larger seeds that would otherwise be inedible. The reverse can also be true: an interaction with the environment that appears to be very important in the short term might become irrelevant over longer periods of time. A trait that allows a population to survive a rare drought might not be relevant to the overall survivability of the species over long enough periods of time.

These generalizations that emerge at different temporal scales are sometimes more or less relevant to ecology or evolutionary biology. The utility of a trait during a rare drought might be indispensable for explaining the present distribution and abundance of a population but ultimately will be no more than a blip in future evolutionary explanations. The utility of a trait under a past set of environmental conditions - which explains its evolution and spread in the population - might be irrelevant to understanding why a population can successfully inhabit some environment in the present if those conditions have changed. The generalizations we can

make when considering narrow temporal windows will generally be more relevant to ecology, while those we can make using wider temporal windows will generally be more relevant to evolutionary biology.

None of this is to say that the narrow temporal windows studied by ecology must be short, or that the wide ones studied by evolutionary biology must be long. We cannot know the role a maternal effect plays in a population's interactions with its environment unless we observe a number of generations. Neither do we have to look across countless generations to see the effects of selective pressures on a population; we can sometimes see it across just a couple, as with Kettlewell's moths or the Grants' finches of Daphne Major. These temporal windows are narrow and wide with respect to one another, not relative to any determinate value.

I do not mean to make a radical claim here, to suggest that there is some fundamental break between ecology and evolutionary biology. They are clearly deeply interrelated and often inseparable. Ecologists are directed by evolutionary concerns, and evolutionary biologists rely on ecology in their study of evolutionary history. I only mean to point towards a way in which they are also different from one another. Where evolutionary biology concerns itself with the effects of past interactions with the environment, ecology concerns the interactions themselves. This leads us to a further distinction internal to ecology that will prove fundamental to much of what I argue later: the interactions studied in ecology can be studied not only as effects, but also as causes or possible causes of future effects. That is, when we know that an attribute is present in a population, we can leverage that fact to make predictions about future interactions and their relation to the future abundance and distribution of a population. Ecology can consider not just the actual environment in a narrow temporal window, but also other possible environments in other such windows.

To illustrate this, let's consider a modified version of the previous toy case.⁸ Suppose that, instead of a beak shape, it's a complex behavior that enables the finches of Fringuello to open the seed pods. They hold the seed pods with one foot to gain leverage and peel them open by grasping them with their beaks. Suppose further that this behavior is *learned*: children learn to open seed pods by observing their parents do so and imitating them rather than by virtue of an innate program. Now imagine that humans have begun importing packaged foods onto the island

⁸ This modified case is loosely inspired by British blue tits learning to open milk bottles to drink the cream off the top of them. See Fisher and Hinde (1949).

and are quite messy about it; they leave whole, unopened granola bars everywhere they go. We want to know whether the finches will be able to eat those granola bars and what effect eating them might have on the population.

Given their dexterity in opening seed pods and that the behavior is learned, it seems likely both that the finches will be able to open the granola bars and that, assuming they like what they find inside, the behavior will spread through the population. We can therefore predict that the finches will likely open and eat the granola bars. We can then use other information about the finches and their environment to predict how this new food source will affect their population's range and abundance. If the granola bars are wholesome for the finches, then the new abundance of food may well cause an increase in population if they have previously been food-limited. The primary selective pressures on the finches might also change. Rather than competition for food, it could be that competition for nesting sites would become the dominant selective force as the population increased. Whatever the actual effect might turn out to be given the particulars of the case, we can see the general shape of the sort of predictions we could make given what we know about the finches in the present.

I want to draw out two points from this example. First, to repeat myself a bit, notice that this analysis can proceed without ever appealing to evolutionary history. That history is often informative, but it isn't necessary to do the work described. That opening the granola bars is "accidental" rather than evolved makes little difference with respect to its effects. Second, and more importantly for us now, we can see here a different sort of question being asked about the finches' interactions with their environment than was posed in the first version of the case. Before, the question concerned how one of their traits influenced their present abundance and distribution. In this modified case, it's instead a question of how a trait could allow them to interact with some novel feature in their environment, and how this novel interaction could affect their future distribution and abundance. In the first case, we were interested in the traits as causes of present effects. In the second, we were interested in traits as present causes of future effects.

A similar point was made by Tinbergen, though in a different context:

Our study always starts from an observable aspect of a life process – in the present case, behaviour. The study of causation is the study of preceding events which can be shown to contribute to the occurrence of the behaviour. In this study of cause-effect relationships the observable is the effect and the causes are sought.

But life processes also have effects, and the student of survival value tries to find out whether any effect of the observed process contributes to survival... It is clear that he too studies cause-effect relationships, but in his study the observable is the cause and he tries to trace effects. Both types of worker are therefore investigating cause-effect relationships, and the only difference is that the physiologist looks back in time, whereas the student of survival value, so-to-speak, looks "forward in time"; he follows events after the observable process has occurred (Tinbergen 1963, p. 418).

We do not need to limit ourselves to studying traits as outcomes. They can also be used as inputs. We can even go further than Tinbergen and I have so far suggested. Besides predicting future effects given the actual present, we can also predict how *possible* alterations to an environment could lead to *possible* futures. We can posit an environmental change and use our knowledge about the traits in a population to predict how that change would affect the population's distribution and abundance. Ecology can be not just forward-looking, but modally forward-looking, considering a number of possible futures. This is, in fact, the work being done by conservation biologists using trait-based models to predict populations' responses to climate change. Given a set of environmental variables, their likely future values, and the present traits in the population, conservation biologists estimate how populations' ranges will shift over time.

I bring this up as a distinction internal to ecology, but it is also a further distinction between the explanatory structure of some parts of evolutionary biology and ecology. Traits and their effects take on opposite explanatory roles between evolutionary explanations, in which they are explananda, and these forward-looking ecological projects, in which they are explanantia. I will cover the latter in great detail in Chapter 1, so I won't belabor the point here. I will, though, offer a preview of things to come. As I said above, teleosemanticists have focused on backward-looking notions of function. This, I will argue, has resulted in a focus on only one explanatory role played by traits and their functions in biology, which contributes to their difficulties handling cases of novelty. Since I have already discussed Swampman in some detail, I'll turn here to indeterminacy of content and Fodor's frogs. Some of the pressure involved in that case seems to come from our desire to explain not just why the frogs have the traits they do, but also how their traits either help or hinder them in various environments. That is, besides wanting to know their evolutionary history, we also want to talk about their visual systems

succeeding or failing to represent something in their environment as good for them to eat *now* and how this affects their success or failure in that environment. Given the focus on traits as explananda in evolutionary biology, it should be little surprise to us that succeeding or failing to fulfill a selected-effects function can be a poor indicator of their *present* success or failure in an evolutionarily novel environment.

3. Conclusion

I've argued here for a distinction in the subjects studied by evolutionary biology and ecology. While evolutionary biology is concerned with explaining the presence and abundance of traits in a population, ecology means to explain the presence and abundance of organisms in environments. These disciplines are not mutually exclusive. Ecologists are often interested in evolutionary questions, and evolutionary explanations rely on claims about an organism's historical ecology. Having said that, there is an interesting difference between the objects they study: evolutionary biology is typically concerned with wide temporal windows, while ecology is concerned with narrow ones, because each discipline is typically more interested in the generalizations that emerge under the windows it studies. A further distinction is internal to ecology. We are sometimes interested in explaining present or past abundance and distribution, but we are also sometimes interested in how present or possible future changes to the environment will change a population's abundance and distribution. This in turn brings into better focus the earlier distinction between the uses of traits and their effects by evolutionary biology and ecology, respectively: the former takes traits to be things one ought to explain, whereas the latter takes traits and their effects as things that explain or predict further facts.

The rest of the dissertation depends on these distinctions. It is because ecology has a unique object of study that ecologists must sometimes depend on a notion of function not yet discussed by philosophers of biology. It is also because teleosemanticists have focused on the role of representation in historical explanations that they have failed to give satisfactory answers to their opponents. With the alternative, ecological explanation in hand, we'll be able to offer more satisfying responses to those critics in chapters 2 and 3.

Chapter 1: Realized Functions in Ecology

1. Introduction

Thus far, I have argued for a pair of distinctions. The first concerns one of the ways in which evolutionary biology and ecology differ from one another. While evolutionary biology is typically concerned with those objects and organizations of objects that appear in relatively wide temporal windows, ecology is typically concerned with those objects and organizations of objects that appear in relatively narrow ones. The second concerns the role of traits in our explanations. They might, on the one hand, be treated as *effects* to be explained by an appeal to past causes, as when we say that a zebra has stripes *because* stripes deter flies. On the other hand, we might treat traits as *causes*, as when we say that a particular hummingbird is able to pollinate a particular orchid *because* it has certain anatomical features. With these distinctions in place, I can focus on the subject of this chapter: the use of traits in ecological explanations as explanantia rather than as explananda in some ecological niche models. In this work, scientists estimate the environmental conditions a population could successfully occupy and then use this information to explain or predict their historical, present, or future range. To return to another variant of our toy example, consider the finches endemic to the islands of Pinzón and Fringuello. In an effort to determine whether the finches will survive in a warmer world, we might investigate their climatic niche (i.e. the set of climatic conditions they can successfully inhabit) and then compare the estimated future climates of Pinzón and Fringuello. Suppose we found that, while Pinzón will likely become uninhabitable with respect to climate, Fringuello is likely to remain within the finches' climatic niche. Given that information, we could predict that the finches' range will shrink as the world warms, eventually going locally extinct on Pinzón, but the species will still survive the changes, since they will be able to carry on on Fringuello.

Though there are approximately 10,000 scientific papers that in some way use or discuss ecological niche models, philosophy has paid them little attention. This is especially a shame given their utility as case studies for the corner of ecological work I'm exploring here. In this chapter, I'll consider a real-world case in order to identify the sort of explanation we get when using traits as explanantia rather than as explananda: the predicted habitat loss and gain for 195 species of sub-Saharan amphibians due to climate change (Garcia et al. 2014). The model uses organisms' traits as inputs in order to help explain or predict why they can or will occupy some environments but not others.

I'll argue that, in this case and many others like it, ecological niche modeling is well understood as implicitly using traits as parts of *functional explanations*. They use what an organism's traits allow it to do – their traits' functions – in order to explain and sometimes predict those organisms' realized niches, or the set of biotic and abiotic conditions that a population interacts with in the environments it actually successfully occupies. When this is the case, these models can be fruitfully understood as using a notion of function not yet discussed in the extant literature, *realized function*, which we can assign in reference to the ways in which an organism's traits help to explain why some part of its range matches onto its population's realized niche. I argue for this in four steps. First, I give an extended treatment of the case study from Garcia et al. (2014). In the process, I show that this model and others like it are well-understood as using functions to predict species' future ranges. Second, I consider what sorts of features the notion of function at play needs to have to do the work so described. I introduce realized functions and show how they can fulfill these explanatory roles. Third, I consider some potential objections to the notion of a realized function. Finally, I conclude with some brief remarks on how realized functions open avenues for novel research.

2. Ecological Niche Modeling in Conservation Biology

Before diving straight into Garcia et al., it's worth taking a moment to discuss the sort of research they are engaged in more generally. Conservation biologists are attempting to do what they can to help a world facing climate crisis. Unique ecosystems and those already threatened will be at high risk in the very near term (IPCC 2022). In fact, the threat is already here. We have already seen many local extinctions linked to climate change (Wiens 2016), and 2023 was the warmest year on record. Barring significant cuts to our current emissions, we can expect approximately 20% of species in tropical and coastal environments to face climatic conditions unlike any they have dealt with in their evolutionary histories (IPCC 2022). They will not have faced selection to survive under these new conditions; they will not have evolved to fit into the world they'll be stuck in. Moreover, it seems likely that species will not be able to evolve in time to fit into the expected novel conditions. While climatic niche shifts appear to be faster than previously estimated, they can nevertheless be more than 200,000 times slower than the rates of change in the relevant climatic variables (Jezkova and Wiens 2016). That doesn't tell us, though, that species are doomed to extinction. Populations do not sit idly by as the world burns around them. Past climate change has often been associated with changes in the range of populations

(Williams and Blois 2018; Wing et al. 2005). Rather than changing to fit the new environment, populations move to an environment that matches with the conditions they have already evolved to fit (Parmesan and Yohe 2003). Species may also find themselves not merely surviving in similar conditions, but thriving in new ones. Introduced species can adapt to climate change as much as a million times faster than native species and ten times faster than climate change (Wiens, Litvinenko, Harris, and Jezkova 2019). This, of course, can lead to other problems. Introduced species can become invasive, destroying the already vulnerable ecosystems they enter.

This situation presents unique challenges for working conservation biologists attempting to predict how a population will respond to climate change. Because so many species are expected to encounter evolutionarily novel environmental conditions, biologists cannot rely as much on history as a guide to the species' future success or failure.⁹ Given also the fact that it seems unlikely most species will be able to adapt quickly enough to climate change – except for in those cases in which their rapid adaptation could be detrimental to the species they bump up against, as with invasive species – conservation biologists must often focus on what a population is capable of doing *now* relative to some set of possible environments in order to predict how they will fare going into the future. Finally, and to make matters even worse, these scientists must make these predictions under conditions of high uncertainty. We do not know precisely what the world will be like as the climate changes, how humans might mitigate or worsen those changes for nonhuman organisms, or how nonhuman organisms could tolerate, adapt to, or plastically respond to them. Conservation biologists have developed many tools and practices for handling these challenges. Here I'm focusing on their use of ecological niche models. As I said above, these models use estimates of a population's niche¹⁰ – the set of environmental conditions they can successfully occupy – in order to explain and predict its range (i.e. the environments it actually inhabits). These have often relied on bioclimatic envelope models, which use the relationship between climate and species occurrence in order to estimate the environments a population can successfully inhabit. These models, however, contribute high levels of

⁹ By this I do not at all mean that history does not or cannot play a role in this work; it is often still incredibly useful. I am here only asserting that we cannot rely on it in the same manner we do in other contexts with fewer novel environmental changes.

¹⁰ Precisely which niche concept is at play can differ from model to model, so I speak here only generally.

uncertainty to predictions about species' future ranges for many of the reasons mentioned above (Garcia et al. 2012). One promising method for decreasing this uncertainty is to couple bioclimatic envelope projections with information about the roles traits play in a population's interactions with its environment. Introducing mechanisms in this way helps us to better align our predictions with biological realities.

Garcia et al. (2014) do just this with their model of the likely range shifts due to climate change of 195 sub-Saharan amphibian species. Their model has two principal parts: (1) a set of climatic envelope models; and (2) trait-based climate vulnerability estimates. The former they constructed using presence-absence data: if a species was present, the climatic conditions (mean temperatures in warmest and coldest months, as well as annual precipitation) in that location were considered part of their viable climatic envelope; if they were absent, those conditions were omitted. They predicted future climatic conditions given a particular emissions scenario and used that information to in turn predict changes in the occurrence of their acceptable climatic envelopes in the world using consensus projections from seven different methodologies. These changes were of three kinds, predicted individually and conglomerated: the loss of currently suitable environments; the fragmentation of currently suitable environments; and the gain of new climatically suitable environments. They also calculated the distance from their presently suitable environment to future suitable environments, assuming unlimited dispersal ability. This all resulted in a map of environments that could be habitable for the species going into the future. As we've already noted, though, this estimate on its own is very likely vulnerable to many sources of uncertainty.

It's here that the second part of their model comes in. Their trait-based vulnerability estimates were taken from the IUCN's trait-based vulnerability assessment for amphibians (Foden et al. 2008; Foden et al. 2013). Garcia et al. suggest that many of these are not, strictly speaking, traits (Violle et al. 2007), but are instead ecological characteristics (i.e. those features that determine their relationships to their environment). I'll return to this point later. For the moment, though, let's consider the relevant traits and how they might affect the future range of each species.

Following on others' work (Williams et al. 2008; Chevin et al. 2010), Garcia et al. discuss three primary risk factors for a species' vulnerability to climate change. First, there's the phenotypic and behavioral plasticity of individuals within a population under changing climatic

conditions, especially the mean temperatures in the warmest and coldest months and annual precipitation. A population composed of individuals with more plastic physiology or behavior will likely better withstand such changes to their home environment, making them less likely to lose their habitable environment than a less plastic one and more likely to gain new habitat. Second, there's a population's actual dispersal ability. A population more capable of dispersing could more easily colonize a newly habitable environments or could better cross between parts of its fragmented habitat. A population with very little ability to disperse would do comparatively poorly under those same circumstances. Third and finally, there are life-history traits. Populations that reproduce early in a season, have high fecundity, or who reach maturity quickly will be more likely to survive climatic changes than those who with the inverse traits.

In order to generate a map of expected gains and losses, Garcia et al. compare their climatic envelope model to the species' traits identified above. Plasticity was used to modulate expected habitat loss, fragmentation, and gain, and dispersal and life history traits were used to help estimate habitat gain. Their results could be quite useful in informing policy. For instance, expected gains and losses often correspond well with geographic regions (e.g., montane regions had higher fragmentation) or present climatic features (e.g., arid regions had species especially vulnerable to changes in precipitation). There may therefore be high-level descriptions of environments that we can use as heuristics for guiding broad conservation policies in some regions.

These specific findings are, however, less important for this discussion than the authors' comments on methodology, which are well-worth quoting at length:

The case of sub-Saharan African amphibians illustrates the framework's application when traits in the strict sense (*sensu* Violle et al., 2007) are largely unavailable, a situation that is common for many taxonomic groups (e.g. Gonzalez-Suarez et al., 2012). Some of the data used were derived from the characterization of known distributions of species (Foden et al., 2008, 2013) as proxies for traits. One example is tolerance to climatic variation, inferred with statistical approaches relating species ranges to climate variables. Whereas previous studies (e.g. Thuiller et al., 2005a; Feeley et al., 2012) used similar approaches, such proxies do not strictly summarize traits but the interaction between traits and the environment. Climatic tolerance inferred with these

approaches may represent under-estimates when climatic niches realized in the present are truncated (Feeley & Silman, 2010). The finding that upper thermal limits tend to be highly conserved while lower limits are highly variable across organisms (Araujo et al., 2013) further indicates that such proxies may be misleading. Only the physiological limits of species could indicate their full capacity to adapt to climatic changes through plastic adaptation...Where available, more precise and reliable estimates of response-mediating traits allow for a closer coupling of bioclimatic envelope models with traits, leading to projections that are more appropriate for conservation planning (p. 731).

There are three points I mean to draw out of this. The first returns us to the point I set aside above concerning whether the “traits” picked out in the model are well understood as such. Many of these, like thermal tolerance, are estimated from present-day presence-absence data and are often taken to *stand in for or summarize* underlying traits. However, Garcia et al. claim that it’s not the traits themselves being summarized, but instead *the interaction between those traits and the environment*. This is because of their interpretation of the particular definition of “trait” they employ, according to which “a trait is a feature measurable at the level of the individual, which does not require additional information from the environment or any other organizational level” (Violle et al. 2007, p. 890). I don’t mean to come down on this one way or another, though. Whether the inputs are genuine traits or not, the point remains the same: the input into the model is a claim about how traits – whether directly or indirectly – enable the population of interest to fit into certain environmental conditions. The inputs in the model are those interactions between an organism and its environment that help to explain why it can successfully inhabit some environments but not others.

