Synthese (2021) 198:3499–3517
https://doi.org/10.1007/s11229-019-02293-8

Against etiological function accounts of interests

Katie McShane[1]

Received: 2 August 2018 / Accepted: 12 June 2019 / Published online: 27 June 2019

Abstract
The etiological account of function defines a part's/trait's function as whatever that part/trait does and was selected for doing. Some philosophers have tried to employ this as an account of biological interests, claiming that to benefit an organism is to promote its etiological functioning and to harm it is to inhibit such functioning. I argue that etiological functioning is not a good account of biological interests. I first describe the history of theories of biological interests, explaining the special role that etiological accounts of function have played within such theories. Second, I explain the problems with allowing etiological accounts of function to play this role and consider objections to my line of argument. Finally, I consider the theoretical alternatives to etiological function accounts of interests and assess their advantages and disadvantages.

Keywords Function · Biocentrism · Interests · Etiological function

1 Introduction

The etiological account of function, sometimes referred to as the "selected effects" account, explains the function of something's parts or traits in terms of their selection history.[1] This kind of account has played an important role in philosophical analyses of biological interests. Biological interests are interests that living things have merely in virtue of being alive, independently of other capacities such as sentience or consciousness. Early accounts of biological interests explained such interests in terms of etiological functioning. John Basl, whose recent book (Basl 2019) offers a more sophisticated account biological interests, also bases claims about such interests on facts about the selection history of their bearer's parts/traits. In this paper, I argue that

[1] Because some authors make these claims in terms of parts and others in terms of traits, I use "parts/traits" hereafter.

✉ Katie McShane
katie.mcshane@colostate.edu

[1] Philosophy Department, Colorado State University, Campus Box 1781, Fort Collins, CO 80523-1781, USA

this strategy is flawed. Whatever its merits as an account of function might be, the etiological account is a bad account of interests.

In what follows I briefly describe the history of theorizing about biological interests, explaining both the early views and Basl's more recent proposal. I then describe the problems with relying on etiological explanation to ground interests in either case. In light of these problems, I argue that etiological accounts of biological interests ought to be rejected. After considering possible objections to my criticisms, I discuss various alternative accounts and note and their advantages and disadvantages. The alternatives, I argue, involve controversial philosophical assumptions about the nature of interests. These assumptions stand in need of further explication and justification.

2 Historical development

In this section, I briefly summarize the development of etiological function accounts of interests, with special attention to Basl's recently developed view. Many of the lines of reasoning I describe here have been subjected to criticism in the literature; my explanations should not be read as endorsements. I explain the reasoning here so that the reader can see why these views took the form they did.

2.1 Biological interests

The idea that nonsentient, nonconscious organisms can be harmed or benefited has seemed intuitively plausible to many people, philosophers and nonphilosophers alike. Claims such as "Injecting poison into the root system is bad for the oak tree" or "It would be good for the ficus to get a bit more sunlight" are fairly common in ordinary discourse, and they don't seem to involve merely metaphorical uses of "good for" and "bad for."[2] However, as theories of human welfare have increasingly grounded welfare claims in facts about people's subjective states (facts about pleasures and pains, preferences, ends, etc.), it has become more difficult to explain how organisms without subjective states could nonetheless possess interests—i.e., have a welfare or a well-being, be subjects of harm or benefit.[3]

In the 1980′s and 1990′s, some philosophers took up the challenge of explaining what the interests of nonsentient, nonconscious organisms (for the sake of brevity, hereinafter simply "organisms") might consist in.[4] Most of these philosophers were biocentrists, theorists who ultimately wanted to argue for the moral importance of such interests. Showing that these interests exist at all, however, was their first task, and they faced a number of difficulties in carrying it out. One that loomed large in the early

[2] See Attfield (1981) for an argument against the claim that such uses are metaphorical.

[3] Here I use the terms "have interests," "have a welfare," and "have a well-being," "be capable of being harmed or benefited" interchangeably. Some authors use "have interests" to mean "take an interest in things," e.g., Taylor (1986, p. 63). I follow Regan (1983, p. 87) in distinguishing between the claims "A is interested in X" and "X is in A's interest." To have interests, as I use it here, refers to the latter usage.