This brings me to the second point. These sorts of stand-ins for more mechanistic explanations for these organism-environment interactions come with drawbacks. Reliance on current presence-absence data may cause us to underestimate a population’s actual future range, since populations in the present day may not occupy all of the environmental conditions they could successfully occupy for any number of reasons. As such, when it’s possible to get it, data closer to the mechanisms themselves or that is more directly supported by experimental evidence can give us a more accurate picture of what a population is in fact capable of. Instead of a general sense that a population appears to require a certain level of annual precipitation, a model

might instead break down that need into things like, for instance, precipitation as a stimulus that triggers reproductive activity (e.g., Chesson et al. 2004). More detailed inputs do not eliminate or dissolve the higher-level phenomenon, though. More detailed inputs just provide us with a better understanding of how it works (cf. Wimsatt 2007). A population doesn't stop needing a certain level of annual precipitation just because we know why it needs that amount.

The third and final point is closely related to but importantly distinct from these others, having to do with what *kind* of data would improve these models. Garcia et al. emphasize that what really matters for getting an accurate prediction isn't, strictly speaking, what we can infer about a population by considering it only as we find it in the world today. After all, looking only at present data can be importantly misleading when a population could, if given the opportunity, successfully occupy a wider range of environmental conditions than we see them occupy today. What we need instead are "more precise and reliable estimates of response-mediating traits," which in turn "allow for a closer coupling of bioclimatic envelope models with traits" (Garcia et al. 2014, p. 731). We need to know what a population is currently capable of not just under its present environmental conditions, but under other possible environmental conditions as well. With that information, we could better predict the sorts of responses the population could mount in the face of climate change. If we knew what sorts of conditions a population's traits could help it to successfully occupy, we could in turn better inform our predictions about the population's future range.

Now, with all of this in place, we can get to the implicit invocation of functions. To summarize, Garcia et al.'s model and others like it use information about the actual and possible interactions between a population's members and certain environmental features that contribute to the population's ability to successfully occupy some environment in order to predict how the population will fare in the face of climate change. Even in a more mechanistic model that further explains these interactions, it is the interactions themselves that explain why a population can or cannot occupy some environment, thereby explaining the continued existence or disappearance of said population (cf. Wimsatt 2007). To improve the quality of our predictions made with such a model, it's important that we consider not just what we presently observe about a population's interactions with its environment, but also what sorts of environmental conditions a population *could* successfully occupy.

Here we ought to pause for a moment to consider what it means for a population to “successfully occupy” some environment. Here’s one gloss: it’s for the population to be able, *ceteris paribus*, to persist there indefinitely. An explanation for this persistence relies on an appeal to the fitness of members of the population in that environment. It’s because members of the population can survive and reproduce that the population persists. The fitness of the population is in turn explained by the interactions their traits enable them to have with their environment. It’s these sorts of interactions that Garcia et al. model, though at a very low resolution and, as they themselves say, the model would be better served by a more mechanistic understanding.

This description fits surprisingly well with standard philosophical analyses of functional explanation. It’s historically been said that the structure of a functional explanation should be this: we appeal to a function in order to explain or help to explain the existence of some trait on the basis of what it does (Hempel 1965). On past readings of this claim, it hasn’t been obvious how what a trait could or will do can help us to answer these questions unless we consider the history of selection on that trait (Garson 2019). But think of Tinbergen again:

The study of causation is the study of preceding events which can be shown to contribute to the occurrence of the behaviour. In this study of cause-effect relationships the observable is the effect and the causes are sought. But life processes also have effects, and the student of survival value tries to find out whether any effect of the observed process contributes to survival...It is clear that he too studies cause-effect relationships, but in his study the observable is the cause and he tries to trace effects. Both types of worker are therefore investigating cause-effect relationships, and the only difference is that the physiologist looks back in time, whereas the student of survival value, so-to-speak, looks "forward in time"; he follows events after the observable process has occurred (1963, p. 418).

We need not explain only a trait’s present existence via a claim about its function; *we might also explain its future existence*. Consider the following general explanatory structure: some trait will exist in a population in the future because it enables members of the present population to interact with certain features of its environment in ways that enable it to survive and reproduce. We can study a trait’s “survival value” by looking forward as well as backward. Returning to the language of successful occupation, a trait’s function can explain why a trait will exist in the

future by explaining why a population can successfully occupy its environment. This is just as much an explanation for a trait's existence as past selection, which takes as its explanans the past effects of a trait. The only difference is where – or, I suppose, *when* – on the timeline the trait of interest is. To borrow a little more from Tinbergen, that we distinguish between these sorts of explanations at all is “an accident of human perception” (1963, p. 418), or perhaps of institutional practices: we perceive biological systems in such a way that we tend to think of their present circumstances as the effects of past causes rather than as present causes of future effects.

It's this sort of function that Garcia et al.'s model is using as an input. The traits it identifies are better understood as either the *functions* of traits or as *traits defined by their function* because they help us to explain and predict why the amphibian species under study will be able to successfully occupy some environments and not others in just the way described above. They give us information about whether the populations will persist by considering whether they will enable those populations to survive and reproduce – to be fit – in likely future environments. Take, for instance, thermal tolerance, which is the result of a suite of physiological traits and is commonly appealed to in explanations for a population's present range. Thermal tolerance is selected as an input in the model because of its present function: it helps to explain the population's range by explaining how it can successfully occupy those environments. The model's output in turn predicts whether the populations' thermal tolerances will help to explain their successful occupation of some environment or not in the future when considered along with other factors, including other possible functions. This is the same sort of explanation we get out of backward-looking functional explanations, but temporally shifted.

But what sort of biological function are we talking about here, anyway? It's time we turn to the general features such a notion of biological function requires and a notion that has them.

3. Realized Functions

There are three key features the notion of function we're after will need to have. First, it must help us both to explain why a population can persist in an environment and *to predict whether* a population can persist in some environment. It won't be good enough if it helps us only to understand the present success of a population. That might work in ecological niche models of a population's present range, but it won't do for predicting possible future success. Second, it must be applicable by considering only a population and its contemporaneous environment, without an appeal to history. This isn't at all to say that history cannot help guide

us in the application of the notion. It's just that, under anthropogenic climate change, some of the environmental conditions we'll need to consider when looking for traits' functions will have no analogue in a population's evolutionary history. We cannot look to history to tell us how a population will respond to a temperature that it has never encountered; we can only consider what its response *would* be. Again, and as discussed by Garcia et al. (2014), we need to know about what an organism is theoretically capable of, not just what it has historically been capable of. This leads us quickly to the third feature. The notion of function must be *modally* applicable. We must be able to say not just what about an organism helps it to successfully occupy its *present* environment, but also what *might* help it to occupy a number of *possible future environments*. We do not know which emissions pathway we'll follow, which environments a population will actually reach, or how exactly those environments will change over time. There are typically a number of possible futures conservation biologists must take into account in order to offer recommendations to policy makers. As such, their models must be able to take in as inputs not just the actual functions of traits, but the possible functions of traits as well. Putting this in more concrete terms, the application of the notion of function must be sensitive to changes in the environment under consideration, since different changes will result in different explanations for why the population can successfully occupy said environment – or for why it will fail to do so, if it cannot!

Our needed notion of function must, then, help us to explain why a population can successfully occupy some environment, and it must be able to do so by considering traits in a quasi-ahistorical way. Let me take this in pieces. First, it will be helpful to further specify *which* environments we ought to consider. We don't want to become overwhelmed imagining an indefinite set of possible ones. At the same time, we don't want to restrict ourselves only to some actual environment or environments. Doing so would hamper our ability to generalize, which is what we need out of these functions in the first place. It's for this reason that I recommend we take a step back from particular environments to instead focus on, as the name for ecological niche models suggests, *the niche*. This has a few major advantages over focusing on the environment itself. First, it provides us with a finer-grained way of describing environments than is possible in merely spatial terms. A population might change its niche without changing its physical range if it begins interacting with conditions in the environment that are either novel or that it previously ignored. Second, it allows us to focus on the environmental conditions

themselves, thereby enabling generalization between environments sharing those conditions. We can look for the features of a population's actual or possible realized niche out in the world in order to predict whether the population will successfully occupy some environment in the future. Third, it focuses us on those environmental conditions that actually matter for determining where a population can successfully live. We don't need to consider every aspect of an environment to predict whether a population can successfully inhabit it, since a population won't meaningfully interact with every aspect of its environment.

Still, the niche concept is itself quite unruly, and variants of it can themselves be too broad (e.g. the fundamental niche) or too narrow (e.g., the climatic niche) for our purposes. It's for this reason that I suggest picking up *the Hutchinsonian realized niche* (hereafter "realized niche"), or the set of biotic and abiotic conditions that a population interacts with in the environments it actually successfully occupies (Hutchinson 1957; Holt 2009; Carscadden et al. 2020). Doing so focuses us on just those environments a population encounters while, at the same time, allowing us a great deal of flexibility in which environmental conditions we might use in the model, depending on which we believe would provide us with the best predictions.¹¹

With the realized niche in hand, I can offer a refined version of what we need out of these functions in the first place. They ought to help us explain why and predict whether a population has a realized niche by an appeal to what an organism's traits allow it to do. This information would then allow us to predict the population's future range. What, though, about the quasi-ahistorical aspect of the notion of function? The appeal to the realized niche can help us here, too. If an organism's traits help us to explain their present realized niche, we don't need to know how the organism ended up with those traits in the first place. We just need to know what its traits allow it to do in the environment. Now, because it ignores history, this means that this notion of function will always provide, at best, an incomplete explanation for a population's realized niche at any given moment. The full explanation would include a story about how the population got there in the first place. However, in these circumstances, the incomplete explanation is all we need. The functions I'm discussing are only one input among many in a model attempting to predict a population's future range by approximating some aspects of reality. To expect more of them in this context would be inappropriate.

¹¹ Which environmental conditions best predict a population's presence or absence turns out to often depend on spatial scale (McGill 2010; Potochnik and McGill 2022).

We need, then, a notion of function that uses how an organism's traits allow it to interact with its environment in order to explain a population's realized niche. I propose the following:

Realized function – Some trait x has the realized function f just in case (1) x allows an organism to f and (2) the ability to f helps to explain why some environment matches onto part of the organism's population's realized niche.

To see how realized functions can fulfill the explanatory and predictive role I've described, let's see how they can be applied to our understanding of Garcia et al.'s model. Let's take, in particular, thermal tolerance once again. An organism's thermal tolerance is the result of a suite of traits. It is one of an organism's most important features when predicting and explaining whether it can successfully occupy some environment or not, especially at the spatial scale that Garcia et al. are using (McGill 2010). Thermal tolerance, then, is a realized function of its underlying traits. An organism can survive in a temperature range, and their ability to do so helps us to explain why some environments match onto the organism's population's realized niche, hence helping us to explain why those environments are part of its realized niche. There are three points about this well worth drawing out in greater detail. First, because of thermal tolerance's importance in explaining a population's realized niche, it's projectible onto other possible environments. We can look around to see if an organism's thermal tolerance could play a role in explaining why it could successfully occupy some possible future environment. Realized functions in general will have this feature, since they first and foremost connect traits to *environmental conditions* which can, of course, reappear across space and time. Realized functions therefore fulfill the first requirement: they both explain and *predict* the persistence of an organism. Second, we can identify thermal tolerance as a realized function *even though we lack a mechanistic connection between it and its underlying physiological and behavioral traits*. We know there must be some connection, but we can study thermal tolerance experimentally without having a determinate answer as to its cause. We can therefore get information about the realized function by considering the effects of the traits without having nailed down precisely which traits possess the realized function. At the same time, and in keeping with Garcia et al.'s own comments on the model, we might still improve our predictive power by searching for more mechanistic explanations. A population might, for instance, be more or less sensitive to temperature changes at different points in its lifecycle – thermal tolerance often varies between gestation, childhood, and adulthood – and this information might provide us with the ability to

more precisely predict how a population will be impacted by a change in the average temperature of its environment. Third, whether an organism's thermal tolerance is a realized function is determined by whether it plays into an explanation for why an environment matches onto its population's realized niche. When we take it as an input, we're not just saying that it's a realized function in the present (though that's part of why we select in the first place). We are also trying to find out if it will be a realized function in the future. We're looking around to see if it will play a role in our future explanations, and then we're using that information to predict the population's future range. This is all to say that we do not have to only talk about or make use of present realized functions; we can also use *possible* realized functions as inputs in the model. That an organism does not, at present, inhabit the full breadth of the temperature range it *could* inhabit does not prevent us from using that full range as the input in the model. Given all of this, realized functions have both the second and third features I identified above. They can be identified without looking at evolutionary history, and they can be *modally* ascribed as well.

It's for just these sorts of reasons that Garcia et al. selected thermal tolerance as one of the inputs in their model: "we selected characteristics that are likely to summarize response-mediating traits under climate change" (2014, p. 727). They aim to use high-yield inputs that they can actually gather information about and that will match onto the predictions they can realistically make about future environmental conditions. The same analysis can be used to understand the other inputs they select: each of them is a realized function with these sorts of characteristics. They play key roles in explaining and predicting why a population is or will be able to successfully occupy some environment *because* the functions help to explain why or predict that an environment matches onto the population's realized niche. This in turn explains or predicts the traits' future existence by an appeal to what they allow an organism to do, thereby fitting into the standard description of functional explanation discussed above.

Still, someone familiar with the literature on biological function may well have some alarm bells going off. Realized functions differ in significant ways from standard notions, and I'd like to address two of the worries those differences might raise now.

4. Objections and Replies

There are two possible objections that I want to tackle head-on here. The first concerns just how *telic* these "functions" really are. Sure, realized functions, as described, help to explain why an environment matches onto an organism's realized niche, but is that something their traits

are *supposed* to do? Perhaps, and similarly to how the objection is leveled against causal role functions, we can talk about the mechanism without ever invoking a sense of telicity or purpose. All the conservation biologist needs is the ability to predict whether a population will be able to fit into an environment or not. They don't care about what the population is *supposed* to do, only what it *will* do. Call this the "no-real-telicity" objection. Second, because the definition of realized function I've offered refers only to a "population," loosely defined, an organism's traits might have and lack a function simultaneously depending on which population we're considering them a member of. Realized functions are relative, perhaps even nonnatural, rather than objective features of the world. They're dependent on the interests of the researcher. Call this the "relativity objection." In what follows, I'll more clearly state and respond to these two objections.

4.1. The No-Real-Telicity Objection

Any given robust notion of function ought not just to assign a positive function, but ought also to be able to differentiate between instances of functioning well and malfunctioning. How something can count as malfunctioning, though, is quite tricky for a naturalistic account. As Garson states the general problem,

How is it possible for a trait to dysfunction? A trait is dysfunctional when it cannot perform its function – everyone knows that, but there's a deeper puzzle. How can a trait possess a function it cannot perform? By virtue of what does a function linger, like a ghost, even when the corresponding ability has vanished?
(2019, p. 28)

Any adequate notion of function will be able to explain this seemingly mysterious feature of the world, that a trait can fail to do something it could not in principle have done.

The no-real-telicity objection comes in two parts. First, it challenges the claim that the input into ecological niche models that I've been discussing needs to have any telicity in the first place. Ecological niche models and the conservation biologists who use them can get by just fine without ever invoking an instance of a trait malfunctioning, and so they do not need to differentiate between successes and failures. Second, even supposing that conservation biologists need telicity, realized functions don't have it. The notion is unable to differentiate between cases of functioning well and malfunctioning because it appeals to the wrong events to ground the functional attribution. Present and future events can't do this; only past ones can.

Let's dive into the first part of the objection. Suppose (most) of what I've written so far is right: ecological niche models rely on inputs that relate traits to environmental conditions via the interactions they enable, and these relations must (1) help to explain and predict a population's persistence through time in varying environments, (2) be accessible without appealing to a history, and (3) must be applicable not just in actual cases but in possible ones as well (i.e. must be modally applicable). Why add to this list that the input is a *function*? It's far from obvious that an input with the listed features must include any sense of *telicity*. Indeed, one might reasonably think that we could get away with discussing only the survival value of a trait without considering what it was *supposed* to do. If we can get at a trait's present or possible contribution to fitness relative to some environmental condition, which, as I've argued, could be totally unrelated to what some trait has evolved to be able to do, *then we've gotten all that the conservation biologist really needs to make their predictions about a population's persistence*. There's no reason to invoke a trait's function in the first place.

It's important here to look back at Garcia et al. (2014) and their discussion of what really goes into their model. "From the data available, we selected characteristics that are likely to summarize response-mediating traits under climate change" (p. 727). They define these response-mediating traits by a reference to Luck et al., who in turn describe them as "characteristics of organisms that may result in changes to, for example, their population size with environmental change" (2012, p. 1066). As stated, this *sounds* like a neutral description of these characteristics. They might either help or hinder an organism or population, and it doesn't matter which. It's here that we need to say a bit more about the notion of a trait in ecology, and especially as the term is used by Garcia et al. Remember that these characteristics are themselves explained by an appeal to the mechanisms underlying them, and the better our understanding of those underlying mechanisms – traits in "the stricter sense" – the better our model will be. Traits, as understood here, are hierarchically arranged, with lower-level morpho-physiological traits "adding up" to traits at higher levels. These lower-level traits are likely best understood as neutral, being mere descriptions of the organism independent of the interactions with the environment they might enable. Not so, however, for the higher-level traits. Below is Figure 3 from Violle et al. (2007), which offers us an example from plant ecology. It describes how morpho-physio-phenological (M-P-P) traits result in performance traits, or those that are a

(more) direct measure of an organism's fitness, which themselves result in an overall performance in an environment and so individual fitness.




FIGURE 3 from Violle et al. (2007). It shows how lower-level traits collectively contribute to an organism's ability to perform well in an environment, thereby contributing to their overall fitness.

The figure is a modification of a framework for understanding animal traits from Arnold (1983).

There are three things to notice here. First, though this is not explicitly stated in the diagram itself, the relation between the M-P-P traits and the performance traits is not constant, but varies with environment. Performance is always relative to some set of environmental conditions; it cannot be considered in isolation from them, unlike the lower-level traits. We will need a different diagram for each (importantly) different environment. Second, these traits are not organized or discussed relative to what they do *badly*, but by what they do *well* relative to the environment of interest. The same is true for response-mediating traits. Being smaller measures of an organism's performance, they "add up" to help tell us about an organism's overall performance – "plant performance" in the diagram – which we can evaluate as either positive or negative by considering its impact on individual fitness. It's for this reason that scientists can say whether an organism succeeds or fails to respond well to environmental changes: they can discuss changes to their fitness relative to different environmental conditions, with some of those changes being positive ones and some of them being negative. Third and finally, these individual-level assessments can be conglomerated in turn, resulting in an assessment of the population overall. That is, many diagrams like the one above could be fit together to feed into an assessment of the average fitness of the population as a whole. This, after all, is the aim of

Garcia et al.'s model in the first place: to use the traits of individuals in a population as predictors of population persistence.

Here is how the telicity gets into Garcia et al.'s model, and, indeed, is indispensable to it. The inputs are functional inputs because of their place in the hierarchy of traits described here. While the results of traits that can be neutrally described, the inputs themselves indirectly impact the fitness of organisms relative to different environments. Indeed, they're selected *because* of the contributions they could make if the organism is in the right sorts of environmental conditions, thereby also setting limits on the conditions the organism can successfully occupy. Think again of thermal tolerance. While it is itself the result of a wide variety of physiological and behavioral traits that can be neutrally described, an individual's thermal tolerance tells us about their performance, which can be successful or unsuccessful relative to the environmental conditions of interest. If an organism has the right thermal tolerance, it's possible it could successfully occupy an environment; if it doesn't, then it couldn't.

The inputs in the model *must be* functional because all of this talk of success and failure that the ecologists are relying on is built into the discussion. To evaluate performance, we must have a sense of what would count as performing well and poorly. It's this that function-talk gets us. Assigning a function to a trait lets us differentiate between cases of success and failure. Without the ability to do so, we would miss out on the whole aim of this predictive enterprise. We would be unable to identify which traits we ought to use as inputs – which are good predictors of future success or failure (cf. Neander 2017). Garcia et al. don't just happen to use functions as inputs. *They have to use functions as inputs.*

But can realized functions really get us this sort of telicity? I have already argued above that the inputs into ecological niche models are well-understood as parts of functional explanations as those explanations have been historically understood. The inputs help explain why the traits that enable the organism to do something will exist in the future via an appeal to what the traits enable the organism to do in the present. Someone with a backward-looking bent might argue, though, that there is an important distinction to be made between functions ascribed on the basis of casual history and those ascribed on the basis of future causal relations. The former are genuine functions, because they explain the present existence of a trait, but the latter are merely *possible functions*, because they do not yet explain the actual existence of a trait.

Certainly Garson (2019) would argue something along these lines, given his understanding of what functions are meant to explain in the first place:

Here's what I take the core puzzle of function to be: How can an effect of a trait play into a causal explanation of that trait's very existence? Immediately, we can see that any theory that doesn't appeal to history is useless to us... That's because a causal explanation for a current event must refer to things that happened in the past (p. 43-44).

According to Garson, then, functional explanations ought fundamentally to concern *present* existence, not future existence. At best, then, realized functions would turn out to be possible future functions, not functions in their own right right now.

Philosophers' focus on history here is, I would argue, an artifact of our past focus on evolutionary biology. By no means was this focus a bad thing – it was, perhaps, even adaptive in the philosopher's niche at the time. Evolutionary biology was a much more developed field than ecology, with more robust scientific practices to investigate. These practices were also more obviously concerned with naturalizing normativity, insofar as evolutionary biologists were constantly concerned with function-language.¹² There were also important social elements that promoted philosophical engagement with evolutionary biology in particular. For instance, many working evolutionary biologists (e.g., Lewontin, Gould, Michod, Levins, and so on) were eager to speak to and work with philosophers. Moreover, ecology was, for a long time, itself deeply preoccupied with making itself more explicitly evolutionary (e.g., Hutchinson 1957 and 1965; Harper 1982). All of this made and makes good sense, and it is little wonder that philosophers ended up so preoccupied with historical functional explanation. Having said that, nothing about functional explanation necessitates that we focus exclusively on historical explanations. We've seen already with Tinbergen (1963) that traits do not have to be considered exclusively as the effects of past causes. They can also be – *and in fact are* – the causes of future effects as well, including the future existence of traits. Had earlier philosophers of biology paid more attention to ecology, they would have noticed this alternative temporal framing for functional explanations. That we didn't is a matter of historical contingency.

¹² This is especially important to keep in mind when we consider why some of these philosophers were interested in naturalizing normativity in the first place, namely, as a way to understand the normativity of representations. I'll discuss this in greater detail in the next chapter.

There are two further points I'll make concerning ahistorical and forward-looking functional explanations. First, if I'm right (1) that Garcia et al.'s model both involves and requires functional inputs and (2) that those functions must be discoverable without an appeal to history given the evolutionarily novel conditions populations will encounter going into the future, then Garson being right would suggest these ecological niche models are destined for a kind of conceptual failure. This, however, doesn't appear to be the case. They are proving extremely useful both in conservation biology and elsewhere in ecology. As such, the claim that functional explanations must only concern the existence of traits in the present must be getting something importantly wrong about the nature of functional explanation.

Second, even if it were the case that the *typical* understanding of functional explanation in philosophy excluded the explanation of the future existence of traits, that still wouldn't mean that *all* functional explanations must take precisely that form. Consider another quote from Tinbergen:

To those, however, who argue that the only function of studies of survival value is to strengthen the theory of natural selection I should like to say: even if the present-day animals were created the way they are now, the fact that they manage to survive would pose the problem of how they do this (1963, p. 423).

Function-talk would not disappear overnight if the theory of special creation turned out to be true. There would still be questions about how an organism or population managed to survive in a hostile, changing environment that would require our attention even if investigating it told us nothing about an organism's origins. Moreover, these open questions would still require that we consider traits not as purely neutral things, but as things that can enable an individual to succeed or fail to survive and reproduce in some set of environmental conditions. To do this, we would need to be able to identify cases of functioning well and malfunctioning in order to know how a population was able to persist in an environment, even if it had no evolutionary history at all. Telicity is around in ecology even in the absence of evolutionary applications.

Supposing I've convinced my interlocutor, a question still remains: do realized functions give us a notion of functioning well and malfunctioning? Can the notion I've put forward tell us how a function can "linger, like a ghost, even when the corresponding trait vanishes" (Garson 2019, p. 28)? Yes, and here's why. The key point is that realized functions are determined relative to the population of which the organism is a member. To quote myself

Realized function – Some trait x has the realized function f just in case (1) x allows *an* organism to f and (2) *the ability to f helps to explain why some environment matches onto part of the organism's population's realized niche.*

An organism's trait's realized function is not determined by what it allows that particular organism to do, but what it allows *some* organism who is a member of the same population to do that helps us to explain the overall population's realized niche. If the presence of the trait in the population and the interactions with some environmental condition or conditions helps us explain why that population has the realized niche it does, then that trait has a realized function even if a particular member of the population can't perform that function.

I want to offer two examples here, the first being a case of a more-or-less static realized function – one associated with thermal tolerance – and the second being one that a single organism could both succeed and fail to fulfill at different points in time – a threat display put on by an adult female bird that wards off nest predators. We've already seen why thermal tolerance counts as a realized function, so I won't reiterate that here. Instead, I want to talk about a case in which an organism might *fail* with respect to thermal tolerance. Imagine that a frog goes into hibernation when the mean temperature dips below some threshold, and that this behavior is instigated by the production of some protein. The behavior's realized function is to help the frogs survive through periods of colder temperatures, since their ability to occupy their realized niche is in part explained by this. Suppose, though, that some individual frog has a mutation in the gene coding for that protein that leads to a delay in the protein's production until the mean temperature is even lower. This causes the frog to go into hibernation slower than its conspecifics and so ends up being exposed to colder temperatures. These colder temperatures lead to a higher chance of death. Because of the mutation, the frog cannot behave otherwise than it does. Nevertheless, their behavioral trait to enter hibernation at some mean temperature threshold has the function of contributing to their tolerance of temperature changes in their environment because that is what it typically does that helps to explain the *population's* realized niche. Therefore, it isn't just that the individual's trait *lacks* the relevant realized function; the individual's trait *fails* to perform it because of the frog's genetic mutation. The function lingers thanks to the population, thereby explaining how the individual's trait can malfunction.

Let's turn now to the second example. Imagine that a female bird responds to the approach of a predator after the eggs in her nest by puffing herself up, flapping her wings, and so

on. Doing so sometimes causes the predator to leave. If it fails to do so, the mother will not directly confront the predator, though she might continue the display. In this way, the threat display helps to explain the population's realized niche: it is in part because the behavior scares away local predators that the bird population can successfully reproduce, thereby enabling the population to successfully occupy their environment. As such, the behavior's realized function is to scare away potential predators. The behavior, of course, does not have a 100% success rate, for one of two reasons. First, the female might do everything right and still not scare off the nest predator. In this case, the failure, if there is one, seems more attributable to the environment than the female bird. Second, though, the behavior might fail because the female's performance is lackluster in some way. She might not be as active, or as puffed up, or any number of other possible aberrant deviations from the ideal behavior, where this deviation isn't inevitable but instead the result of some passing state. In a case like this, the female's behavior is failing. In this second case, it's not as if the behavior has lost its function. It is instead *malfunctioning*.

Garcia et al.'s ecological niche model both uses and requires telic inputs. This is because it needs to be able to differentiate between cases of success and failure, and the response-mediating traits uses as inputs (or, at least, the outcomes of those response-mediating traits) are organized in accordance with their potential positive impacts on fitness. I've also argued that functional explanation can be just as much forward-looking as backward-looking, but even if it isn't as described in the literature, we still need to differentiate between successes and failures outside of a backward-looking context. Finally, realized functions really can support the sort of telicity we need here. They tell us not just what a trait does for an individual, but what it should be able to do given its role in our explanations for a population's realized niche. As such, a trait's realized function allows us to say when an individual's trait is functioning well or malfunctioning. There really is telicity in these sorts of models, and realized functions have got it.

To finish out this response and to further clarify the sort of telicity at play here, it would be helpful to consider how realized functions draw the line between functions and accidents when compared to the selected effects (SE) notion of function. Here I'm using "accident" as a term of art. In the context of evolutionary biology, "function" is typically used to refer to what has been selected for and "accident" to other features that are just "along for the ride." It might be nomologically necessary that they come along, too, but it's not *them* that are being selected

for. For instance, it's because the heart was selected for its ability to pump blood that it has the function of pumping blood. That it also makes interesting rhythmic noises is just an accident. Perhaps, given the way the world is, the heart *must* make the rhythmic noises it makes, but those rhythmic noises don't figure into our evolutionary explanation for the heart.

The function/accident distinction has played a key role in evolutionary biology, in large part because it turns out that it has sometimes been extremely difficult to differentiate between the two. For instance, in historical debates over mutualisms between species (e.g., between a pollinator and a flowering plant), folks have argued over whether pollinators have faced selection not just to be able to collect nectar and pollen, but also to move that pollen from one flower to another. The question has also reared its ugly head in debates over adaptationist methodology, which takes for granted the traits we see in an organism *must* have been selected for to exist in the present. Not all traits are like this, though. As Gould and Lewontin (1979) put it, some of them are like the spandrels of cathedrals, which exist because they have to in order to support the architecture, not because architects chose to put them there for ornamentation. To offer Gould and Lewontin's biological case, certain aspects of the structure of mollusk and some brachiopod shells, while ubiquitous, appear to be there not because they are themselves adaptive, but because other traits that have been selected for just happen to produce them.

SE functions seem to cleanly divide the true functions from the accidents, in the senses of those terms used above. The definition of a realized function that I've put forward, though, provides us with a very different dividing line. Rather than being concerned with what a trait has historically been selected for, it considers the role a trait plays in explaining the population's realized niche. Again, let's consider the heart. In the modern world of cardiologists, the heart may well have the realized function of making certain rhythmic noises. After all, those noises help to diagnose all sorts of ailments, thereby enabling doctors to tend to them. To expand on this a bit, there is a cohesive and generalizable story for us to tell here. We take advantage of the fact that our heart makes certain noises in order to better diagnose when things are going well or poorly with it, and these diagnoses are made possible by nomological relations between the way the heart works and the sounds it makes. The heart makes the sounds it does because of the way it pumps blood, and specific variations of those noises are related to specific variations in that activity. Contrast this with a true accident. I might avoid being struck by lightning while walking with a group of friends through a thunderstorm because I'm shorter than my friend who was

struck. That doesn't mean, though, that my shorter height has the realized function of helping me to avoid getting struck by lightning, because that fact doesn't play into explanations for why humans fit into our realized niche!

The notion of realized function draws a very different line between functions and accidents when compared to the distinction drawn by evolutionary biology. If I'm right about what ecological niche models need as inputs, then this is just the line we need drawn. The models cannot make do with only those things deemed functions by evolutionary biology, since history, as we've seen, isn't enough. We need more than just historically oriented criteria; we also need functions that can be picked out by considering a population relative to its present and its possible future environments, and the functions we pick out there are nonaccidental in that they sustain the generalizations about the relevant populations that we need for the work in virtue of certain features in those populations.

4.2. The Relativity Objection

Calling on the role of a population may have generated telicity and helped maintain the function/accident distinction, but it also brings up the second potential worry: just which population ought we to use in assigning realized functions? Consider the Western honey bee. We could subdivide all of the living members of the species in any number of ways, considering those in Asia or Europe, in Europe alone, in Germany, or in Saxony. Considering these different populations may well result in our assigning different functions. Even besides the scaling effects mentioned above, a local population likely has adaptations to its particular surroundings (e.g., adjusting colony behaviors to best sync with the timetable of local flowering plants) that explain why those environments match onto that particular population's realized niche. Those local adaptations are therefore the realized functions of some of their traits. Another population in another locale, though, will lack those realized functions because they *don't* help to explain *its* realized niche. But what do we get if we conglomerate a bit? When we consider both populations together, we find that we need to appeal to the realized functions of the one subpopulation in order to explain the overall population's realized niche. But now we've done something odd. We've said that the bees both have and lack those realized functions.

There's a cheap objection here accusing this notion of function of producing contradictions. Certainly, on its face, it seems to, given the situation I've just described. Remember, though, that realized functions are assigned relative to a population. As such, we

aren't actually saying *P* and its negation when we assign different functions to honey bee traits at different scales and locations. We're instead saying that those traits play an explanatory role in one context and that they don't play that role in a different one. We're not straightforwardly contradicting ourselves, then, but we might be doing something just as bad for a naturalist. We might be making nonnatural claims. This objection will be familiar to those involved in debates about teleosemantics, as it mirrors one raised by Dretske (1986) and elaborated on by Fodor (1992).¹³ If it turns out that which function we assign depends not on the world, but instead on the interests of the researcher, then the choice between functional ascriptions comes out *arbitrary*. The notion of function at play is not fully naturalistic. Seeing as realized functions are *meant* to be naturalistic, that would prove a problem for me.

My response here hinges on an argument from the introduction of the dissertation. Wimsatt (1972; 2007) claims that systems can be organized in more than one way at once. These different organizations arise because an object can be divided in different ways into parts that are not spatially or, as I extended the argument, temporally coincident. To reuse my past example, a population considered over an evolutionary timescale is not the same as that population considered over an ecological timescale: the former concerns more objects across a wider temporal window, while the latter concerns relatively fewer objects across a narrower temporal window.

One consequence of this proliferation of organizations is that a system we might initially believe is well-understood from a single spatial or temporal perspective can actually only be understood or explained if we take on a variety of such perspectives. This in turn means that an explanation or prediction at one scale can be (quasi-)independent of an explanation or prediction at another. Not to sound like a broken record, but why some trait evolved in the first place can be independent from what good that trait does an organism in the present day. Now, a diehard reductionist might object to all of this on the basis of their metaphysics. Sure, we don't know how to explain everything from some one true perspective, but that doesn't mean that one true perspective doesn't exist. I'm happy, though, to rely here only on the epistemic point. As it stands, it doesn't appear that we can do anything but act as if more than one scalar perspective is

¹³ Dretske, of course, was concerned with selected-effects functions, but the objection could be applied here, too, *mutatis mutandis*.

legitimate. Perhaps we will someday have a completed science with a desert ontology, but, for now, as Wimsatt colorfully puts it, we're stuck with a rainforest (Wimsatt 2007).

To repeat myself a bit more, in the case of ecological niche models, different realized functions explain different things depending on the scale. Just which characteristics explain or predict a population's ability to occupy an environment, for instance, changes with spatial scale: presence or absence is effectively random when considering square feet, is often determined by interspecies interactions at slightly larger scales, and then by suitable habitat and climate as you go larger still (McGill 2010). What exactly does the explaining can also vary with location, as illustrated by the honey bee case that started this section. Which traits matter for our explanation depends on which environment we're considering the honey bees in relation to. Finally, temporal scale can play a significant role. Whether a population will be able to occupy some environment or not depends, in part, on when it encounters it. As climate change progresses, some environments will become climatically suitable and others will stop being so, as we saw with the amphibian species in Garcia et al. (2014). In order to account for all of these different cases and contexts, then, realized functions *must* be able to change as the scale or location under consideration does.

Nothing about this variation is either arbitrary or nonnatural. All of the facts grounding the functions are out there in the world. Which facts are doing the grounding varies with scale and location, but that's no surprise. If we want to know how all the honey bees in Germany will do under a changing climate in the next forty years, we can't very well avoid asking about honey bees at those spatial and temporal scales. Moreover, once we're looking at that particular scale, *all the facts are set*. We have no say in which features of the world matter for the prediction, only which prediction we're trying to make. Perhaps there's something arbitrary in our choice to ask a question in the first place, but that's far removed from our *answer* being so. In any case, it would be a strange argument indeed to say that an answer to some scientific question was arbitrary merely because we decided by chance to ask it.

5. Conclusion

I've argued that realized functions help us to predict and explain whether a population can occupy an environment, that they don't require that we appeal to history, and that they can be ascribed modally. As such, they have the right characteristics to act as inputs in some ecological niche models, like the one used by Garcia et al. (2014). They therefore have a vital role to play in

certain areas of conservation biology. But they can do more work for us besides. In the coming chapters, I'll show how they can be used to help the teleosemanticist respond both to the problem of indeterminacy and to Swampman. For the moment, though, I want to talk about their possible utility in some another area of biology and philosophy of biology.

Though I cannot explore the issue in detail here, I suspect that realized functions could be used to at least partially account for what ecologists mean when they discuss ecological functions, often glossed as the contributions of populations to the persistence of the ecosystem of which they are members. There is an ongoing and lively debate over the exact nature of these functions. Some have argued that they are best understood as causal-role functions concerning the contribution of an organism or population to the ecosystem of which it is a member (e.g., Dussault 2018; Dussault 2022; Morrow 2023). Others have attempted to subsume ecological functions to selected-effects functions (e.g., Millstein 2020), arguing that they can be best-understood as coevolved phenomena. I would suggest that at least some ascriptions of ecological function, especially those closely tied together with claims about traits, are likely well-understood as realized functions on the basis that they explain why an environment matches onto a population's realized niche via a claim about the contribution of the population to the ecosystem of which it is a member. Honey bees, for instance, do not just fit into their environment because they can collect nectar and pollen, but because they pollinate some of the flowers they visit, thereby helping their food source to propagate. Realized functions could thereby explain both what good it does the population for its members to perform some activity and the significance of that activity for the ecosystem without having either (1) to avoid connecting function-talk with fitness or (2) to ascribe any story about coevolution, which may well not be truthful.

There is much more to say, but I ought to set strict philosophy of biology aside for now so that we can orient ourselves towards the second aim of this dissertation: the defense of teleosemantics. We'll start in the next chapter by reexamining the problem of indeterminacy as posed by Fodor's frogs. Realized functions, as identified and developed here, will prove useful in resolving at least one variant of this problem.

Chapter 2: What Fodor's Frogs(' Visual Systems) See

1. Introduction

We've talked already about the ways in which cases of malfunction may appear mysterious for a naturalistic account of biological function. As Garson put it, "How can a trait possess a function it cannot perform? By virtue of what does a function linger, like a ghost, even when the corresponding ability has vanished" (2019, p. 28)? If a trait cannot, even in principle, do something, then it is difficult to see what facts could ground our claim that it *ought* to be able to. But this problem is really just one instance of a broader class concerning the difficulty of making claims about normativity in a fully naturalized vision of the world. Here I want to turn our focus towards another instance of the problem: representation, or, more precisely, *misrepresentation*.

It is, perhaps, clear enough how we could tell a naturalistic and causal story connecting a true representation and that which it represents. The representation could be tokened *because* that which it represents is out in the world, however we might end up caching out the causal relation. It's much more difficult to see how this could work in the case of misrepresentation, though. Or, rather, it's unclear how anyone ever could misrepresent (Fodor 1984). Consider, for instance, a representation of a horse. We might suppose that the representation of a horse is representing correctly when it bears the right causal relation to a horse in the world. So far so good, but now suppose that, in the dark of night, one misrepresents a cow *as* a horse. You think you've seen a horse, but really it's just a cow. How could this occur on a causal account of representation? Well, the cow apparently played a similar enough causal role that seeing it lead to the tokening of a representation of a horse, but then what's so special about a horse? If a cow can do just as well – if things other than the content of the representation can play the same causal role – and it's the causal role that decides the content of the representation, then nothing about the causal role tells us whether the tokening of the representation in response to a cow is a case of misrepresentation or not. There can therefore be no misrepresentation on a causal account. There can only be an increasing number of disjuncts in the content of the representation as we learn that more and more things can play the right causal role.