[4] Central texts in these discussions are Johnson (1991), Goodpaster (1978), Attfield (1981), Rolston (1988), Taylor (1986), and Varner (1998). See also Cahen (1988), Thompson (1990), Nolt (2009), Holm (2012, 2017), Basl and Sandler (2013a, b), and Basl (2019) for critical discussions of this early literature.

literature was finding a basis for interest-attribution according to which organisms would have interests but nonliving things would not.[5] The concern seemed to be that if the biocentrists' preferred theory of welfare also attributed interests to cars, rocks, etc., then this would be considered a *reductio ad absurdum* of it as a theory of welfare.[6]

In order to address this issue, some theorists claimed that what distinguishes organisms from nonliving things is the fact that organisms have a good of their own, while nonliving things do not.[7] We can think of nonliving things as divided into two categories: nonartifacts (things not made by people: e.g., rocks, water, flames) and artifacts (things made by people: e.g., cars, pencil sharpeners, thermostats).[8] Nonliving nonartifacts, theorists claimed, do not have a good of their own because they do not have a good at all. If one steps on a rock and breaks it in half, they argued, one does not harm the rock; the rock is not worse off when fragmented than it is when whole. The rock doesn't have an interest in being in one piece rather than two; in fact, the rock does not have an interest in being in any particular state or another. Theorists explained this fact by noting that rocks are not "teleologically organized": they do not engage in goal-directed activity; their parts/traits do not have functions or purposes.[9] By contrast, organisms are teleologically organized. They do engage in goal-directed activity, and their parts/traits do have functions or purposes. A plant's vascular system, for example, has the biological function of circulating water, dissolved minerals, and sugars to the plant's various structures.[10] Stepping on the plant could damage its vascular system, causing the system to malfunction or to cease functioning altogether, in which case the plant itself would die. In the case of both the rock and the plant, stepping on it would alter its structure. However, in the case of the plant but not the rock, this alteration would interfere with the processes by which the plant keeps itself alive.

The distinction between things that are teleologically organized and things that are not does not help in the case of nonliving artifacts, however, since sophisticated machines can also be teleologically organized (Taylor 1986, pp. 123–125;

[5] Sune Holm (2017) refers to this as the "problem of scope."

[6] See, Varner (1990, p. 251) for one of the more explicit examples of this rationale.

[7] Taylor (1986, pp. 18, 123–124) and Rolston (1988, pp. 94–125) were the most explicit in drawing this distinction, though it is in the background of other accounts as well.

[8] While the distinction between things made by people and things not made by people correctly captures early biocentrists' interest in contrasting the artifactual with the natural, it is worth noting that "made by people" is not the same as "the product of intentional design." On the latter definition, beaver dams and bird nests might well count as artifacts, though not ones made by people. I use the former definition here for reasons of historical fidelity to the literature I am discussing.

[9] An important exception to the generalization that teleological organization was considered the important difference between living things and nonliving nonartifacts is Robin Attfield's view. Attfield (1981, p. 42) argued that the good of a plant is a matter of the fulfillment of its nature and that neither artifacts nor "things lacking inherited capacities" are capable of natural fulfillment. The nature of an organism, he claimed, is a matter of its "essential capacities," i.e., "capacities in the absence of which from most members of [its] species that species would not be the species of [that organism]." More recent "natural goodness" accounts (see, e.g., Thompson 2007; Hursthouse 1999; Foot 2001; for critical discussion see Odenbaugh 2017) bear some similarities to Attfield's view in this regard.

[10] See Taylor (1986, pp. 121–123) and Rolston (1988, pp. 96–104) for arguments of this type. For discussion, see Thompson (1990), Samuelsson (2010), Cahen (1988), and Holm (2012). As Varner (1990, pp. 256–257) points out, there are important differences between claims about goals and claims about functions.

Rolston 1988, pp. 104–106). A car's structures and processes are goal-directed, and its parts/traits also seem to have functions or purposes. The goal of a car's engine-lubrication system, for example, is to prevent the engine from overheating; the function of the oil pump within that system is to circulate oil around the engine's moving parts. Early biocentrists wanted to be able to claim that plants have interests but cars do not; this required them to show why damage to a plant's vascular system would count as a harm to the plant, but damage to a car's engine-lubrication system would not count as a harm to the car. Theorists addressed this matter by arguing that the car's good is not a good *of its own* (Taylor 1986, pp. 123–124; Rolston 1975, pp. 104–106; see also von Wright 1963, pp. 50–51). What is good for a car is just a matter of how humans want to use it or what state we would prefer it to be in. What is good for the plant, however, does not depend on human interests or intentions in this way. The claim was that this is true of all artifacts and all organisms. The good of an artifact is really the good of something else, or at least derived from the good of something else; it is not a good that the artifact has in its own right.[11] This distinguishes artifacts from organisms, which do have a good of their own.

The position of these early biocentrists, then, was that organisms but not nonliving things should be counted as possessing interests because organisms are teleologically organized, and their good is a good of their own. That which furthers an organism's good-of-its-own is good for it; that which frustrates an organism's good-of-its-own is bad for it. Still, an account of what an organism's good-of-its-own consists in was needed, and this is where accounts of biological function came in.