This is the disjunction problem: how can we differentiate between a truthful representation with a disjunctive content and a misrepresentation when telling a naturalistic, causal story about reference? In the last chapter, we saw how attributing a function to a trait

allows us to identify cases in which that trait is malfunctioning; a function tells us what a trait ought to do, and so provides us with a sense of malfunction, as well. A set of philosophers began serious work in the 1980s trying to apply naturalistic notions of function to solve the disjunction problem. If we can cache out the causal story explaining a representation's content not just in terms of any old causal story, but in terms of the representation's *function*, and that function is determinate, then we can definitively say which tokens of a representation are *misrepresenting* the world. Knowing a representation's function would allow us both to identify a special class of external inputs – those that align with the proper functioning of the tokened representation – as the inputs associated with the content and to discount other inputs as candidate contents on the basis that the tokening of the representation in response to them would be a malfunction. We just need a robust, naturalistic notion of function to get us there.

Millikan (1984) and Neander (1983) independently developed versions of this theory of meaning – teleosemantics – in which the kind of function relevant for this story is a biological one.¹⁴ Both of these authors, as well as many others inspired by them, have further argued that the relevant biological functions are selected-effects (SE) functions, where a trait's SE function is what it has historically faced selection to be able to do. A token representation's content is therefore dictated by what the representing machinery has historically faced selection to be able to represent.

It's easy to see the appeal of a theory like this. We typically think of selection as, well, *selective*. The heart didn't evolve the way it did in order to make certain rhythmic sounds, but in order to pump blood. Selection, then, ought to provide us with determinate functions for our representing systems. These determinate functions can then ground determinate contents of the representations produced by those systems. An appeal to an SE function, then, gives us the ability to differentiate between true representations and misrepresentations. A tokened representation is true when it accurately represents what it has faced selection to represent, and it is otherwise misrepresenting. These accounts of SE function, though, face a problem. Under the right circumstances, they seem to fail to differentiate between cases of successful representation and of misrepresentation due to their inability to supply a determinate function. Functions, the story goes, can't fix contents, because something's SE function is indeterminate (Dretske 1986).

¹⁴ Millikan has said this is a misnomer (CITE), but the gloss works well enough for us here that I'll stick with it.

If performing x always coinstantiated with performing y in an organism's environment of evolutionary adaptation such that the effect of performing one was indistinguishable from the effect of performing the other, then at best we can say that an organism has the function of performing “ x or y .” We get stuck with the same nasty case of, to quote Neander (2017), “disjunctionitis.”

The problem is even worse when we consider the evolution of representational systems. The fundamental problem is that we can very often represent the same thing in the world in a variety of different ways. That is, the same object can be the referent of a wide variety of contents that refer to different properties of that object. For instance, I can pick out my computer by referring to “this computer,” “the thing on the desk in front of me,” “this set of shapes and colors,” and so on and so forth. Insofar as different contents successfully refer to the same thing in the world, the choice between these contents will be invisible to selection (Burge 2010). The problem Dretske raised runs rampant, then, when we consider representations.

Fodor (1992) put this point in special relief. Suppose that frogs evolved in an environment in which:

- (1) flies were their only food source;
- (2) the only flies they saw were also small, dark, moving things (SDMs) (i.e. they only ever saw flying flies); and
- (3) the only SDMs they saw were flies.

In the imagined scenario, three different things have always coinstantiated in the frog's evolutionary history: flies, frog food, and SDMs. Selection for its nervous system to represent one of them would do just as well as selection to notice any of the others. The frogs have faced selection to be able to snap at the flies in their environment, but what they represent the flies *as* doesn't matter so long as the content gets the job done. Because the fitness effects of representing these different contents are equivalent in the frog's historical environment, a story about selection only tells us that the frog's internal state picks out “flies, or frog food, or SDMs.”

But it seems like we ought to be able to choose between these, because doing so is necessary for us to give an adequate account of *misrepresentation*. To see why, imagine that a frog is suddenly moved from its natural environment, in which (1)-(3) held, to a laboratory. In this lab, researchers sometimes release flies for the frog to eat, but they also sometimes throw little black BBs by the frog instead. In this new environment, then, (1) and (2) hold, but (3) no

longer does. Some SDMs are not flies. Whether it is exposed to real flies or to BBs, the frog responds identically: it tries to eat them. Our question is, when it sees and eats a BB, has its nervous system represented its environment veridically or not? The answer depends on what the frog is representing. If its nervous system is trying to pick out either flies or frog food, it has obviously failed, since it has labeled a BB a fly or an instance of frog food. If, on the other hand, its nervous system is picking out SDMs, it has represented its environment veridically. Perhaps we can get around this by an appeal to researcher interest. We could, say, narrow the function down to “detecting flies” or “detecting frog food” because evolutionary biologists are interested in the features of the environment relevant for the frogs’ historical success. Flies and frog food are so relevant, but SDMs aren’t. But now the choice between these options, the objection goes, comes out arbitrary. It is dependent on the interests of the researchers, not on the facts about the frog and its environment, since, from the frog’s perspective, representing any of (1)-(3) would do just as well. As such, teleosemantics fails to naturalize the determinate content of the frog’s cognitive state, and so fails in turn to provide us with an adequate account of misrepresentation. It does not provide us with the telicity we need to get the job done.

The problem of indeterminacy, as the grammar of the name suggests, has typically been discussed as if it were a single, unified issue. There are, however, a pair of logically distinct problems that arise from Fodor’s case, and these different problems require different solutions. I argue that the teleosemanticist’s failure to adequately respond to the challenge has been, in part, the result of their attempt to address both of these problems with the same solution, namely, some refinement of a backward-looking notion of function. These different refinements do a better and worse job addressing variants of the frog case, but none of them can offer us a solution for all the variants. What the teleosemanticist really needs isn’t a better definition of selected-effects function, but an additional notion of biological function that can better address the ecological version of the problem.

In §2 of this chapter, I distinguish between two versions of the problem of indeterminacy. Among these will be the problem of *functional* indeterminacy, which was originally raised by Dretske and that we will see arises most obviously in the context of problem-solving in the biological sciences. In §3, I further distinguish between two variants of this problem that arise in evolutionary biology and ecology respectively. In §4, I argue that only one of these problems can be solved by an appeal to backward-looking functions. Though any given backward-looking

solution will likely solve the problem posed for ecology in some contexts, none of them can in principle solve it across *all* contexts. In §5, I show how the problem of functional indeterminacy for ecology can be solved by an appeal to realized functions. Finally, in the conclusion of the chapter, I sketch out how realized functions might help us to solve another problem of indeterminacy: the one concerning *content matching* with which Fodor was predominantly concerned. Realized functions can often provide narrower content attributions in virtue of their sensitivity to the particulars of the environment an organism is inhabiting along with that organism's interactions with its features, as we saw in the previous chapter.

2. Two Variants of the Problem

Just what the frogs' visual systems represent seems to be indeterminate, but just what problem this indeterminacy leads to in our practices is less obvious than it at first appears. Here I want to identify two ways in which indeterminacy could be problematic. Because it is genuinely unclear whether both of the variants of the problem of indeterminacy I discuss here are applicable to the frog case,¹⁵ it's best if we look to humans for the first problem, which I'll call the problem of *content matching*. We'll return to the frogs to illustrate the second, the problem of functional indeterminacy.

To start, consider a general feature of representations we often take for granted in psychological explanations: that they have contents that link up neatly with contents represented in natural language. When someone thinks something in, for instance, English, we – mostly – take for granted that they are thinking about the same thing represented in English itself. Now, I've put that in folk-psychological terms, but it's true that a great deal of cognitive science assumes that it's possible for us to link, at some explanatorily useful level of description, the contents of our cognitive states to contents that are expressed in natural language (e.g., in mathematical reasoning, when I think about the number one, I'm thinking about just the same thing as is expressed by the English language phrase “the number one” or in mathematical terms by the Arabic numeral “1”). Natural languages are, of course, amazingly precise. We can differentiate between any number of not just nomologically coinstantiated properties, but metaphysically coinstantiated ones as well. I can talk about, and thereby differentiate between,

¹⁵ More particularly, what I'll call the problem of content matching may not arise with Fodor's frogs, as it only matters if a psychological explanation is necessary for some phenomenon. I don't particularly want to argue about whether we need to appeal to psychology to explain frog-phenomena, so I use a different example to illustrate the problem.

“triangles,” “shapes that, in Euclidean geometry, have internal angles adding up to 180°,” and “trilateral figures.” Mental representations that match onto these English phrases can thereby do the same. If teleosemantics is unable to provide us with functions that differentiate between representing nomological covariants, let alone *metaphysical* covariants – if anything would do just as well for selection, surely a metaphysical covariant would – then a teleosemanticist’s efforts to explain how a human mental representation could misrepresent is doomed. It’s this problem that Fodor seems to specifically have in mind: “In effect, Darwinian explanations treat reliably coextensive representations as synonymous; whereas, of course, psychological explanations don’t” (1992, p. 78).

Call this the problem of content matching. While this might at first seem identical to the problem of indeterminacy as originally stated, it’s worth differentiating because it requires, as we will see, a different solution than the other variant does. To resolve the problem of content matching, we don’t just need any old determinate content. We instead need a content that matches onto something represented by a *particular bit* of natural language. It won’t do to get something “close enough,” or a mere synonym. Some of our cognition makes use of natural language, so any content less hyper-specific than natural language just won’t be good enough.

Let’s turn now to the second variant, the problem of functional indeterminacy, and contemplate the frogs once again. While put in terms of the content of the representation produced by their visual system, the problem arises because we can’t decide whether their visual system is functioning well or malfunctioning when it triggers the reflex to eat a BB. If the visual system is meant to represent a fly or frog food, then it’s failing; if it’s meant to represent a small, moving dark object, then it’s succeeding. If we can’t decide between the possible contents, then we can’t decide whether it’s functioning well or malfunctioning. Frogs’ visual systems, though, do not represent things in English, or, to the best of my knowledge, any other natural language. We don’t, then, need to find an exact match. All we need is something that is a good enough synonym that allows us to differentiate between cases of proper functioning and malfunctioning. We are not constrained to make it a neat or tidy natural language predicate; we just need it to do the right explanatory work in biology. Even a disjunction could do just fine.

It’s this problem – the problem of functional indeterminacy – that I mean to address here, and I hope that I’ve shown that solving it could be quite a bit simpler than solving the problem of content matching, so long as biology cooperates with us. All we need is *some* determinate

content. Still, a solution is not so easily grasped as this might imply. While going in for a strange (in English) content eliminates the indeterminacy – “the presence of an item having certain biochemical properties, those biochemical properties that make flies nutritious to frogs” (Price 1998, p. 71) comes to mind – it won’t do so in a way that is satisfactory across all explanatory contexts in biology. More particularly, because we want different things out of functional explanations in evolutionary biology and in ecology, no one functional attribution can do all of the relevant work. A solution to the problem of functional indeterminacy needs to explain things not just for evolutionary biologists, but for ecologists, too.

Before moving on to discuss these conflicting demands, though, it’s worth taking a step back to consider the sorts of contexts in which the problem of functional indeterminacy, and so of misrepresentation, arises for teleosemanticists in general. The issue is not, typically, with fixing a content given some evolutionary context. We can always find a content to do the trick. The problem appears when we move the frogs into the novel environment in which what have been nomological covariants break apart. The problem is, fundamentally, associated with *novelty*. This is nothing new for teleosemantics, or for the notion of an SE function. Many of the oldest and most persistent objections to these theories have had to do with novel environmental features (Lyons 1995; Garson and Papineau 2019). I’ll return to this later in the chapter, and again in the next when discussing Swampman. Here I just want to start us down this path, that SE functions and the theories of content that depend on them have a general problem with novelty.

3. Two Questions about the Frogs’ Visual Systems

Before showing why a single function won’t do for the frogs’ visual systems, it’s worth spending a moment to review the two sorts of functional explanations we find in evolutionary biology and ecology, respectively. We’ve seen already in the introduction to this dissertation that these two branches of biology can sometimes use traits in very different ways. In evolutionary biology, traits are the thing to be explained. Why does this trait exist in this population at this frequency? Because it enabled ancestors of the present population to interact with their environment in certain ways. In ecology, traits are often used as *inputs* into explanations. Why is this population able to successfully occupy the environments it does? Because the traits in the population presently allow them to interact with those environments in certain ways. In the last chapter, we saw how this distinction plays out with respect to functional explanation. A trait’s

function in evolutionary biology is what it does that helps to explain its present existence and distribution. In ecology, it can be used to explain and predict a population's realized niche. While a single function can sometimes do all of this explaining, there are times when it will fail to. Most prominently, we saw this occur when a population encounters evolutionarily novel environmental conditions. In those cases, what some trait was selected for can come apart from how that trait contributes to a population's ability to occupy the novel environment. If a population has never had to deal with certain climatic conditions, for instance, knowing how its traits enabled it to survive in past climatic conditions won't give us enough information to explain and predict its ability to survive in the novel ones. It was for this reason that I needed to identify and develop the notion of a realized function in the first place. Explaining and predicting populations' responses to anthropogenic climate change in conservation biology required that there were functions we could discover by considering only the population now and in the future, not in the past.

This is all to say that, when we ask about a trait's function, we might be asking one of two questions: what explains the trait's present existence and distribution; or what role does the trait play in our explanation for the population's ability to occupy its realized niche? When we ask about the function of the frog's visual systems, then, we might be asking one of two questions. First, we might ask what their visual systems allowed their ancestors to do that contributed to their fitness, thereby explaining the present existence and distribution of the trait¹⁶ in the population today. Second, we might ask how their visual systems help to explain why they have the realized niche they do.

Fodor argues that the problem of indeterminacy arises for the frogs under precisely the same circumstances under which I've argued that the answers to the above two questions come apart. When the frogs enter the evolutionarily novel lab setting, we can't tell whether they are functioning well or malfunctioning when they snap at BBs. The issue is *not* identifying a content via a function, but identifying a content by relying only on a notion of function – SE function – that cannot adjust in the face of novelty. It's the problem of novelty with another face.

I have briefly argued above that, so long as we can attribute a determinate content, however odd it might be in natural language, we can say that the visual systems are functioning properly or malfunctioning when they snap at a BB. This is not, however, the only worry we can

¹⁶ Or, perhaps better, the collection of traits making up their visual system.

raise about the content. First, as I've just articulated, even supposing that's right, we might still worry that one functional ascription can't do all of the work we need it to across both evolutionary biology and ecology. Can it both explain the present existence and distribution of the frogs' visual systems and help to explain why the frogs have the realized niche they do? Second, whatever an organism evolved to be able to do might just feel wrong to us if we want functions to match onto what would be good for an organism to do in the present. Given the ink spilled by both teleosemanicists and evolutionary biologists trying to explain why what is good for an individual can come apart from why a trait evolved in the first place (e.g., Neander 2017; Gould and Vrba 1982), I suspect that something like this question is lurking behind many answers (e.g., Millikan 1990) and objections to answers to the problem of indeterminacy. It's not because there's an SDM that it's good for the frogs to snap, and so SDM can't be the content of the representation.

In the next couple of sections, I'll do two things. First, I'll show that, while many SE functions can do a lot of the explanatory work in both evolutionary biology and ecology, *no* SE function can do all of it, for the same sorts of reasons I discussed in the last chapter. As a result, we need to attribute more than one function to the frogs' visual systems: we need (at least) SE functions to do the evolutionary work and realized functions to do the ecological work. Second, in showing how we can attribute a realized function to the frogs' visual system, I hope to also assuage worries about whether the function matches properly onto our question about what's good for the organism to do. It's because SE functions ignore this issue that they are unsatisfactory answers given the theory's opponents' intuitions about the case. Admitting both SE and realized functions into our ontology enables us to give a more intuitively satisfying response to this worry.

4. SE Functions and the Ecological Question

To briefly restate the original problem, it seems as though Fodor's frogs would have done just as well in their environment of evolutionary adaptation had they faced selection to see any of the following:

- (1) flies were their only food source;
- (2) the only flies they saw were also small, dark, moving things (SDMs) (i.e. they only ever saw flying flies); and
- (3) the only SDMs they saw were flies.

As such, teleosemantics cannot tell us whether the frog's visual system is representing well or misrepresenting when it triggers the reflex to snap at a BB. The theory therefore fails to deliver an adequate account of misrepresentation. While I am confident that we can deliver *some* determinate answer as to the frogs' visual systems' SE function, I am much less confident in which of the available answers is the right one. As a result, I have little interest in taking on the defense of any of the solutions on offer or in litigating between them. Instead, my aim here is just to show why none of them can provide a complete answer to the ecological question; none of them can give a satisfying explanation across ecological contexts as to why the frogs have the realized niche they do. With that in mind, I'll begin by offering up a schematic of an argument against SE functions as answers to that question. I'll then show how that schematic applies to Neander's (2017) SE solution to the problem of indeterminacy to give a sense of how it would run for any given solution on offer. I conclude with a brief summary of how the objection applies to a number of other proposed SE solutions.

To state the problem most generally, any given SE solution to the problem of indeterminacy will be *stuck in the past*. In short, and as we've seen time and again, if the present environment differs in ecologically significant ways from a population's environment of evolutionary adaptation, then what some trait evolved to be able to do can differ from its present contribution to the population's ability to occupy the environment. Let's walk through the frog example one more time to see why the change in the environment is so disruptive. Before the change occurred, whatever the frogs' visual system's SE function was would also explain the population's present success. After all, what ancestors of the population did that explained their survival and reproduction would be just the same as what the present population does that explains its survival and reproduction. The SE function serves just as well in both the evolutionary and ecological contexts. But then the environment changes; the link between nomological covariants *breaks*. Where once a fly, a piece of frog food, or an SDM were always guaranteed to be the same as the others, suddenly they can come apart. Besides a fly, an SDM can be a BB. Supposing the frogs can survive in the lab – the BBs are rare enough that the population won't simply die off – the SE function may or may not still explain their realized niche. It depends on what their visual systems were selected for. If their visual systems were selected to detect frog food or flies, then things would perhaps be fine, but if they were selected to track SDMs, then tracking SDMs wouldn't be what explained their present success, since

there would be SDMs that hurt them when they ate them in their present environment. Moreover, even supposing that their SE function *did* perfectly align with what explained their new realized niche, their environment could change in any number of other ways that would mean the SE function *stopped* working in this way. An SE function might succeed as an answer to the ecological question in some environments, but it would inevitably fail with respect to others.

Here, then, is the general template of the objection: while any given SE function would provide us with an answer to the ecological question in some contexts – call these “nomologically consistent contexts” – no SE function can provide answer across all ecological contexts – call the contexts in which an SE function fails “nomologically inconsistent contexts”. There will inevitably be some possible environment in which the SE function will come apart from the realized function. While this might appear a philosopher’s idle worry to someone interested only in evolutionary biology, we’ve already seen that it isn’t one in ecology. We anticipate that many populations will face evolutionarily novel environmental conditions going into the future, and some of these will, inevitably, mean that our answers to the evolutionary and ecological questions concerning function will come apart. Even if we don’t, though, the problem would remain. To beat the drum again, ecology is its own discipline with its own questions and problems that need to be addressed. Evolutionary biology is often intertwined with it, but it need not be. We need to be able to pursue answers to at least some ecological questions independently of evolutionary ones.

With the schematic in hand, let’s now see how we can apply it to a particular case. Neander argues that our best answer to the problem of functional indeterminacy is to say that the frogs’ visual systems represent SDMs (2017). Again, I won’t worry here about why she argues for this – her argument hangs on constraints drawn from cognitive science that give us a content as stuck in the past as any other – but instead will take her at face value to show why the answer just won’t work out in practice. We can see straight away that this is an instance of the second variant of the schematic argument. Snapping at BBs does not help to explain why the frogs have the realized niche that they do, because doing so isn’t good for them. Their visual systems triggering the reflex does not play into our explanation for the frogs’ new realized niche; if anything, they can survive in the lab *despite* the harm this does. But the BBs, for all their faults, are still SDMs. As such, Neander’s preferred SE function cannot do the right ecological

explanatory work: it does not line function up with what is good for the frogs in their present environment.