Theories of biological function emerging from philosophy of biology at the time offered biocentrists the hope of a naturalistically acceptable conception of biological function—i.e., an explanation of the "teleological organization" possessed by organisms in virtue of which they have not only a good, but a good of their own. The etiological account of function from Larry Wright seemed to be especially congenial to the biocentrists' project.

2.2 Etiological functions

When someone asks what is good for a plant, one reply we might plausibly give is that the plant's good consists in its parts/traits functioning properly—i.e., doing what they're supposed to do, working as they should. However, to say this requires an account of what "functioning properly" or "doing what they're supposed to do" amount to. How do we distinguish between functioning and malfunctioning, between parts/traits doing what they're supposed to do and doing something else? Etiological accounts of function offered answers to these questions. On an etiological account of function, "what the parts/traits are supposed to do" becomes "what they were [through the evolutionary process of natural selection] selected for doing." When the parts/traits do what they were selected for doing, they are functioning; when they do not, they are malfunctioning.

[11] See Varner (1998, p. 66) for a criticism of the claim that the goals of artifacts cannot be explained without reference to the goals, intentions, or purposes of their human designers. For further discussion, see also Holm (2017).

Etiological accounts of function were first introduced by Larry Wright in 1973. According to Wright,

"The function of X is Z means

(a) X is there because it does Z,

(b) Z is a consequence (or result) of X's being there" (Wright 1973, p. 161).

Wright intended this account of function to provide a unified explanation of both what he calls "conscious functions" (e.g., "The function of the sweep-second hand on a watch is to make seconds easier to read") and what he calls "natural functions" (e.g., "The function of the heart is pumping blood") (Wright 1973, pp. 142, 139). The explanation of "is there because" in criterion (a) will be different for conscious functions and natural functions. The sweep-second hand is there because it makes seconds easier to read in the sense that it was put there by the watch designer because it makes seconds easier to read. The heart is there because it pumps blood in the sense that it was selected for through the evolutionary process of natural selection because it pumps blood.

Wright's account of natural functions was of great interest to biocentrists. On Wright's account, plants and other organisms have parts/traits that have natural functions, and what those functions are is a matter of not only what they currently do (b), but also what their evolutionary history is (a). If natural functions could serve as the basis for interest-attributions, such that promoting this functioning would count as benefiting the organism and interfering with this functioning would count as harming the organism, then we would have an explanation of what the good of a plant, for example, consists in. It is an explanation that comports well with ordinary claims about what is good or bad for plants: most ways of harming plants do amount to preventing their parts/traits from performing their natural functions properly, i.e., from doing whatever they evolved to do. For example, when I step on the plant and damage its vascular system, I am preventing that system from doing what it was selected for doing, namely circulating water, dissolved minerals, and sugars. This account of functions is also a naturalistically legitimate one; it doesn't require one to appeal to divine intentions, attribute mental states to plants or otherwise anthropomorphize them, or posit forces outside of those well-established within evolutionary biology in order to explain what would count as functioning and malfunctioning. Finally, the distinction between conscious functions and natural functions offers hope for excluding artifacts as bearers of interests. If only *natural* functions are determinative of welfare, then only entities that have functions established through natural selection will count as bearers of interests. That claim might well seem plausible, since conscious functions seem to be a matter of the intentions of an object's designer rather than the welfare of the object designed. The view that artifacts do not have a good of their own thus seemed explicable in terms of how their traits/parts came to have the functions they have. The hope, then, was that we would end up with a theory according to which plants have a welfare—they can be harmed or benefited—but rocks and cars do not.[12]

[12] I deliberately omit discussion of Agar's (2001) version of biocentrism here. While it does rely on claims about etiological function to argue for a version of biocentrism, its reliance departs substantially from that

2.3 Basl's account

The rise of "synthetic biology" and the prospect of organisms that are entirely the product of human design and engineering has made trouble for some of these early biocentric theories.[13] In the relevant sense, synthetic organisms are artifacts: their parts/traits are what they are and do what they do because humans designed them that way.[14] As a result, their functions count as conscious functions, not natural functions. Attributing interests only to things possessing natural functions would mean that synthetic organisms could not have interests. If we had an evolved plant and a synthetic plant, both of which were morphologically and behaviorally identical in every respect, the former would possess interests while the latter would not (Basl and Sandler 2013b). If I were to step on the evolved plant and destroy its vascular system, doing so would harm the plant. If I were step on the synthetic plant and destroy its vascular system, doing so would not harm the plant. This implication has struck many theorists as counterintuitive (Basl and Sandler 2013a, b; Attfield 2012).