This isn't to say that this will, inevitably, be the case. Alternative circumstances would line up the content "SDM" with an answer to the ecological question. If, for instance, the scientists threw bits of hamburger instead, "to represent SDMs" could, perhaps, do just fine as both an SE and a realized function. But a determinate SE function, no matter which we subscribe to and however ideal it is as an answer to the evolutionary question, won't be able to do all of the explanatory work we need it to in ecology. It will likely work as both an SE and a realized function in some evolutionarily novel circumstances, but will fail to do so in others. What we need, then, isn't one function, but two: an SE function to handle the problem of indeterminacy in evolutionary biology and a realized function to handle the problem in ecology. I'll tackle the latter in a moment.

For the moment, though, I want to further emphasize the generality of this argument using the following chart of SE function contents suggested in the literature along with corresponding nomologically consistent and inconsistent contexts.

SE Function Content	Nomologically Consistent Contexts	Nomologically Inconsistent Contexts
"Frog food" (Millikan 1990)	In an environment that includes things that aren't flies but are frog food, "frog food" would help us to explain why the frogs had the realized niche they did.	In an environment in which there is plenty of frog food but none of it flies around and so remains effectively invisible to the frogs, "frog food" wouldn't explain the frogs' failure to thrive in that environment.
"Something at a certain time and place affording no-choice catching and swallowing" (Millikan 2023, p. 60)	In an environment in which hamburger bits are thrown past the frogs, this functional ascription would also help to explain why the frogs were able to successfully occupy the environment.	In an environment with BBs, this function would not help to explain how they successfully occupied the environment, supposing eating BBs was bad for them.
"Fly" (McGinn 1989; Sterelny 1990)	In an environment in which a novel species of fly is introduced that behaves in the	In an environment in which non-fly frog food is present, representing "fly" would not

	same manner as the flies the frogs evolved to catch, the content “fly” would help to explain the frogs’ ability to occupy the environment.	help to explain the full range of their visual systems’ contributions to the frogs’ ability to successfully occupy the environment.
“The presence of an item having certain biochemical properties, those biochemical properties that make flies nutritious to frogs” (Price 1998, p. 71)	In an environment in which non-fly frog food is present that was sufficiently chemically similar to flies, this content would help to explain the full range of their visual systems’ contributions to the frogs’ ability to successfully occupy the environment.	In an environment with frog food that is not biochemically similar to flies, this content would not help to explain the full range of their visual systems’ contributions to the frogs’ ability to successfully occupy the environment.

TABLE 1 A table showing various contents from various possible SE functions of the frogs’ visual systems available in the literature along with an environment in which it succeeds as an answer to the ecological question and one in which it fails.

These examples are not cherry-picked. Rather, they all follow the same general formula that can be applied to *any* SE function. We can find some circumstance under which the link between what had been nomological covariants breaks that will, in turn, make the SE function less explanatorily useful for understanding the frogs’ realized niche. In general, we can look for a context in which either the chosen feature or features are not sufficiently narrow as to explain the frogs’ realized niche (see “frog food”) or we can find a context in which snapping at something other than the ascribed content *does* help to explain their realized niche (e.g., “the presence of an item...”).

So, what do we do? How can we offer a satisfying answer to the ecological question? Let’s turn now to the frogs’ visual systems realized functions to see how this can be done.

5. The Determinate Realized Function of the Frogs’ Visual Systems

Let’s get that definition one more time.

Realized function – Some trait x has the realized function f just in case (1) x allows an organism to f and (2) the ability to f helps to explain why some environment matches onto part of the organism’s population’s realized niche.

Something I have not yet emphasized is that the above definition arms us not only with a way of recognizing realized functions, but also with the means to rule out many candidate functions. If

some ability does not help to explain why an environment's features correspond to a part of the population's realized niche, then a trait cannot have that ability as its realized function. The reverse is also true. If some environmental feature doesn't appear in the population's realized niche, then a trait can't have a function to interact with it. Here I follow Hutchinson (1957) concerning what features appear in a realized niche: we should abstract out from the particulars of the environment to a higher-level description of the role an environmental feature plays in the target population's ecology. Since the aim of attributing a realized function is to help explain its realized niche – why some part of a population's environment matches onto that realized niche – a realized function should refer to this higher-level description of the feature.

Now, with these constraints in mind, we can look at the frogs' visual systems' realized functions in their environment of evolutionary adaptation. The first thing to note is that, in describing their realized niche, "SDM" won't cut it as a content. The reason for this is straightforward: SDM isn't sufficiently explanatory to be a part of their realized niche. Triggering the reflex to snap at SDMs isn't what helps to explain why the environment matches onto a part of their realized niche because knowing something is an SDM doesn't tell us why snapping at it helps them to successfully occupy the environment. We also need to know that the SDM is *nutritious*, since it's in virtue of being nutritious that snapping at them is good for the frog. We can run a similar line with "fly." Because only flying flies appear in their environment and because the frogs only snap at flying flies – indeed, we're led to believe that they wouldn't snap at a stationary one! – "fly" isn't sufficiently explanatory, either. Even if we add "flying" to "fly," though, it won't quite do. What matters for the success of the frogs isn't, after all, that they're snapping at flies, but that they're snapping at *frog food*. Flies play the role of food in the frogs' ecology, so it's this feature that appears in their realized niche. Though, really, that isn't quite right, of course: as with "fly," it's not stationary frog food that does the trick, but flying frog food.

This leaves us, for now, with a pair of possible functions: detecting flying frog food, and detecting nutritious SDMs. I'll admit that I am not certain how we might decide between the two of these, supposing that "nutritious" is indexed to frogs. This is in part because they seem to me to be roughly equivalent contents. Certainly they are different *in English*, but, as I noted in §2, the frogs don't speak any language at all. Our content attribution will be, at best, an approximate translation, and quibbles over which we ought to choose feel a bit like a projection of worries

about synonyms that would be impossible to even articulate in terms of the mental representations available to the frog. Even if we couldn't decide between them, it shouldn't give rise to the problem of *functional* indeterminacy. If we end up with a closed disjunction, the frogs' visual systems can still malfunction – can still misrepresent – if they snap at something that isn't flying frog food or a nutritious SDM.

Still, I'll do my best to argue that we should endorse the content “nutritious SDM” over “flying frog food.” The reason is that “flying frog food” might itself not be sufficiently explanatory. If the frogs are unable to detect flying frog food that isn't, for instance, dark, then their visual systems couldn't have the function of detecting “flying frog food” *simpliciter*, since it wasn't flying frog food in general but flying frog food of a certain kind that appeared in their realized niche. This, of course, all depends on empirical details that are scarce in an imaginary case, and it could lead to a particularly bizarre content from our perspective. If such a content were demanded by the science, though, then that's just what we'd have to accept. So, I'll proceed supposing that the frogs' visual systems have the realized function of detecting nutritious SDMs, and so produce representations with the content “nutritious SDM.”

Let's move on now and suppose those dastardly scientists have moved the frogs into the lab and started throwing BBs. We have a break in the tie between what had been nomological covariants: fly, frog food, and SDM no longer constantly coinstantiate. Per the case, eating BBs does not help to explain why the frogs have the realized niche that they do, because doing so isn't good for them. (If anything, they might have their realized niche despite this failure.) Supposing the lab is still a part of the frogs' realized niche,¹⁷ then, their visual systems would be malfunctioning when they snapped at the BBs because they would be failing in their function of snapping at nutritious SDMs.¹⁸ In this version of the story, functional indeterminacy never rears its ugly head.

¹⁷ If it isn't a part of their realized niche, then the frogs in the lab would have no realized functions at all, at least if we consider the frogs in the lab alone a population. They would still have realized functions relative to the larger population of frogs, of course. If this seems like an odd consequence, I would direct my reader to chapter 1, §4.2, “The Relativity Objection.”

¹⁸ Or perhaps they are functioning properly in a new way! If scientists only keep the frogs around that snap at BBs, the frogs snapping at BBs might explain why they can successfully occupy that environment: because the scientists only keep the frogs alive that do so. Regardless, the realized function would differ from some possible SE functions.

We can run much the same analysis for other cases in which an SDM is introduced into the frogs' environment that's harmful for them to eat. Inevitably, the frogs' visual system's realized function will remain the same as it was in their environment of evolutionary adaptation such that the visual system will malfunction when it triggers the reflex to snap at the novel SDM. That takes care of many novel environmental features, but I still need to show how the notion of a realized function can handle novel environmental features that are good for the frog to snap at. Any case of a novel food source sufficiently similar to the flies – thrown bits of hamburger, a novel prey species – will be subsumed under “nutritious SDM,” and so won't require an adjustment to the realized function. More exotic cases might, though. Suppose, for instance, that the frogs begin snapping at small pellets that turn out to aid in the digestion of other food, as small stones do in the gizzards of some species. In that case, their visual systems *would* have a new realized function, namely, to represent either nutritious SDMs or the small pellets that aided in digestion, since it would be in virtue of detecting these two different things that their visual systems contributed to their ability to successfully occupy the novel environment. The realized function would change with the environment, and this change in realized function would allow us to do the relevant explanatory work.

We have here, then, a way of discovering the frogs' visual systems' realized functions that allows us both to avoid the problem of functional indeterminacy and to answer the ecological question. The realized niche does the work that selection was supposed to for backward-looking teleosemantics: it rules in and out a wide variety of possible contents by an appeal to biological facts on the ground. It also is able to handle the cases that were problematic for those backward-looking solutions because realized functions are meant to handle cases of novelty. Moreover, the realized functions so ascribed tell us what is good for the frogs to do now. We are not stuck saying only that the frog is doing something well when it snaps at and eats a BB. At the same time, we aren't forced to say its visual system wouldn't, under those circumstances, be doing what it had been selected to do, supposing that our preferred SE function is to detect SDMs. We can, at least in principle, have it both ways. The trick, it turns out, is for the teleosemanticist to embrace pluralism about biological function.

This pluralism about function, and so content, might initially seem suspicious. We might even raise the old worry that our choice of which function to ascribe is importantly dependent on researcher interest, and so subject to the same old objection: teleosemantics is not genuinely

naturalistic. I already have an answer to this challenge on the table, though. All of the functions, and the associated contents, are there at the same time, and they are dictated to us by the world. They are so dictated because we cannot help but explain how organisms are able to successfully occupy environments – or why they *fail* to occupy some environments – at different spatial and temporal scales. The same is true of our interests. I cannot avoid saying that Fodor’s frogs represent one thing or another once I’ve focused on any particular scale. Realized functions are naturalistic, and this naturalism extends to a version of teleosemantics that makes use of them.

6. Conclusion

In this chapter, I’ve argued that the problem of indeterminacy is better understood as a set of problems rather than a singular one. We might worry about the problem of content matching or of functional indeterminacy. It’s important that we keep these different problems separate because each of them can be solved in a different way. While the problem of content matching requires that we decide on a content as specific as we get in natural languages, the problem of functional indeterminacy does not. It can instead be solved with less specific content ascriptions. It does turn out, though, that no singular function can do all of the right work, because we need biological functions to play different explanatory roles in evolutionary biology and in ecology. No SE function can do all of the relevant ecological work, so we need to posit a realized function alongside the SE function; the teleosemanticist needs to be a pluralist about biological function and so about content.

All of this leaves open, though, the *other* problem. Here, I want to sketch out the start of a response to the problem of content matching that uses realized functions. My hope is this: because realized functions are sensitive not to historical circumstances, but to present ones, and because the realized niche can be described in a relatively fine-grained way if our various interactions with our environment are similarly fine-grained, a realized function is more likely to deliver a content specific enough that it could match onto a representation in natural language. If our behaviors differentiate between nomological covariants, then so, too, will those behaviors’ realized functions.

This is nothing like a proof; it’s closer to a prayer, if anything. I’ll explore the possibility in greater detail, though, in the next chapter. By considering how we might ascribe realized functions to Swampman, I’ll show how at least some of those functions deliver fine-grained content ascriptions.

Chapter 3: Swampman and the Problem of Novelty

1. Introduction

All backward-looking teleological theories of content share an objection: Swampman. Imagine that Donald Davidson is walking through a swamp when he is completely disintegrated by lightning that strikes the stump next to him (Davidson 1987). By complete coincidence, the lightning strike also reorganizes the matter of that stump into a perfect physical replica of Davidson the moment before his disintegration. This new being, Swampman, continues – “continues” – Davidson’s walk. He heads “home,” eats dinner, loves “his family,” writes papers, and so on and so forth. In short, Swampman behaves exactly as if he has all the same representational states as Davidson did and would have had, and we have good reasons to want to endorse this claim. After all, we can predict and, seemingly, explain Swampman’s behavior by making all sorts of claims about his representational states. The relevant behaviors can be as simple or as complicated as you’d like. He goes to the fridge to get a sandwich because he’s hungry and believes a sandwich is in the fridge. He defends a Tarskian theory of truth because he believes it’s true.

Yet a backward-looking teleosemantics will have to deny that any of these claims can be properly made of Swampman, since he has no history to refer to. While Swampman appears to have representational states, he in fact has nothing of the sort,¹⁹ because he has no *functions* that could ground the contents of those representations. This, of course, conflicts with both our intuitions and, seemingly, with our actual explanatory practices. Outside of evolutionary psychology, you don’t often find a working psychologist appealing to evolutionary history in order to explain or predict a behavior, especially when the representations appealed to concern things that are quite distant from anything in our environment of evolutionary adaptation (e.g., a telephone). Teleosemantics, then, faces a problem. They must explain why our psychological explanations for his behavior would work so well if they were utterly unfounded.

This isn’t, though, a new problem for us. It is, instead, another variant on a common theme: teleosemantics and the backward-looking notions of function that typically underpin it have a problem with novelty. In fact, they have the problem *of* novelty: how can we understand evolutionarily novel functions and, by extension, evolutionarily novel contents? We’ve seen this

¹⁹ At least, this is true in the moment he comes into existence. Some teleosemanticists can say that he will come to have representational states by acquiring functions in ontogenetic time. More on this later

play out elsewhere already. Looking to history wasn't enough for the ecological niche models used by conservation biologists because that history was not enough to determine how a population would fare under novel environment conditions. Similarly for Fodor's frogs, a novel environment in which the link between nomological covariants broke lead to an inability to answer a question about the frog's ecology – how did their visual system help to explain their realized niche under the novel conditions? – using only their evolutionary history. Swampman poses the same problem again, but at its most extreme: it seems like we ought to be able to explain things about him by an appeal to his representational states, *but everything is novel for him*, so our explanation for how he has those representations in the first place cannot ever appeal to history. Considering Swampman gives us a chance to ask how the teleosemanticist can best answer the problem of novelty in its most crystalized form.

I propose that, as in past cases, our best solution is to say that, while Swampman really does lack the representations theorized about by the backward-looking teleosemantics, he does have different sorts of representations: representations with teleosemantic contents in virtue of their realized functions. Seeing as that's a mouthful without a good acronym, though, and in the spirit of keeping my terminology parallel with realized niches and functions, I'll call these "realized representations." Embracing pluralism about function and representation helps us to more fully explain Swampman's behaviors, and, along with them, our own. I argue for this in four steps. First, in §2, I present the problem of Swampman in more detail. Rather than using him as an appeal to our intuitions, however, I instead show his importance as a test case for a theory of representation in the context of ecology. I pinpoint some particular representations that he must have in order for us to explain and predict certain ecological facts about him. In §3, I consider Papineau's response, which I take to be representative of many teleosemanticists' replies, in which he argues that, contra my position in the last chapter, representations as we find them in the actual world are inherently historical. As such, the backward-looking teleosemanticist does not need to explain how Swampman could represent, even supposing he does, because his representations would be different in kind from the ones the teleosemanticist is explaining. While he gets something importantly right, Papineau's response nevertheless falls short. No matter how we finagle things, SE functions just can't do the right explanatory work when we ask about the role of representations in Swampman's ecology. Then, in §4, I show that Swampman would have realized functions that would give him realized representations. He

won't have just the same set of representations as Davidson, but he'll have a set that both explains why so many of our explanatory and predictive practices work so well for him despite the significant differences between him and Davidson.

This solution is not just an answer to the challenge posed by Swampman, though. It is also an answer to the problem of novel contents more generally. I finish out §4 with a brief discussion of how the solution generalizes in this way. In §5, I consider perhaps the most developed backward-looking response in the literature from Garson and Papineau (2019). I show why the proposal nevertheless falls short. While extremely promising, this solution faces problems similar to those that I've identified elsewhere for backward-looking notions of function. This isn't to say that it *fails*, though. It's just that, as I've said before, it is *incomplete*. Finally, in §6, I conclude with some general remarks on where I take the debate over novel contents in teleosemantics to stand in light of the preceding discussion, including a discussion of a variant of Swampman that my proposal cannot explain.

2. A Dilemma

Let's get started. Suppose Swampman – Davidson's duplicate – really did pop into existence. He's a perfect physical duplicate, and therefore carries on exactly as Davidson would have. His next step will be the same one that Davidson would have taken. Supposing Davidson was in the middle of a thought, Swampman will "finish" it. Why, though, would we think that he has any representational states to begin with?

There are at least three routes²⁰ an argument could take here. First, we might appeal to our folk-psychological intuitions about the case. If we encountered Swampman in the wild, we'd assume he had the same cognitive apparatus as any human; we'd take his psychological life for granted. Supposing our intuitions are well-founded, this gives us an immediate reason to think Swampman proves something is going wrong with teleosemantics. I won't pursue this argument, though. For one, I think it is unclear whether we *should* trust our intuitions in a case like this. Swampman is so far from our daily lives that it wouldn't surprise me to learn that our intuitions about him, which were developed in a very different context, were deeply misleading. More importantly than this, though, it's unclear whether a teleosemanticist needs to take our intuitions here seriously given the nature of the theory for reasons that I will discuss in §3. To convince a

²⁰ I am only considering naturalistic explanations for Swampman's representational states here, since my most relevant interlocutors in this debate are themselves naturalists, or at least hope to be.

teleosemanticist that they need to take Swampman seriously, we'll need to beat them at their own game. Second, if we believe that representational states arise from something other than a trait's functions (i.e. if we aren't teleosemanticists), then Swampman would have the same representational states as Davidson in virtue of having just the same nonfunctional states, whatever those might be. I'll ignore this possibility here, both because I am a teleosemanticist myself and because I mean to address my fellow teleosemanticists in this chapter. I'll focus instead, then, on a third argumentative thread: our explanations and predictions about Swampman's behaviors necessarily depend on appeals to Swampman's representational states, so our theory of representation had better say he's got them. To make this point more clearly, let's consider a dilemma (Sebastián 2017; but cf. Dennett 1995).

If Swampman existed, then we would presumably be able to explain all of his behaviors by appealing to the sciences. After all, as another part of the causal order, there's nothing particularly mysterious or nonnatural about him; it's just that his appearance in the world is random. There are at least three questions we ought to ask here. First, just which facts about Swampman would we be trying to explain and predict? Second, when would those explanations and predictions carry over to Davidson? Third, just which sciences will we need to appeal to in order to do all of that explaining and predicting? As my reader might expect, I'll focus here on broadly biological questions. On the one hand, we might ask what I called in the introduction an evolutionary question, but that we might pose more generally as a historical one: how did Swampman come to have the particular set of traits that he has? Most²¹ of our answers on this front will be simply that he has his traits by chance. After all, his coming into being was itself a chance event. As such, our answer to this question will come apart radically from our answer with respect to Davidson, whose traits are explicable in terms of his evolutionary and ontogenetic history.