To address this problem, John Basl (writing together, in some places, with Ron Sandler) has proposed a new view, which amends earlier biocentric theories to include artifacts as bearers of interests (Basl 2019; Basl and Sandler 2013a, b; see also Basl 2012). Basl agrees that the interests of nonconscious, nonsentient things must be a matter of their teleological structure, understood etiologically. (Basl does not refer to "etiological functions," since he does not want to take a stand on the correct view of functions; he talks instead of etiological teleology.) Unlike earlier theorists, however, Basl argues that there is no reason to require that the etiology grounding interest-attributions involve *natural* selection. What Wright would call "conscious functions," Basl argues, can serve equally well to ground the attribution of interests to artifacts. This modification solves the problem posed by synthetic biology: both the synthetic organism and the evolved organism will count as having interests (both in virtue of having a selection history: via conscious selection in the former case and natural selection in the latter case). It achieves this result, however, at the cost of attributing interests to nonliving artifacts—the very result that earlier theorists were working to avoid. Basl argues that this is not a problem, since he argues further that the interests of artifacts

Footnote 12 continued
of other views. Agar claims that entities with intrinsic value are those possessing contentful representational states. He uses etiological functions as an account of the content of mental states, notes that etiological functions can also account for the content of the nonmental "biopreferences" possessed all by living things, and on this basis argues that living things should be considered intrinsically valuable. Because he uses etiological functions to explain the content of representational states, and then uses the possession of representational states to attribute intrinsic value to entities (rather than using etiological functions directly to describe the content of interests), Agar's view is not really an etiological function account of interests, but rather an etiological account of representation and thereby intrinsic value. It is, however, an important development in the literature on biocentrism, one that merits a separate consideration.

[13] There are further concerns about such theories beyond the ones I discuss here. See, e.g., Basl (2017, 2019) and McShane (2014) for criticisms related to the units of selection problem.

[14] I follow most writers on this issue (e.g., Basl and Sandler 2013a) in distinguishing between (a) humans affecting the characteristics of naturally evolved organisms, e.g., through selective breeding or genetic modification; and (b) humans designing and creating entirely new organisms. The latter is distinctive of synthetic biology. In the former case, one might still claim that the organism has parts/traits that are the product of natural selection. One cannot make that claim in the latter case. For a further discussion of this distinction and its ethical implications, see, e.g., Preston (2018), Attfield (2012), and Basl (2019).

have no moral importance (Basl and Sandler 2013a, p. 704; Basl 2019, chapter 4). Nevertheless, the attribution of interests to items such as cars, pencil sharpeners, and thermostats marks a significant departure from prior versions of biocentrism.

3 Criticisms

In evaluating etiological accounts of interests, it is helpful to begin with earlier accounts, which grounded interests only in *natural* etiological functions, since some of the problems with these accounts inform worries about Basl's more recent account. According to the earlier view, the biological interests of an organism consist in its parts/traits continuing to perform (or at least being able to perform) the functions that they were selected for performing. However, there are good reasons for thinking that this is not true, i.e., that it is not always good for an organism to have its parts/traits perform (or even be able to perform) the functions that they were selected for performing.

A part's/trait's behavior is selected for when it is adaptive, i.e., has a positive impact on fitness within an environment.[15] However, as I have noted elsewhere (McShane 2014), what worked well for an organism's ancestors might not work well for the organism itself, especially if the environment has changed. Imagine a plant with stomata that are structured in a particular way, producing a high rate of transpiration. This type of plant has evolved in a warm, moist climate, where a high rate of transpiration keeps it cool, and there is little risk of water deficit. In such an environment, stomata with this particular structure were selected for because of the evolutionary advantage conferred by this high rate of transpiration. Now imagine that the climate changes from warm and moist to cool and dry. The etiological function of the stomata would still be to produce a high rate of transpiration, yet doing so in the new environment might be quite bad for the plant. Cooling is less of an advantage in the new environment, and water deficit is very likely. Even if these parts/traits exist as they do because they facilitate a high rate of transpiration, we cannot conclude that continuing to perform this function will be good for the plant. We cannot even conclude that continuing to *be able to* perform this function will be good for the plant. Retaining the ability to release lots of moisture into the air doesn't benefit a plant at all in a dry, cool climate. The ability to perform this function will not have any positive effect on the plant in this environment, and it will most likely come at some biological cost.[16] Unless we stipulate that retaining all evolved etiological functions *just is* what's good for a plant, which would entirely beg the question, there seems to be no reason for thinking that retaining this function would be in the plant's interest.