On the other hand, we might ask what I called an ecological question: how is it that Swampman manages to successfully occupy his environment?²² Consider, for instance, the following scenario. Swampman approaches the crosswalk of a crowded intersection on foot. Cars are moving through the intersection at a rapid pace. Swampman stands still at the crosswalk until

²¹ Some of his more plastic traits will likely be explicable by an appeal to history once enough time has passed, but the vast majority will not be so explicable.

²² We can take for granted that he would have, since he would behave just as Davidson would have, and Davidson would have carried on successfully occupying his environment!

a crosswalk sign lights up green, at which point he crosses the street with a number of other pedestrians, making it to the other side of the street. We then ask: how is it that Swampman managed to safely navigate the intersection? As we have also already seen, our answer to this need not appeal to Swampman's history at all. Ecology is different from evolutionary biology, and would require that we seek out answers even if evolutionary theory were false (Tinbergen 1963).

The teleosemanticist here runs into a dilemma (Sebastián 2017; but cf. Dennett 1995). The first horn goes like this. If the teleosemanticist is right – Swampman has no mental representations to speak of – then we won't need to appeal to any representations or representational contents to fully explain and predict his behaviors, including in situations like the one described above.²³ If we can explain and predict all of Swampman's behaviors purely by an appeal to nonrepresentational states, though, then, given that Davidson is just as much Swampman's physical duplicate as Swampman is Davidson's, it ought to follow that we can explain most²⁴ of his behaviors in exactly the same manner as we explained Swampman's. That is, we could explain Davidson's behaviors without ever appealing to representational states.

But this can't be right. First, teleosemantics is meant to explain representational states. If the teleosemanticist's position on Swampman requires that they give up on those states as useful theoretical posits for explaining and predicting such a wide range of behaviors, then they also undermine the motivation for their own position. Second, if we had been watching Davidson crossing the street instead of Swampman, we *would* have appealed to all sorts of representational states to explain the behavior. Davidson behaved the way he did because he meant to cross the road safely. He saw that there was a great deal of traffic moving through the intersection and knew it would be dangerous to cross while that was going on. He waited for the crosswalk sign to turn green because he knew that it represented that it was now safe for him to cross. This explanation is riddled with talk of representations, correctness conditions, and so on. The very structure of the explanation supposes it's possible for Davidson to be in a situation in which it makes sense to talk about him succeeding or failing. The function-talk is, seemingly, ineliminable, for reasons similar in kind to those I gave for functions being necessary to Garcia et al.'s model. Because our description of the situation relies on our framing Davidson's actions

²³ In fact, a teleosemanticist relying on SE functions wouldn't be able to appeal to much of any biology in general either, since Swampman wouldn't have any biological functions to speak of.

²⁴ I would never claim *all*, for reasons that will become apparent below.

in terms of their success or failure, we must suppose that his various behaviors are *meant to* lead to certain outcomes rather than others.

This pushes the teleosemanticist toward the second horn: Swampman really does have representational states, because we can only explain and predict *Davidson's* behaviors by appealing to those same representational states. The backward-looking teleosemanticist, though, can't explain how those representations work, because they evidently aren't the result of Swampman's history. Here are the rock and the hard place. Either the backward-looking teleosemanticist needs to deny that semantic properties play a role in answering ecological questions like the one posed above, or they need to admit that Swampman has some kind of representational content after all, and their theory of representation is incomplete. Perhaps there's a way to thread the needle, though. Let's turn now to Papineau's response to Swampman.

3. Papineau²⁵ on Representations

At the heart of Papineau's argument is a claim about what teleosemantics is meant to explain, namely, animals' representational systems and representations.²⁶ It has been expressly formulated as an *a posteriori* reduction of representation as we find it in the world, not as theory of representation however it could *in principle* exist. Moreover, the reduction teleosemantics offers is one that claims representations are inherently historical, because it's their shared history that explains their shared features²⁷ across tokens of the type. Where you find one of the features of a representing system, you typically find the others. "As with any natural kind, we can ask what grounds this clustering of properties. Why will any system that displays one typically display the others? And the answer seems obvious enough. Representation is a functional kind" (Papineau 2022, p. 14).

Now, part of what's at stake here is precisely whether representations are well understood as functional kinds as described by Papineau, so it's worth unpacking the obviousness here. First,

²⁵ I focus here on Papineau, but similar arguments have been run by Millikan (1984; 1996; 2001; 2010; 2017), Neander (1996), and Dennett (1996), among others.

²⁶ I should include a bit of a disclaimer here. As stated, Papineau's (2022) argument spends a lot of time on the metaphysics of natural kinds. I ignore that discussion here because the core of his argument does not depend on it so much as it does on an analysis of what teleosemantics is meant to explain in the first place.

²⁷ Papineau lists five of these shared properties (Papineau 2022, p. 13-14) and states that, where you find one, you typically find the other four. The particular properties won't matter either for his own argument or for my response to it. All that matters is that there is *some* cluster that could plausibly be called a natural kind. As such, I've omitted the properties from this text.

let's grant that animal representation really does hang together in the way Papineau describes.²⁸ There are many different representing systems that operate, at the right level of description, in the same way. Second, what, if anything, acts as the explanatory property explaining this hanging together? Animal representations do not all work in exactly the same way, and they do not share the same phylogeny. The ability to represent has evolved independently across the tree of life. As such, animal representations cannot be explained by an appeal either to some singular intrinsic property or by an appeal to shared ancestry. Why, then, do these properties tend to cluster? Because all of these different animals have faced the same selective pressures “to enable animals to tailor their behaviour to their circumstances” (ibid, p. 14). Hence, animal representation is a functional kind explained by its history of selection.

Here, then, is Papineau's answer to Swampman. The reason we can ignore the thought experiment is because, whatever Swampman would have going on in his head, it wouldn't be a part of the functional kind “animal representation.” Even supposing he did genuinely represent, the manner in which he did and our explanation for how he did would differ from our explanation for how animals do so, including ourselves. Supposing this is true – that our explanations really would be different in the way he discusses – Papineau would be able to comfortably live with the second horn of the dilemma raised above. Teleosemantics does not have to be a theory about every sort of possible representation, even supposing other sorts of representing systems than the ones we find in the actual world *are* possible. It only aims to explain how we work, and how organisms relevantly similar to us do.

²⁸ Though I'd still like to take a moment to note my dissension from the claim that selective pressures really are the explanatory property Papineau takes them to be, even for representations and representing systems that are explained by their history, broadly construed. I think they aren't; to understand the clustering of properties Papineau discusses, we would also need to discuss the dynamic interplay of a population with its environment. The development of an organism depends just as much on its interactions with its environment as it does on its inherited genotype (see Lewontin and Levins 1997; Oyama 2000; West-Eberhard 2003; Laland et al. 2011; Sultan 2015; Duckworth 2019). Moreover, it seems unlikely that selective-pressures are anywhere near as clearcut as I suspect Papineau needs them to be. Selection is an absolute mess, and it's unlikely that we could easily find selective pressures that explain the evolution of a representing system across as many contexts – across the whole kingdom Animalia! – that Papineau is imagining (Christie et al. 2021; Dewey forthcoming). Supposing this is right, though – and this is a big “supposing,” deserving at least a chapter in its own right to follow through on – it might still not follow that teleosemantics would need to explain Swampman's supposed representational states. Swampman would not, after all, have undergone any development; his phenotype would not have been the result of any interplay with the environment.

Papineau and other backward-looking teleosemanticists running this line get something deeply right, for two reasons. First, though I'll leave this mostly speculative here, it seems plausible that representing systems could be multiply realized not just at the level of the physical mechanism implementing them (e.g., neurons in humans, chips and logic gates in an AI), but also at the level of the mechanism discussed by the teleosemanticist.²⁹ It might be that something other than a history of selection could give rise to the ability to represent. At the least, in the absence of a negative argument, I don't see any *prima facie* reason to deny the possibility. Second, as I alluded to above when I mentioned the evolutionary question, it's true that Swampman differs from Davidson in ways that matter a great deal for our explanations about him. Many features of Swampman's representational system, supposing he has one, won't be explicable in the same terms as Davidson's. This is because there are significant biological differences between Swampman and the rest of us, namely, a difference in evolutionary history, and this difference matters for our explanations of Swampman's behaviors versus our own. For instance, while Swampman would appear fearful when he ran into a snake on his way out of the swamp, the explanation for that fact would only be able to rely on facts about his present state and his interactions with the environment around him. This same explanation would be woefully incomplete when applied to Davidson, though, since in his case it would need to include the evolutionary history of his nervous system to prompt fear-behaviors in the presence of snakes. As such, the teleosemanticist is right to insist that Swampman is *different*, and any theory trying to make sense of him will need to respect this fact.

Having said that, Swampman and Davidson nevertheless share many similarities with respect to ecology. To better understand this, let's start by considering Swampman alone. I may have appeared to take the following for granted above: we can properly describe Swampman as succeeding or failing to succeed at fitting into his environment. I should make good on the claim here, then. As I've discussed at length, we can ask questions about how an organism fits into its environment without ever discussing their history. Indeed, to extend the argument from Tinbergen (1963), an organism's ability to successfully occupy its environment would be a genuine subject of scientific study *even if they had no history to speak of*. There would still be something to explain. For Swampman, we would still need to explain why and how he was able

²⁹ I'll discuss this in more detail below, but I've also already done so. The last chapter, reframed, was an existence proof for pluralism about representations.

to successfully navigate his environment. More than this, and in parallel to the discussion of traits in ecological niche models in chapter one, our answers to ecological questions about Swampman would appeal to the ways in which his underlying physiological and behavioral traits “add up” to an overall description of his performance, and we would select those traits on the basis of their contributions to that performance. As we saw in chapter one, the telicity here is baked into the discussion in virtue of the explanatory role of Swampman’s traits when we are answering ecological questions. We cannot get around saying his traits are good or bad for things, and so we cannot get around assigning them functions.

Let’s now consider Swampman and Davidson together. Their identical physical properties would guarantee identical physiological and behavioral traits, and they would also interact with exactly the same environmental features. As such, our ecological explanations for the two of them would also be exactly the same. Whatever explanation we decided on for one of them would carry over to the other. With this point in hand, we can return to a modified version of the original dilemma. On the one hand, the teleosemanticist could deny that we need to appeal to representations in order to answer the ecological question. I take it that this option is a nonstarter. First, Davidson *does* represent. We need to use representations in order to explain how he can successfully occupy his environment. For one, we need to talk about his representational states in order to explain how he successfully navigates a busy intersection. Even setting that aside, though, if we need to appeal to representations in order to explain how Fodor’s frogs succeed or fail to occupy various environments, then we surely need to do so to explain the same for Davidson.

This brings us to the second horn of the dilemma: Swampman really does represent. Papineau was previously able to accept this horn because he could still say that those representations were explained differently in virtue of differences in history. That line of argument is no longer available to him, though, because there is no relevant difference between Swampman and Davidson that could ground any such difference in our explanations for their ecological successes and failures. The ecological explanations must be the same. Backward-looking teleosemantics, then, fails as a theory of content not because it doesn’t provide us with an explanation for how Swampman represents, but because it can’t provide us with an adequate account of how *we* do. It can’t tell us how the representations that feature in ecological explanations come to be.

There's more work to do, as there was with Fodor's frogs. I, of course, have a proposal. We can better respond to the problem of Swampman by appealing to his realized functions, which can explain how it is that he can have realized representations. This, in turn, explains how *we* have such representations. Teleosemantics can be a stronger theory if we just go in for pluralism.

4. Swampman's Realized Representations

We'll start again with the definition:

Realized function – Some trait x has the realized function f just in case (1) x allows an organism to f and (2) the ability to f helps to explain why some environment matches onto part of the organism's population's realized niche.

So, for Swampman to have realized functions at all, he'll need a couple of things: to be a member of a population; and to be a member of a population with a realized niche. I argue that both of these are true. In fact, Swampman could be well-understood as a member of many different populations,³⁰ some with corresponding realized niches. First, he might be considered the sole member of the population of Swamppeople. One might worry, though, whether such a population could be well-understood as having a realized niche. Being incapable of sexual reproduction, it seems reasonable to think it wouldn't, since there would be no sense in which the population could successfully occupy the environment indefinitely into the future.³¹ Fortunately for us, then, we can also count Swampman as a member of the *human* population. He therefore shares humanity's realized niche. He wouldn't be a member in virtue of species membership, of course – he has no species, at least considering the phylogenetic species concept (De Queiroz 2007)³² – but he could be considered one in virtue of the sorts of relations he would bear to other humans. On the less strenuous side of things, Swampman would interact with humans just like any human would. He would eat the same foods, take part in commerce, speak with others, have friendships and rivalries, and so on and so forth. Per the thought experiment, he would be so

³⁰ See chapter one, §4.2. for a discussion of the implications of one organism being a member of multiple populations simultaneously.

³¹ I don't explore the possibility in this dissertation, but I suspect one could construct a notion of function similar to the notion of a realized function that would rely only on claims about how an *individual* interacted with their environment that could do work here. I remain unsure, though, how one would decide naturalistically what counted as a success or failure. In §5, I consider the possibility of using Garson's (2017) generalized selected-effects theory of function to do this.

³² Though he would if we applied the biological species concept, being capable of interbreeding with humans! Though perhaps that is just another strike against it...

integrated into human society that it would be impossible for someone who didn't know about his origin to tell that he wasn't just another human. Any scientist studying him who was ignorant of his history would treat him as another member of the human population, and doing so wouldn't throw off their findings for most studies.

These all seem like good reasons to endorse the claim that Swampman could be well understood as a member of the human population, but we can also give a more rigorous analysis. Godfrey-Smith defines a Darwinian population as “a collection of particular things...that has the capacity to undergo evolution by natural selection” (2009, p. 6). The features required to be a part of a Darwinian population aren't historical, but mechanistic. The individuals need to be able to successfully reproduce, they need to vary, and there needs to be the possibility of random changes between generations that could lead to differences in the fitness of individuals. Swampman, being physically and behaviorally identical to Davidson, would have the requisite features to be able to be a part of the human Darwinian population that Davidson had been a member of. Moreover, because he would be a member of the human population, he would also have the same realized niche.

So far so good: Swampman has a realized niche – the same one as the human population – and so can also have realized functions. His traits will, in fact, have the same realized functions as everyone else. There are just two further points I ought to emphasize before getting into the meat of things. First, our discussion here doesn't depend so much on what his traits do that's good *for him* as it does on what his traits do that generally contribute to humanity's ability to occupy our realized niche. Our description of Swampman's realized functions will therefore be precisely the same as our description of our own, a result that is in line with my earlier argument that our answers to ecological questions concerning Swampman would be identical to our answers concerning Davidson. Second, as I showed in the first chapter, we don't need to know the precise mechanism underlying a realized function in order to say that some part of an organism has that function. Consider again the case of thermal tolerance. It's often enough to say that a population has a set of traits that give it some thermal tolerance as their realized function without having much of any clue about the identity of that set of traits (Garcia et al. 2014). Getting more mechanistic is often better, of course, since an explanation or prediction using data closer to the ground will be more accurate, but that sort of precision isn't strictly necessary.

To show that Swampman has realized representations, then, I need only show that *humans* have realized representations.³³ With that in mind, let's return to the case of an individual crossing a busy intersection. I have already argued in some small detail both that our best explanation for the success of such an individual involves their representing certain features of the environment and that our discussion of success conditions in the case is ineliminable. Let's make a quick detour back to ecology before getting more specific about the realized functions and representations relevant to the case.

Many arguments for representations being ineliminable come from an appeal to the lawlike generalizations they allow us to make about ourselves in psychology, which in turn enable us to keep track of systematic regularities in our behavior (e.g., Fodor 1974; 1997). Here I mean to add to those discussions with a focus on the importance of representations as theoretical posits that give us correctness conditions in ecology. As I argued in chapter one, we need to know how an organism's various traits allow it to interact with its environment in order to explain and predict its population's realized niche. Representations and representing systems are, as we've seen, just further traits we can appeal to in this effort. For instance, by saying that a frog's visual system is trying (or "trying") to pick out nutritious SDMs in its environment, we're able to support further claims about whether or not it is succeeding in that effort, and these in turn help us to explain and predict the frogs' realized niche. In the last chapter, I discussed how we must be able to decide on correctness conditions in order to attribute a content to a representation, but the reverse is also true. Success and failure are always relative to some activity. In order to say the frog is succeeding or failing when its visual system triggers the reflex to snap at a BB, we need to know that it is trying (or, again, "trying") to represent the BB as something in the first place. As Churchland (1986) puts it, the benefit of representing the world around us is that it allows us to coordinate our actions with the conditions we find in it. We think that Fodor's frogs are representing because their visual systems do just this. Something crossing the frog's visual field acts as a stimulus that triggers its reflex to snatch the thing out of the air, and, if things are going well – if the frogs are representing veridically – then they have coordinated that reflex with the presence of a nutritious SDM.

³³ I ought to also point out that the discussion here will concern the content of particular representations rather than our more general ability to represent. This focus helps to more particularly address the problem of attributing a content to a particular representation.

We can now return to Davidson at the crosswalk and run much the same line of argument. We want to know how it is that Davidson and humans more generally can successfully occupy their environment. (Notice again that the telicity here is baked-in. We cannot get around it if we are interested in explaining our success!) That environment includes a lot of fast-moving cars in spaces we occupy, and so we need to be able to coordinate our movements in such a way as to avoid those cars. We do, as a matter of fact, mostly succeed on this front, and this fact helps to explain why we have the realized niche we do. We would not, after all, survive in our own environments if we couldn't respond appropriately to cars at crosswalks. We can therefore posit a mechanism that causes certain behavioral changes in response to the presence of cars and other other things we find around crosswalks, since, in the absence of any such mechanism, our success would be a mystery. This mechanism, whatever it is,³⁴ has the realized function of keeping track of a condition in the environment – moving cars – in order to trigger certain avoidance behaviors appropriate to the situation in response to those cars, because doing so helps to to explain our realized niche. In other words, the mechanism has the realized function of representing moving cars because doing so helps to explain our ability to respond appropriately to moving cars in our environment, which in turn helps to explain why we have the realized niche we do.

Humans have realized functions that give some representations their content, and so Swampman, being a physical duplicate of a human, and so a member of the human population, would have those same realized representations for the same reasons. Swampman really can represent. We can even say this while still preserving the backward-looking teleosemanticist's insight that Swampman is importantly different from Davidson precisely because he lacks his evolutionary history. Many historically oriented functional and representational explanations that apply to Davidson fail with respect to Swampman. So, we have an answer to the problem of Swampman. More than that, though, we have at least part of an answer to the problem of novelty. The description above did two things on this front: first, it attributed the content "moving car" without ever appealing to history; and, second, it attributed an evolutionarily novel content – we didn't, after all, evolve any moving-car detectors. As such, an organism can have the realized function to represent something even brand new to them so long as doing so helps to

³⁴ Remember that realized functions can be assigned even if we don't know the particular trait or set of traits responsible for the performance of the function. See §3 of chapter one.

explain their realized niche. We have, in fact, seen this already with Fodor's frogs. Their visual systems could come to represent novel contents if their ability to do so helped to explain their new realized niche.

Swampman must have representations because we need to appeal to them to explain his ecological successes. Teleosemantics, then, must reckon with Swampman. I've argued here that it can do so via an appeal to his realized representations – indeed, by an appeal to *our* realized representations. The same argument applies to us with respect to the problem of novelty, and so can help to resolve the more general issue as well.