[15] The term "behavior" is used very broadly here; the term "effects" could be substituted with no change in meaning.

[16] I assume throughout this paper that if anything is bad for living things, dying is. Thus if continuing to do what its ancestors did causes a plant to die, then continuing to do what its ancestors did counts as bad for the plant. Someone could, of course, challenge the assumption that dying is bad for living things. However, if we are to test the adequacy of theories of welfare, we must rely on some independent sense of which welfare claims are plausible. I take these to be the least controversial of such claims.

The etiological function account of interests says that what's good for an organism is for it to keep doing whatever helped its ancestors survive and reproduce. As a claim about welfare—even biological welfare—this is peculiar. It expresses a deeply conservative view: that what worked in the past is what's best (for the organism) for the future. A principle such as this might work in a world where the environment never changes from generation to generation, but that is not the actual world.[17] Perhaps this example reveals one difference between ascriptions of interests and ascriptions of functions. While we might still want to insist that evolutionary history determines what the function of a part/trait is, there is less reason to think that evolutionary history determines what will be good for an organism. Along similar lines, we might think that it can be good for an organism for its parts/traits to do something unprecedented, thereby helping it to survive in a novel environment; however, it might be more difficult to accept that doing something unprecedented is the function of those parts/traits.[18]

Basl's view is that both conscious selection and natural selection generate interests. His view thus inherits the problem just mentioned, since it claims that natural selection determines the interests of evolved organisms. In addition, his view faces problems stemming from his further claims about conscious selection. Even if we were to accept that designers' intentions determine what the *purposes* of an object's traits/parts are, there is no reason for thinking that fulfilling these purposes must be good for the designed object. To see why, imagine that I design a car that will explode when I press a certain button. (It has a delay function, giving me plenty of time to escape before it explodes.) It might be true that the *purpose* of the button, as I have designed it, is to cause the explosion. However, it is difficult to see how fulfilling this purpose benefits the car. On the contrary, it would seem that if anything is bad for a car, being blown up is. One needn't agree with this claim about automotive welfare, however, to see the main point here, which is that there needn't be any relationship between a designer's intentions and the welfare of the object designed. Designers might intend to harm the objects they design and include features that have this purpose. If Basl's view is correct, however, designers cannot do this. The purposes that they bestow on the parts/traits of the objects they design *become* what is good for those objects. The peculiarity of this conclusion suggests that there is also a difference between ascriptions of purposes and ascriptions of interests. While it might be the button's purpose to trigger an explosion, there is no reason for thinking that it benefits the car for the button to do so, or even to be able to do so.[19]

[17] As I note elsewhere (McShane 2014), climate change is making such change more common.

[18] Godfrey-Smith (1994, p. 353) notes that the past-directedness of function claims comes from their use in biology in "explanations of why the functionally characterized entity exists, or exists in the form it does." In discussing functions, he claims that "if the explanandum is how things are now, nothing present or future can be the explanans. Only the past will do." This does not seem to be true for the explananda of interests.

[19] Note that this problem will arise even if we claim that the function of parts/traits is not a matter of the designer's intentions, as Basl claimed, but rather a matter of which features caused people to produce more cars with this feature. This is a different account of conscious functions that makes conscious functions more like natural functions: whatever is causally responsible for the reproduction of a part/trait in subsequent versions is what the function is. If my exploding car became popular, such that more car-makers designed cars with self-destruct buttons, causing the car to blow up would count the function of the button according to this new account, but fulfilling that function still would not be good for the car.

I conclude, then, that both the earlier etiological function accounts of interests and Basl's more recent account are flawed because they posit a conceptual relationship between facts about selection history and facts about interests, when in fact there is at best a contingent overlap, and only in favorable circumstances, between these things. Etiological accounts might or might not be good accounts of functions or purposes, but they are not good accounts of interests.