Perhaps, though, there's another way. As I've discussed at length, different patterns arise at different timescales, and these different patterns often require their own proprietary explanations. SE functions as I've discussed them won't solve the problem of novelty because there weren't busy intersections in our environment of evolutionary adaptation, but this might be more a problem with the scale discussed rather than with the theory of function itself. We need to reckon, then, with the possibility that an SE theory focusing on a different timescale could better resolve the issue. Of special interest is the possibility that selection processes at play over ontogenetic time could give rise to representations with novel contents, even for Swampman. There has actually been some work on this front recently. Let's look at it now.

5. A Selective-Effects Solution to the Problems of Swampman and Novelty?

Garson and Papineau (2019) have recently argued that an appeal to ontogenetic selection processes can be used to explain the acquisition of new representations over the lifetime of an individual. This is not a new argument by any means, having been proposed by Millikan in the same text in which she proposed biosemantics (1984). It is, however, a more developed and systematic answer to the question of how this might work. Whereas Millikan suggested that certain processes can lead to *derived* proper functions (i.e. some processes can have the function of producing new functional things), Garson and Papineau suggest that a modification of the SE theory of function can best do this work. The argument, to briefly summarize, runs like this. First, it endorses Garson's generalized selected-effects (GSE) theory of function,³⁵ according to which entities can be thought of as genuinely selected for not just if they are differentially reproduced, but also if they *differentially persist* (2017; 2019). For instance, consider a

³⁵ I am somewhat skeptical of this account of biological function, but am happy to admit it into the discourse here for the sake of seeing what sort of work it can do for the backward-looking teleosemanticist.

population of synapses connecting neurons in the brain. Over the course of development, some of those synapses persist whereas others disappear. Insofar as the synapses that persist do so because of some effect that they have while the other synapses disappear because they do not have this effect, we can say that the persisting synapses have been selected for their effect. If that effect was to represent something, then a neural mechanism can acquire the function to represent something. This process is called synapse selection, and it seems to be a biologically plausible explanation for a variety of neural phenomena (Schultz and Dickinson 2000; Garson 2012).³⁶

I won't go into mechanistic details here about just how the process is meant to work, though. Nothing I say will depend on the specifics on that front. Instead, I want to talk a bit about the explanatory structure at play and then to apply that structure to Swampman. GSE functions provide us with standard functional explanations as they have historically been understood by philosophers of science: they give us an explanation for the existence of a trait by an appeal to what that trait has done in the past (Hempel 1965; Garson 2019). It is because a set of synapses had the effect that they stuck around. Their function – that effect – explains their present existence.

This is a very promising start, for a few reasons. First, it seems clear enough that an ontogenetic selection process could, at least in principle, provide us with content attributions. If representing something helps to explain the persistence of a set of synapses, then their function would be to represent that thing. Second, while a theory appealing to the differential persistence of neural parts, whichever ones we might be talking about, still clearly relies on a discussion of a population, it can actually do away with the need to discuss a population *outside of* that individual. Unlike, for instance, my own notion of a realized function, a GSE function could arise without ever needing to identify a larger population of interest. It can therefore give a highly personal account of the content of representations. Third, a GSE function can be attributed without, at least in principle, having to say that representing some content is fitness-enhancing. So long as there is *some* selection-process at play, it doesn't matter whether that process is

³⁶ It's worth mentioning that this account of synapse selection is not, according to Garson and Papineau, meant to account for associative learning, which they say is more similar to something like the development of a derived proper function. I'll remain agnostic on this point, but would point towards Bucker's (2022) description of "tier 2" contents as a possible alternative description of the situation. In fact, I generally find Bucker's account of content more convincing than Garson and Papineau's, but this isn't particularly the place to run that argument given my intended audience.

selecting for some fitness effect. This could help the account deal with objections of the sort raised by Burge (2010) that *natural* selection doesn't care about truth. Fourth, someone advocating for this use of GSE functions could avoid the sorts of worries about telicity and the "proper" understanding of functional explanation that I had to ward off in chapter one. GSE functions get the telicity of SE functions, and they explain the present existence of a trait.

So let's see just how far GSE functions can get us with Swampman, since he is, after all, the most extreme variant of the problem of novelty. His neural system, being the same as Davidson's, would be subject to just the same selection pressures that Davidson's would have been. An ontogenetic selection process could therefore explain how Swampman could learn to represent new things.³⁷ If Davidson had somehow never learned about cotton candy, Swampman could come to represent it as he learned more about it over time. The same would be true for any other cases in which learning could explain his ability to represent something evolutionarily novel. Given that Swampman will presumably have some time before he reaches any busy intersections, it's conceivable that we could attribute the relevant representations necessary to explain that case by an appeal to this variety of ontogenetically acquired functions to represent.

Somewhat ironically, though, the account runs into problems when we consider the theoretically simpler case of a representation that *isn't* learned, and so not subject to differential selection. Think, for instance, of the fixation reflex. We seem to represent movements seen in our peripheral vision as things worthy of our attention – perhaps as possible threats – and orient ourselves towards such perceived possible threats as a result. The same will be true for Swampman. If he sees a sudden movement in his peripheral vision while at the crosswalk, his visual system will represent it as something deserving of his attention, and he'll be prompted to direct his attention to the moving object, just as Davidson would be. This fixation reflex, though, seems to be *built in*. Our cognitive system doesn't learn to represent movement in our peripheral vision as something we ought to pay attention to; it represents it as such because of the way our visual system is built. As such, there can be no process of selection to produce a representation that triggers said reflex.

We don't even need to posit innateness to get this sort of case off the ground, though. For any representations Davidson had that were no longer subject to selection because they no longer

³⁷ Supposing, of course, that this learning process was well-understood as a selection process, a point I'm happy to grant for the sake of this argument.

had any competitors – the relevant neural apparatus had already outcompeted its “conspecifics,” so to speak – there would be no ongoing selection process we could appeal to in order to attribute the relevant content to Swampman. The sorts of representations that we can attribute to Swampman are therefore seriously limited.

There is a possible response to this worry worth exploring, though. Wideman et al. (2023) have recently argued for the notion of a conserved function. These functions are attributed on the basis of, among other things, selection *against* alternative versions of the relevant traits in each of those lineages. This sort of negative selection has often been ignored by philosophers of biology. Perhaps the GSE theorist could appeal to selection against alternatives to explain how Swampman could come to represent in these cases. Neural systems are, after all, dynamic systems. In the absence of some causes preventing it from happening, we should expect Swampman’s neural mechanisms to change slightly over time. If we instead observe a continuity in these various neural systems, then we might be able to explain their persistence by an appeal to negative selection against alternatives that did not correctly represent. Swampman could therefore still come to represent things that Davidson had *already* learned to represent.

This story is convincing me to. We could just as easily appeal to negative selection against alternatives to explain the persistence of a neural mechanism as we could positive selection. It remains unclear to me, though, whether negative selection could explain the persistence of *innate* neural mechanisms. If these mechanisms don’t change over time, then there would be nothing to select against. That argument, though, hangs on empirical details that I am not prepared to explore here. In any case, there is a second problem I want to bring up that is significantly more threatening to this account of Swampman’s and our own novel representations. Remember that the GSE theorist does not just need to say that Swampman *has* representations. They need to also show that the representations underwritten by GSE functions can answer our ecological questions about Swampman. I argue that they can’t. We can’t use GSE functions to explain why a population has the realized niche that it does. To see why, think again about the general explanatory structure at play with GSE functions. For Garson and Papineau, a biological function is inherently historical: it explains the present existence of a trait by an appeal to what it has caused to occur in the past. This is to say that the facts that explain the function of a trait are what that trait has done in the past, not what it does in the present. In the case of synapse selection, a set of synapses have the function of representing something because doing

so explains their present existence via their differential persistence when compared to other sets of synapses.

It's not, however, just things that occurred in the past that explain the present realized niche of a population.³⁸ A population can't successfully occupy its present environment just because the ancestors of that population were able to. Its present members can do so because of what they can do *right now*. As I argued with Tinbergen (1963) in the introduction, ecology answers questions like this by an appeal to traits as causes of future effects, not as the effects of past causes. Insofar as answering the ecological question relies on an appeal to the effects of a trait rather than its causes, no account of function that relies only on its causes can fill this explanatory role. A trait's GSE function can only do so superficially and in virtue of matching the relevant trait's realized function. A GSE function to detect fast-moving things doesn't explain Swampman's present realized niche. Rather, his realized function to do so does.

Maybe, though, someone advocating for the use of GSE functions in this context could appeal to something a little different. We might still be looking at the wrong part of the timeline. Rather than saying that the GSE function explains the present realized niche of a population, perhaps we could instead say that a prediction about its *future* GSE function could do so. For instance, it isn't Swampman's synapses ability to detect fast-moving cars that explains his ability to occupy his realized niche; instead, it's that the synapses *will continue to have* that ability. Putting that a little differently, we can explain Swampman's and our own ability to indefinitely persist in our present environment by an appeal to the contributions of his and our traits towards our fitness going into the future. Notice, though, that we're now talking about Swampman's present traits as causes of future effects. *We are, in essence, talking about their realized functions*. We're interested in the contributions of those traits to the population's ability to persist indefinitely in its present environment. Moreover, we likely have to look for present realized functions in order to predict the future GSE functions of interest. Think again of Haeckel's definition of ecology discussed in the introduction: "in a word, ecology is the study of all those complex interrelations referred to by Darwin as the conditions of the struggle for existence" (Allee et al. 1949). To know what will be selected for over time, we need to know what good various traits do already. Realized functions become SE functions – and GSE functions – over

³⁸ Of course, as I've said before, some things in the past – past movements of the population, for instance (e.g., Wiens et al. 2019) – are necessary to get a *complete* explanation, but they still won't be enough on their own.

time, and our predictions about the process depend on our knowing the realized functions already. It's still the realized functions of the relevant traits, then, that can properly be said to answer ecological questions. GSE functions, including predictions about them, do so only derivatively.³⁹

We cannot replace realized functions with GSE functions. While GSE functions are useful tools for explaining how Swampman and ourselves can develop novel representations over time, they still can't provide us with adequate answers to ecological questions. While impressively dynamic, they remain as stuck in the past as their standard SE counterparts. As in previous cases, though, this doesn't tell us that GSE functions don't have a role to play here. We likely still need them to make sense of many backward-looking functional explanations about ourselves. The problem is that there are more functional explanations than these, and so we need another notion of function to make sense of *them* – the notion of a realized function. To best answer the problem of novelty as teleosemanticists, we must be pluralists.

6. Conclusion

In this chapter, I've argued that, contrary to the standard teleosemanticist response to Swampman, our theory of content does need to be able to attribute contents to his representational states. It must do so because many explanations for Swampman's realized niche would depend on appeals to his ability to represent, and those explanations would be identical to our explanations for our own ecological successes. Fortunately, we can get these explanations off the ground by appealing to Swampman's realized functions. In the process of showing how we could discover these functions, I also argued that the same moves could be used to solve the more general problem of novelty. Because a trait's realized function can change with its environment and because new realized functions can arise with new environmental conditions, we can naturalistically attribute evolutionarily novel contents to ourselves by an appeal to those functions. Finally, I considered a recent backward-looking response to the general problem of novelty from Garson and Papineau (2019) and argued that it still falls short.

It's not just a purely backward-looking teleosemantics that does so, though. Consider another version of Swampman, *Swampmartian*. Like Swampman, Swampmartian is behaviorally identical to Davidson and is indistinguishable from him from the perspective of those interacting

³⁹ This actually leads to a further and – for us – older problem. If we can't adequately answer ecological questions about a population by an appeal to their GSE functions, then we cannot adequately answer the problem of functional indeterminacy, either

with him. Like Swampman, he will carry on just as Davidson would have, and, as a result, many of his behaviors will only be explicable by an appeal to his representational states. To explain how he successfully navigates a busy intersection, we'll need to talk about the role of representations in his cognitive system.

Whereas Swampman was a physical duplicate, though, Swampmartian's similarity is only skin deep. His internal structure differs radically from our own. Of special importance is that he is incapable of reproducing with a human, even in principle. This isn't caused by a "problem" with his reproductive system, but by an incompatibility between it and our own. He differs significantly enough from us that he couldn't produce viable offspring with a human. As a result, Swampmartian can't be considered a member of the human population. He can be thought of only as a population of one. As I discussed in §4, this means that he doesn't have a realized niche, since his population of one couldn't persist in the environment indefinitely. This, in turn, means that he has no realized functions to give content to his representations. We cannot, then, use realized representations to explain how he navigates a busy intersection.

This problem, exemplified by Swampmartian, is not just some fringe case. It is, instead, an especially dramatic instance of a general phenomenon. In a warming world, many populations are losing their realized niches. Climatic changes are making the world uninhabitable for them. As a result, they have no realized functions. They do all sorts of things that a biologist would still say were functional – they have strategies for acquiring food, avoiding predators, and so on and so forth – but my account of realized functions can't make sense of those claims. This is the case by design. The ecological niche models I discussed in chapter one use realized functions as a means for predicting a population's range. If that population is destined for extinction, then the models ought to say that the traits in the population have no realized functions, since doing so explains why no environments will match onto their realized niche in the future.

GSE functions fare comparatively better with Swampmartian and his real-world equivalents. Since he has some underlying mechanisms that behave in the same ways as our neural systems do, those mechanisms are subject to the same sorts of selection processes as ours. They can come to represent things in virtue of their acquired GSE functions. This doesn't solve everything for Swampmartian, though, or for a member of a population without a realized niche – GSE functions would still have the same limits that I identified above – but the case illustrates the importance of *pluralism* for my account. That realized functions can't do all of this work isn't

a reason to reject them, because they were never meant to do it. They have a particular domain in which they're applicable, just like SE and GSE functions do. I've adopted pluralism about biological function for precisely this reason. We need all of these different notions of function – and, evidently, more besides, given the limits of both realized and GSE functions when applied to Swampmartian and similar cases! – to explain and predict things about phenomena we find in the world around us.

What sort of pluralism is this, though, and how do its various parts connect up? Let's now move on to the conclusion of this dissertation, in which I'll address these issues in more detail.

Conclusion: Other Games in Other Towns

The central thesis driving much of this dissertation is that different biological sciences are, well, *different*. They study different subjects and have different explanatory and predictive practices. As a result, our theorizing in the philosophy of biology and related fields must be sensitive to the distinctions between these sciences. We cannot assume that a theory about one of them will work for another, even if it seems to us that one of them is more central or basic.

In the age of the philosophy of the special sciences, it might seem as though this lesson doesn't need to be reiterated – we've all read Fodor by now – but, here we are. Garson has recently argued that only the SE theory of function can adequately account for the use of function-talk in biology and for functions' explanatory power. It is, he says, “the only game in town” (2019, p. 62). In this dissertation, I've shown that he is mistaken. While SE functions are indispensable for certain explanatory projects in evolutionary biology, there are nevertheless other games in other towns. Biology is not monolithic, and telicity does not arise exclusively as a result of a certain kind of history. Instead, we find other notions of function in use in other biological sciences. I've focused here on the use of functions in ecology, especially conservation biology. Some ecological niche models implicitly rely on what I've called realized functions, which are the effects of traits that help to explain a population's realized niche. We need these realized functions to explain how it is that organisms can be rightly described as succeeding or failing to fit into their environments even when they have not evolved to do so. Despite Garson's and others' (e.g., Mossio et al. 2009) recent attempts to show otherwise, we must still be pluralists about biological function.

What *sort* of pluralist ought we to be, though? There have been a number of efforts to make sense of how the various notions of function at play in biology might be connected. Perhaps there is a sort of hierarchy, each notion a special case of a broader category, such that the apparent pluralism turns out to be monism in disguise instead (Walsh and Ariew 1996; Saborido 2014). Alternatively, we might be disjunctive pluralists and hold that a variety of notions of function show up in a variety of biological sciences, and never the twain – or however many there turn out to be⁴⁰ – shall meet (Godfrey-Smith 1993). SE functions have their place in

⁴⁰ I'd rather remain coy on this point, but would likely admit at least causal role (Cummins 1975) and GSE functions (Garson 2017) to the pantheon in addition to those I've explicitly discussed.

evolutionary biology; realized functions have their own in ecology. There's little to nothing we can say, though, to connect them up or to make sense of their places in the natural order.

To get a better handle on the situation, let's look again at the practices of biologists. Different biological sciences study different things, but they are by no means independent of one another. We've seen this already in the discussion of ecology and evolutionary biology in the introduction. To quote Haeckel one last time:

By ecology we mean the body of knowledge concerning the economy of nature – the investigation of the total relations of the [organism] both to its inorganic and to its organic environment; including above all, its friendly and inimical relations with those [other organisms] with which it comes directly or indirectly into contact – in a word, ecology is the study of all those complex interrelations referred to by Darwin as the conditions of the struggle for existence (Allee et al. 1949).

Ecology and evolutionary biology are intimately intertwined, and each of them is, in a sense, directed by the other. I've already discussed how studies of the evolution of a trait often begin with an observation about their present survival value (see Davies, Krebs, and West 2012). Let's now consider a quick case study going in reverse – from an observation of selection to survival value – to better appreciate this point.⁴¹ Suppose that, using molecular methods, we discover a history of selection on a segment of a bacterium's genome. We don't yet know what this bit of the genome does, though, so we use a knockout method to study the effect of its absence on the organism. We find that it was responsible for the production of an amino acid by the bacterium. It can no longer produce the amino acid on its own, and when placed in an environment without the amino acid, the bacterium dies. We then hypothesize that the gene was selected for its role in the production of that amino acid.

In this vignette – which is, I take it, representative of a great deal of real-world work – we can see how ecology and evolutionary biology are both necessary for our understanding of the system. The study begins with discovering a fact about the evolutionary history of the bacterium – one of its genes has been selected for over time – and then investigates the role played by the relevant gene in order to understand why it was selected for in the first place. More than just the

⁴¹ Again, my thanks to Renee Duckworth for our conversations about what, exactly, we can directly observe about a biological system and how those observations play into both evolutionary and ecological explanations.

generic necessity of these two sciences, though, we have here an example of the necessity of multiple notions of *function*. We suspect that the gene in question has some SE function because we know it has been selected for, but we don't yet know what that function is. In order to discover it, we look for what good the gene does for the organism in the present, its *realized* function. We assume that its realized function, the production of an amino acid, is the same thing it was selected for in the past, and so conclude that the gene's SE function is to produce that amino acid. The gene's present existence is explained by an appeal to this historical effect.

It's because of the joint necessity of these sciences and their respective notions of function that I advocate here for the same sort of integrative pluralism about function developed by Cusimano and Sterner (2019). While there may not be a clear metaphysical relation between these various notions of function, they are nevertheless *practically integrated* in our scientific practices. We use a combination of them in the same study to more fully understand the system we're investigating. Putting this a bit more technically and to offer something like a definition, integrative pluralism about some set of notions of function claims that:

- (1) the various notions of function are at least sometimes all applied to the same system;
- (2) the system in question can at best be only partially understood by the application of a single one of these functions; and
- (3) we do not, at present, have the capacity to unify the notions of function at play.

(1) and (2) distinguish this sort of pluralism from *disjunctive* pluralism. The notions of function do not stay isolated in their own sciences. (3) distinguishes it from any sort of hierarchical monism. If we cannot reduce the notions to one another or to some more fundamental sort of function, then there's no sensible hierarchical structure between them to describe.

I am no disjunctive pluralist. I've already shown that SE and realized notions of function live up to (1) and (2) – we had to make use of both of them to better understand the bacterium in the toy case above – but there are two other cases from cognitive science that I ought to flag as well, namely, Fodor's frogs and novel contents. The central arguments of each of those chapters were not that we needed to *replace* SE functions with realized ones, but that *we needed both of them* to better understand the subjects of interest. Backward-looking teleosemantics ran into trouble because it considered only one of these, and I showed how we might better respond to the

objections by embracing pluralism. Our best explanation for how and what Fodor’s frogs’ visual systems represent relies on claims about both their SE and realized functions because their visual systems represent more than one thing at once. The same is true for our own representational systems. We can only understand our evolved representational capacities by an appeal to their evolutionary history, and we can only understand their roles in explaining our present realized niche by an appeal to their realized functions. A full understanding of our own cognitive systems requires that we make use of all of these notions of function at the same time. I have, it turns out, been advocating for integrative pluralism about function all along.