4 Objections

But perhaps this is too fast. One might worry that my description of the relevant functions or purposes in the examples above were designed to cause problems. There might be other equally legitimate descriptions of the function of the stomata or the purpose of the car's button that would have a better result for the accounts I am criticizing. Basl (2019, p. 95, fn. 60), in fact, argues that this is the case.[20] Following Goode and Griffiths (1995), he argues that organisms have multiple "ends" (in my terminology, functions) and that some of these are "nested." In the above example, we can say that while one description of the etiological function of the plant's stomata is to facilitate a high rate of transpiration, another description of their function is to produce a rate of transpiration appropriate to their environment; yet another description of their function is to help the plant survive and reproduce. All of these, Goode and Griffiths argue, are legitimate descriptions, and we should not see them as in competition with one another. It is true that if the high rate of transpiration weren't appropriate to the environment, stomata with this structure would have never been selected for in the plant's ancestors. It is also true that if transpiration appropriate to the environment hadn't enhanced the plant's ancestors' ability to survive and reproduce, they never would have been selected for. However, Goode and Griffiths argue, we don't want to say that the function of all traits is merely survival and reproduction. Lower-level realizations of these higher-level traits, they argue, are essential to biological explanation. The function of the stomata, for example, isn't just to help the plant survive and reproduce; it's to do so in a particular way. Basl argues that we should see these multiple correct descriptions as "nested ends," and he leaves it an open question how we are to prioritize among them. For purposes of biological explanation, Goode and Griffiths argue that different function-descriptions are appropriate at different levels of biological explanation, and so we should choose which level to prioritize by looking at which phenomena we are trying to explain. This leaves it open that we might choose a function-description that does not cause problems for attributions of interests in the current environment.

Let us consider how an etiological account of functions that includes Goode and Griffiths' pluralism about function ascriptions might work. In explaining an organism's interests, on such a view, it is unclear how we could avoid giving priority to the highest-level (i.e., most general) ends, survival and reproduction.[21] As we have seen, behaviors that interfere with these ends might still be considered functions of an organism's

[20] Thanks to an anonymous reviewer for pressing this point.

[21] See the next section, however, for an argument against including reproduction.

parts/traits, but they cannot be considered interests of the organism. In this regard, the priorities that are appropriate to an account of welfare might differ from those that are appropriate to biological explanation in general.

On this reformulated view, lower-level functions could only count toward an organism's welfare insofar as they are compatible with—perhaps even contribute to—this highest-level function. In the above example, this would yield the following alternative function-descriptions of the stomata:

i. Contributing to the plant's survival and reproduction, by achieving the level of transpiration appropriate to the environment, by achieving a high level of transpiration; or
ii. Achieving a high level of transpiration, but only if doing so is appropriate to the environment, and if appropriateness to the environment increases the ability to survive and reproduce.

With these alternative descriptions of the function of the stomata in hand, we can see the objector's argument more clearly. Contrary to what I argued above, on this view when the plant moves to the cool dry environment, its stomata either cannot perform [in the case of (i)] or do not have [in the case of (ii)] a function that involves achieving a high rate of transpiration.[22] In either case, achieving a high rate of transpiration is ruled out as an etiological function of the stomata, and thus achieving a high rate of transpiration will not count as a benefit to the plant in this environment.

There are problems with this reformulated view, however. It is still impossible for novel behaviors to be good for the plant, even in cases where they would have a significantly positive impact on the plant's ability to survive and reproduce. As long we are still working within an etiological account, interests will be limited to those behaviors that have a certain selection history. Unprecedented behaviors—new ones caused, for example, by a random mutation in a particular individual—cannot be among the plant's interests. This will be true even if we include as functions all possible etiological function-descriptions at all levels of generality. The etiological account still requires that the part/trait be here *because* it performs this function. In natural selection, this means being selected for; in conscious selection, it means being chosen by designers for this reason. Novel behaviors caused by random mutations (in natural selection) or unforeseen effects (in conscious selection) will not meet this criterion. However, as we saw above, while this might be a legitimate a constraint on function-attributions, it is more difficult to see why it should be a constraint on interest-attributions. Mutations can still be good for an organism, unforeseen effects can still be beneficial to designed objects, even if we are not willing to count them as the functions or purposes of the parts/traits involved.

The problems with etiological accounts seem to result from the backward-looking nature of etiological function attribution. What counts as etiological functioning is a matter of something's origin story: the explanation of how a thing's parts/traits came

[22] Under reformulation (i), we would say that it is impossible for the stomata to perform this complex function in the new environment, since it cannot achieve the level of transpiration appropriate to the environment *by* achieving a high level of transpiration. Since etiological functions must be something that the part/trait actually does [recall Wright's condition (a)], it will not count as an etiolotical function. Under reformulation (ii), we would say that the condition for "achieving a high rate of transformation" to count as a function of the stomata has not been met, and thus doing so is not a function of the stomata.

to be here as they are. While what to count as an etiological function is a historical matter, what to count as beneficial or harmful to a thing seems to be a matter of what happens to the thing in the present and the future. I conclude that this reliance on history, which is the central and defining feature of etiological accounts, is precisely the feature that should lead us to reject them as accounts of interests.

5 Alternatives

One reason for the appeal of etiological function accounts of interests seems to be a grave concern about the alternatives.[23] If we give up on the etiological function account of biological interests, what are we left with? In this section, I review the main alternatives and consider their advantages and disadvantages.

If the etiological function account of biological interests fails, we must reject one or more of the following claims:

(1) that the etiological account of function is the best account of function;
(2) that the relevant kind of teleological organization is the possession of parts/traits with functions;
(3) that having a good of one's own is a matter of being teleologically organized in some way;
(4) that the distinguishing mark of things that have biological interests is having a good of their own; or
(5) that biological interests do exist.

This list is in descending order, from the narrowest claim to the broadest. From the biocentrist's point of view, giving up (1) would require the smallest change to existing theories; giving up (5) would require the biggest change. Let us consider the implications of giving up each of these claims, starting with the smallest change to existing theories and working our way up to the biggest change.

1. The etiological account of function is the best account of function.

The rejection of (1) is compatible with preserving the bulk of the biocentrist's analysis—i.e., it would not require us to reject (2) through (5). All that we would need to do is find a better analysis of function to substitute for the etiological account. The question, of course, is whether a better analysis exists. Can we find an alternative account of function that would do a better job than the etiological account did as an account of biological interests?

The main early competitor to Wright's etiological view of function was Robert Cummins' (1975) causal role account. On Cummins' view, a part's/trait's function is its contribution to a specified capacity of a specified containing system. Since a part/trait can have different containing systems, each of which might have many different capacities, Cummins' view only allows one to talk about *a* function of a part/trait, not *the* function of a part/trait. This makes the account very flexible; it allows the same part/trait to have different functions within different containing systems or relative to

[23] Basl and Sandler (2013a, p. 704) are explicit about this fact. They describe the alternative as being "left with nothing but speculation and arbitrary claims about the good of non-sentient living things."

different capacities. However, this very flexibility is a problem for those who want to explain biological interests as a matter of functioning well, since the account itself does not privilege any particular containing system or capacity. A tree's xylem might have the function of transporting water and dissolved minerals from the roots to the canopy relative to the circulatory capacity of its vascular system. It might also have the function of making woodworkers rich relative to the profit-generating capacity of the capitalist economic system. Each of these is equally a function of the xylem. Neither has a claim to be *the* function of the xylem. A theory that identified the tree's well-being with the functioning of its parts would then have to say that a well-off tree is one in which the xylem both transports water and dissolved minerals effectively and makes woodworkers rich. It is difficult to see why latter would be beneficial to the tree.[24]

Another alternative account of function has been proposed by John Bigelow and Robert Pargetter (1987). Their "propensities account" focuses on the same aspect of organisms that the etiological account makes central, i.e., being subject to selective forces. However, the propensities account avoids the etiological account's reliance on history. According to the propensities account, the function of a part/trait is whatever increases its bearer's propensity to survive and reproduce in the future. On this view, functions aren't a matter of what increased an organism's ancestor's reproductive success in the past, but rather a matter of what will increase the organism's own reproductive success in the future. Insofar as the problems with etiological accounts described above were due to their backward-looking nature, Bigelow and Pargetter's view might seem to be the perfect solution.[25]

The etiological account and Bigelow and Pargetter's account, however, share a common problem. Both end up making biological fitness—i.e., reproductive success—the single value that establishes facts about welfare. There are reasons for thinking that this is a mistaken view of welfare. While an organism's survival and its reproduction are often related (an organism must survive, at least for a while, in order to reproduce), the goals of survival and reproduction are distinct from one another and can even compete with one another. In considering such cases, it becomes clear that the role of reproduction as a component of individual welfare is dubious. If a plant were to respond to drought by putting all of its energy into producing seeds rather than into maintaining its own structures, it might fail to survive yet increase its reproductive success. This would be a self-sacrificing strategy on the plant's part: decreasing its own survival prospects for the sake of producing more offspring. If it is possible for an organism to sacrifice its own welfare in order to increase its reproductive success, then we cannot identify an organism's welfare with its reproductive success. For these reasons, theorists such as Sune Holm (2012) and John Nolt (2009, pp. 149–151) have been inclined to rule out

[24] See Dussault and Bouchard (2017, p. 1120) for a similar criticism regarding ecosystem functioning.