What, though, about (3)? I’ve already argued that realized functions are not reducible to SE functions in the previous chapters, but maybe we can go the other way around.⁴² SE and realized functions – along with evolutionary biology and ecology – can be thought of as, in part, looking at the same system from different temporal angles. A trait’s realized function will eventually⁴³ become its SE function. Perhaps, then, we could do the same work that SE functions do by an appeal to a trait’s *historical* realized functions. Consider again the bacterium. At some point in the past, its ancestors acquired the ability to produce the amino acid thanks to the gene of interest. At that point, it became a realized function of the gene to help produce that amino acid, because its ability to do so helped to explain why the relevant population had the realized niche that it did. This in turn helped to explain the population’s persistence, and so the continued existence of the trait into the present day.⁴⁴ Under this description, it seems as though we have an explanation that doesn’t, at any time, appeal to an SE function.⁴⁵ We appear to have reduced the SE function to a realized one stretched across the timeline.

But let’s take a step back to consider a point from Wimsatt (2007). Reductions do not, in all cases, make the reduced object “disappear” from our ontology. Rather, many mechanistic reductions are better understood as explanations for how the higher-level thing works. In explaining how it works, though, we may not have actually gotten rid of it. So long as certain

⁴² My thanks to Luke Goleman for initially pushing an argument like this in conversation.

⁴³ Exactly how long this will take differs between particular accounts of SE function, though it’s worth noting that no account, to the best of my knowledge, gives any definitive timeline. While Godfrey-Smith’s (1994) modern history theory of function says that recent selection is most important for understanding the function of a trait, the exact timeline for the acquisition of such a function remains ambiguous.

⁴⁴ See §2 of chapter one for a more in-depth discussion of and argument for this point.

⁴⁵ Someone impressed by Christie et al.’s (2021) argument against SE functions as explanations for the present existence of a trait might be especially interested in this line of argument.

properties of the system only come into focus once we're talking about all of the parts and their interactions, those properties will still remain attributable primarily to the higher-level thing and not to its parts alone.

I propose that something like this is at play with realized and SE functions. Our explanation for the present existence of a trait in a population only comes into focus when we consider, to again call back to the introduction, the wider temporal window. That wider temporal window is itself made up of many narrower ones containing organisms interacting with their environments. It's only when we add them all up and consider the myriad and possibly changing effects of a trait through time that we have an explanation for that trait's present existence and distribution. The same would be true even if the role played by a trait were the same at every timestep. If we are interested in the role a trait has played in its own replication over time, then we can't very well get an explanation going that considers only the narrow temporal windows. SE functions do not disappear from a theory that endorses realized functions.

(3) – we can't, at present, unify the notions of function – is therefore satisfied, too, and we ought to be integrative pluralists about SE and realized functions. The two notions are jointly necessary for making sense of the biological world. Pluralism about anything, though, always invites the question: is the identified set not just necessary, but also *sufficient*? In the last chapter, I argued that a set of cases in which an organism or population has no realized niche shows they are not. If we use only SE (or GSE) and realized functions, then the successes and failures of Swampmartian, as well as of the members of a population destined for extinction, remain mysterious. We cannot rely on realized functions when there isn't a realized niche, and SE and GSE functions will still struggle with cases of novelty. At the same time, I've argued that these limitations, properly understood, are both the strength of these notions and the motivation for pluralism in the first place. Realized functions shouldn't be attributed to traits in a population headed for extinction given the work they do in ecological niche models. SE and GSE functions are meant to be primarily concerned with the past because that's the subject of evolutionary biology.

Like the ecologists of the 19th century discovering and trying to make sense of the net of interrelations between myriad species and their environments, I have done my best to identify some of the kinds of telicity we find in the biological world, along with the ways in which they interact. But like them, too, I know that I am seeing only a small part of a larger and more

complex world. I believe that there remain as of yet undiscovered notions of function appropriate to other biological sciences, and I am confident that their discovery will yield further theoretical fruits, just as realized functions have in this dissertation. I leave to other researchers, though, the task of discovering these still other games in other towns.

Works Cited

- Allee, W. C., A. E. Emerson, T. Park, O. Park, and K. P. Schmidt. 1949. *Principles of Animal Ecology*. Saunders.
- Araujo, M.B., F. Ferri-Yañez, F. Bozinovic, P. A. Marquet, and F. Valladares. 2013. Heat freezes niche evolution. *Ecology Letters* 16: 1206–1219.
- Arnold, S. J. 1983. Morphology, performance and fitness. *American Zoologist* 23(2): 347-361.
- Buckner, C., 2022. A forward-looking theory of content. *Ergo an Open Access Journal of Philosophy* 8: 37. doi: <https://doi.org/10.3998/ergo.2238>
- Burge, T. 2010. *Origins of Objectivity*. Oxford University Press.
- Carscadden, K. A., Emery, E. C., Arnillas, C. A., Cadotte, M. W., Afkhami, M. E., Gravel, D., Livingstone, S. W., and Wiens, J. J. 2020. Niche breadth: causes and consequences for ecology, evolution, and conservation. *The Quarterly Review of Biology* 95(3): 179-214.
- Chesson, P., R. L. E. Gebauer, S. Schwinning, N. Huntly, K. Wiegand, S. K. M. Ernest, A. Sher, A. Novoplansky, and J. F. Weltzin. 2004. Resource pulses, species interactions and diversity maintenance in arid and semi-arid environments. *Oecologia* 141: 236 - 253.
- Chevin, L.-M., R. Lande, and G. M. Mace. 2010. Adaptation, plasticity, and extinction in a changing environment: towards a predictive theory. *PLoS Biology* 8: e1000357.
- Christie, J. R., C. Brusse, P. Bourrat, P. Takacs, and P. E. Griffiths. Forthcoming. Are biological traits explained by their selected effect function? *Australasian Philosophical Reviews*.
- Churchland, P. 1989. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. MIT Press.
- Cummins, R. 1975. Functional analysis. *The Journal of Philosophy* 72(20): 741-765.
- Davidson, D. 1987. Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association* 60(3): 441-458.
- Davies, N. B., Krebs, J. R., and West, S. A. 2012. *An Introduction to Behavioral Ecology, 4th ed.* Wiley-Blackwell.
- Dennett, D. 1995. The unimagined preposterousness of zombies. *Journal of Consciousness Studies* 2(4): 322-326.
- Dennett, D. 1996. Cow-sharks, magnets, and Swampman. *Mind and Language* 11(1): 76-77.
- De Queiroz, K. 2007. Species concepts and species delimitation. *Systemic Biology* 56(6): 879-886.

- Dewey, A. R. Forthcoming. Pressing Christie et al.'s (2023) objection: Why single out selected effects? *Australasian Philosophical Review*.
- Dreske, F. 1986. Misrepresentation. In R. Bogdan (ed.), *Belief: Form, Content, and Function*. Oxford University Press.
- Duckworth, R. 2019. Biological dynamics and evolutionary causation. In T. Uller & K. Laland (Eds.), *Evolutionary Causation: Biological and Philosophical Reflections* MIT Press. Pg. 153.
- Dussault, 2018. Functional ecology's non-selectionist understanding of function. *Stud Hist Philos Biol Biomed Sci*. 70: 1-9.
- Feeley, K.J., and M. R. Silman. 2010. Biotic attrition from tropical forests correcting for truncated temperature niches. *Global Change Biology* 16: 1830–1836.
- Feeley, K.J., and M. R. Silman. 2011. The data void in modeling current and future distributions of tropical species. *Global Change Biology* 17: 626–630.
- Fisher, J. and R. A. Hinde. 1949. The opening of milk bottles by birds. *Brit. Birds*, 42: 347-357.
- Foden, W.B., G. M. Mace, J.-C. Vie, A. Angulo, S. H. M. Butchart, L. DeVantier, H. T. Dublin, A. Gutsche, S. N. Stuart, and E. Turak. 2008. Species susceptibility to climate change impacts. *Wildlife in a changing world: an analysis of the 2008 IUCN Red List of Threatened Species* (ed. by J.-C. Vie, C. Hilton-Taylor and S.N. Stuart), pp. 77–88. IUCN, Gland, Switzerland.
- Foden, W.B., S. H. M. Butchart, S. N. Stuart, J.-C. Vie, H. R. Akcakaya, A. Angulo, L. M. DeVantier, A. Gutsche, E. Turak, L. Cao, S. D. Donner, V. Katariya, R. Bernard, R. A. Holland, A. F. Hughes, S. E. O'Hanlon, S. T. Garnett, C. H. Sekercioglu, and G. M. Mace. 2013. Identifying the world's most climate change vulnerable species: a systematic trait-based assessment of all birds, amphibians and corals. *PLoS ONE* 8: e65427.
- Fodor, J. A. 1974. Special sciences (or: the disunity of science as a working hypothesis). *Synthese* 28(2): 97-115.
- Fodor, J. A. 1984. Semantics, Wisconsin-style. *Synthese* 59(3): 231-250.
- Fodor, J. A. 1987. *Psychosemantics: the Problem of Meaning in the Philosophy of Mind*. Bradford Books.
- Fodor, J. A. 1992. *A Theory of Content and Other Essays*. MIT Press.

- Fodor, J. A. 1997. Special sciences: still autonomous after all these years. *Philosophical Perspectives* 11: 149-163.
- Futuyma, D. J. 2005. *Evolution*. Oxford University Press.
- Garcia, R. A., M. B. Araujo, N. D. Burgess, W. B. Foden, A. Gutsche, C. Rahbek, and M. Cabeza. 2014. Matching species traits to projected threats and opportunities from climate change. *Journal of Biogeography* 41: 724-735.
- Garcia, R. A., N. D. Burgess, M. Cabeza, C. Rahbek, and M. B. Araújo. 2012. Exploring consensus in 21st century projections of climatically suitable areas for African vertebrates. *Global Change Biology* 18(4): 1253-1269.
- Garson, J. 2012. Function, selection, and construction in the brain. *Synthese* 189: 451-481.
- Garson, J. 2017. A generalized selected effects theory of function. *Philosophy of Science* 84(3): 523-543.
- Garson, J. 2019. *What Biological Functions Are and Why They Matter*. Cambridge University Press.
- Garson, J. 2020. There are no ahistorical theories of function. *Philosophy of Science* 86(5): 1146-1156
- Garson, J., and D. Papineau. 2019. Teleosemantics, selection and novel contents. *Biology and Philosophy* 34(3): 36.
- Gibbs, H. L., and Grant, P. R. 1987. Oscillating selection on Darwin's finches. *Nature* 327: 511-513.
- Godfrey-Smith, P. 1993. Functions: consensus without unity. *Pacific Philosophical Quarterly* 74(3): 196-208.
- Godfrey-Smith, P. 1994. A modern history theory of function. *Noûs* 28: 344-362.
- Godfrey-Smith, P. 2009. *Darwinian Populations and Natural Selection*. Oxford University Press.
- Gonzalez-Suarez, M., P. M. Lucas, and E. Revilla. 2012. Biases in comparative analyses of extinction risk: mind the gap. *Journal of Animal Ecology* 81: 1211-1222.
- Gould, S. J., and R. Lewontin. 1979. The Spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society B: Biological Sciences* 205(1161): 581-598.
- Gould, S. J., and E. S. Vrba. 1982. Exaptation—a missing term in the science of form. *Paleobiology* 8(1): 4-15.

- Harper, J. L. 1982. After description. In E. I. Newman (ed.) *The Plant Community as a Working Mechanism*. Blackwell Scientific Publications.
- Hempel, C. G. 1965. *Aspects of Scientific Explanation*. Free Press.
- Holt, R. D. 2009. Bringing the Hutchinsonian niche into the 21st century: ecological and evolutionary perspectives. *Proceedings of the National Academy of Sciences*: 106(supplement 2): 19659-19665.
- Hutchinson, G. E. 1957. Concluding remarks. *Cold Spring Harb Symp Quant Biol* 22: 415-427.
- Hutchinson, G. E. 1965. *The Ecological Theater and the Evolutionary Play*. Yale University Press.
- IPCC. 2022. *Climate Change 2022: Impacts, Adaptation, and Vulnerability*. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [H.-O. Pörtner, D.C. Roberts, M. Tignor, E.S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, B. Rama (eds.)]. Cambridge University Press. Cambridge University Press, Cambridge, UK and New York, NY, USA, 3056 pp., doi:10.1017/9781009325844.
- Jezkova, T., and J. J. Wiens. 2016. Rates of change in climatic niches in plant and animal populations are much slower than projected climate change. *Proceedings of the Royal Society of London* 283:20162104.
- Kaufmann, S. A. 1970. Articulation of parts explanation in biology and the rational search for them. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1970: 257-272.
- Kingsland, S. E. 2005. *The Evolution of American Ecology, 1890-2000*. JHU Press.
- Krebs, C. J. 2005. *Ecology: the Experimental Analysis of Distribution and Abundance, 5th ed.* Benjamin Cummings.
- Laland K. N., K. Sterelny, J. Odling-Smee, W. Hoppitt, and T. Uller. 2011. Cause and effect in biology revisited: is Mayr's proximate-ultimate dichotomy still useful? *Science* 334(6062): 1512-6.
- Levin, S. A. 1992. The problem of pattern and scale in ecology: the Robert H. MacArthur Award lecture. *Ecology* 73(6): 1943-1967.
- Levins, R., and R. Lewontin. 1985. *The Dialectical Biologist*. Harvard University Press.

- Lloyd, E. 2015. Adaptationism and the logic of research questions: how to think clearly about evolutionary causes. *Biological Theory* 10(4): DOI: 10.1007/s13752-015-0214-2.
- Luck, G.W., S. Lavorel, S. McIntyre, and K. Lumb. 2012. Improving the application of vertebrate trait-based frameworks to the study of ecosystem services. *Journal of Animal Ecology* 81: 1065–1076.
- Lyons, J. 1995. *Linguistic Semantics: An Introduction*. Cambridge University Press.
- McGill, B. J. 2010. Ecology: matters of scale. *Science* 328: 575.
- McGinn, C. 1989. *Mental Content*. Blackwell Publishers.
- McIntosh, R. P. 1985. *The Background of Ecology: Concept and Theory*. Cambridge University Press.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. MIT Press.
- Millikan, R. G. 1990. Compare and contrast Dretske, Fodor, and Millikan on teleosemantics. *Philosophical Topics* 18(2): 151-161.
- Millikan, R. G. 1996. On swampkinds. *Mind and Language* 11(1): 103-117.
- Millikan, R. G. 2010. On knowing the meaning; with a coda on Swampman. *Mind* 119(417): 43-81.
- Millikan, R. G. 2017. *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford University Press.
- Millikan, R. G. 2023. Teleosemantics and the frogs. *Mind and Language* 39(1): 52-60.
- Millstein, R. L. 2020. Functions and functioning in Aldo Leopold's land ethic and in ecology. *Philosophy of Science* 87(5): 1107-1118
- Morrow, K. H. 2023. A causal-role account of ecological role functions. *Philosophy of Science* 90: 433–453.
- Mossio, M., C. Saborido, and A. Moreno. 2009. An organizational account of biological function. *The British Journal for the Philosophy of Science* 60(4): 813-841.
- Neander, K. 1983. Abnormal psychobiology: a thesis on the 'anti-psychiatry debate' and the relationship between psychology and biology. Thesis.
- Neander, K. 1996. Swampman meets Swampcow. *Mind and Language* 11(1): 118-129.
- Neander, K. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. MIT Press.
- Oyama, S. 1985. *The Ontogeny of Information*. Duke University Press.

- Papineau, D. 2022. Swampman, teleosemantics and kind essences. *Synthese* 200(6): 1-19.
- Parmesan, C., and G. Yohe. 2003. A globally coherent fingerprint of climate change impacts across natural systems. *Nature* 421(6918): 37-42.
- Potochnick, A., and B. J. McGill. 2022. The limitations of hierarchical organization. *Philosophy of Science* 79 (1):120-140.
- Preston, P. 1998. Why is a wing like a spoon? A pluralist theory of function. *Journal of Philosophy* 95(5): 215-254.
- Price, C. 1998. Determinate functions. *Noûs* 32(1): 54-75.
- Saborido, C. 2014. New directions in the philosophy of biology: a new taxonomy of functions. In Galavotti MC, Dieks D, Gonzalez WJ, Hartmann S, Uebel T, Weber M (eds) *New Directions in the Philosophy of Science*. Springer International Publishing, Cham, pp 235–251.
- Sebastián, M. 2017. Functions and mental representation: the theoretical role of representations and its real nature. *Phenomenology and Cognitive Science* 16: 317–336.
- Schultz, W., and A. Dickinson. 2000. Neuronal coding of prediction errors. *Annual Review of Neuroscience* 23: 473-500.
- Simon, H. A. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* 106(6): 467-482.
- Sterelny, K. 1990. *The Representational Theory of Mind: An Introduction*. Blackwell.
- Sterner, B., and S. Cusimano. 2019. Integrative pluralism for biological function. *Biology and Philosophy* 34(6): 1-21.
- Sultan, S. E. 2015. *Organism and Environment*. Oxford University Press.
- Tinbergen, N. 1963. On the aims and methods of ethology. *Zeitschrift für Tierpsychologie* 20: 410-433.
- Thuiller, W., S. Lavorel, and M. B. Araujo. 2005a. Niche properties and geographical extent as predictors of species sensitivity to climate change. *Global Ecology and Biogeography* 14: 347–357.
- Violle, C., M. L. Navas, D. Vile, E. Kazakou, C. Fortunel, I. Hummel, and E. Garnier. 2007. Let the concept of trait be functional! *Oikos* 116(5): 882-892.
- Walsh, D. M., and A. Ariew. 1996. A taxonomy of functions. *Can J Philos* 26(4):493–514
- West-Eberhard, M. J. 2003. *Developmental Plasticity and Evolution*. Oxford University Press.

- Wideman, J. G., S. Elliott, and B. Sterner. 2023. An account of conserved functions and how biologists use them to integrate cell and evolutionary biology. *Biology and Philosophy* 38(5): 1-23.
- Wiens, J. J. 2016. Climate-related local extinctions are already widespread among plant and animal species. *PLoS Biology* 14:e2001104.
- Wiens, J. J., A. Camacho, A. Goldberg, T. Jezkova, M. E. Kaplan, S. M. Lambert, E. C. Miller, J. W. Streicher, and R. L. Walls. 2019. Climate change, extinction, and Sky Island biogeography in a montane Lizard. *Molecular Ecology* 28: 2610-2624.
- Williams, J. E., & J. L. Blois. 2018. Range shifts in response to past and future climate change: Can climate velocities and species' dispersal capabilities explain variation in mammalian range shifts?. *Journal of Biogeography* 45(9): 2175-2189.
- Williams, S.E., L. P. Shoo, J. L. Isaac, A. A. Hoffmann, and G. Langham. 2008. Towards an integrated framework for assessing the vulnerability of species to climate change. *PLoS Biology* 6: e325
- Wimsatt, W. C. 1972. Complexity and organization. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1972: 67-86
- Wimsatt, W. C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press.
- Wing, S. L., G. J. Harrington, F. A. Smith, J. I. Bloch, D. M. Boyer, and K. H. Freeman. 2005. Transient floral change and rapid global warming at the Paleocene-Eocene boundary. *Science* 310(5750): 993-996.