[25] Bigelow and Pargetter's own view runs into a slight complication here. As Peter Godfrey-Smith (1994) has noted, the reproductive success of an organism depends on the environment it is in, and Bigelow and Pargetter specify that the relevant environment is the organism's "natural habitat," which they understand as a matter of (recent) history (1987, p. 192). This brings back the same problems we saw with previous etiological accounts, since if the environment changes, what worked well in the organism's "natural habitat" might not work well in its new habitat, and so might not be in its interest. In assessing a thing's interests, the relevant environment would have to be whatever environment currently occupies and will occupy in the future, i.e., during the time period for which the consequences for welfare are being assessed.

reproduction altogether as a component of welfare, claiming that welfare is a matter of survival, but not reproduction. The etiological and propensities accounts, however, do not distinguish between parts/traits that did (or will) increase fitness in virtue of enhancing reproduction and parts/traits that did (or will) increase fitness in virtue of enhancing individual survival. This is a problem for both accounts, and it might reveal another difference between function-attributions and interest-attributions. While we might want to say that it is the function of a part/trait to increase the organism's reproductive success, even at the cost of the organism's own survival, it is more difficult to see how sacrificing its survival for the sake of its offspring is in the interest of the organism. In understanding evolution and natural selection, reproductive success might well be of primary interest, in which case survival matters only insofar as it is needed to ensure the success of an organism's reproductive endeavors. However, in understanding what is good for an organism, the opposite seems to be true: survival is of primary interest, and reproduction matters only insofar as it has an effect on an organism's prospects for survival.

2. The relevant kind of teleological organization is the possession of parts/traits with functions.

There are other accounts of function that construe functions as derivative of some other, more basic type of teleological organization. On this kind of view, functions are simply behaviors or patterns that contribute to or maintain this other type of teleological organization. This other type of teleological organization has been described in different ways by different authors, but most regard it as involving a kind of goal-directedness, often with the specific goal of self-maintenance (Mossio et al. 2009; Holm 2012; see also Nolt 2009). While this type of theory may appear to involve only a rejection of (1), since it offers an alternative account of function, these accounts in fact constitute a rejection of (2). While they do offer a different account of function, the possession of parts/traits with functions is not what is important for the possession of interests. What is important is the other more basic kind of teleological organization in virtue of which their parts/traits come to have functions.

Sune Holm has offered an account of this type, defining self-maintenance as "a property of systems that are able to exert a causal influence on their surroundings in order to maintain (at least some of) the boundary conditions required for their own existence" (Holm 2012, p. 537). Holm claims that a thing's biological interests are a matter of self-maintenance: that which promotes its self-maintenance is beneficial to the organism, and that which interferes with its self-maintenance is harmful to the organism.[26] However, as Holm himself notes, living things do not seem to be the only self-maintaining systems that exist. Flames and hurricanes are also self-maintaining systems, he admits, and thus they would possess interests on his account. Like Basl, Holm thinks that we ought to accept that nonliving things can have interests, since the

[26] John Nolt (2009, p. 149) identifies a thing's biological interests with its autopoietic functions, i.e., "those that establish, maintain, or enhance [its] survivability." However, he presupposes an etiological account of function. While he thus excludes etiological functions that increase fitness in virtue of their effect on reproduction, the remaining functions—the ones that do determine a thing's interests – are determined in the same backward-looking way as the earlier etiological accounts described above. For this reason, his account inherits the problems we have seen with those views.

mere possession of interests doesn't entail that those interests are morally important (Holm 2012, pp. 539–540). Nonetheless, the idea that flames and hurricanes have interests is a challenging one.[27]

More important, however, are the philosophical questions raised by Holm's account. On Holm's view, self-maintenance is the crucial capacity of systems that makes them eligible for ascriptions of welfare and that determines what their interests are. But why should we think that only things that maintain themselves (in the sense described above) have interests? On Holm's view, if A maintains itself, then A can be harmed, but if B maintains A, then A cannot be harmed. It is difficult to see why this should be the case, i.e., what the degree of one's ability to self-maintain has to do with whether one can be harmed. In thinking about welfare in other areas, a thing's ability to maintain itself does not play the same role: we don't say that only people who are economically self-maintaining can be harmed economically, or that people who are socially and emotionally maintained by others are thereby not susceptible (or less susceptible) to harm. Typically, ethicists take the opposite view: that those who are entirely dependent for their survival on others have stronger interests as a result of their inability to self-maintain. Whatever one thinks of these comparisons, organizational approaches such as Holm's seem to have an unmet justificatory burden. They need to explain why we should think that self-maintaining systems are special, such that damage to them counts as a harm, while damage to a system that is maintained by forces outside of itself does not count as a harm—i.e., why this particular kind of causal independence from others is a marker of things that can be harmed.

3. Having a good of one's own is a matter of being teleologically organized in some way.

Holm's view and Basl's view share a common feature: they both accept that the marker of things that have interests is that those things are teleologically structured. Holm and Basl disagree over the correct account of this teleological structure, Basl claiming that it is a matter of having a certain selection history and Holm claiming that it is a matter of being organized as a self-maintaining system. For both of their theories, we have seen reasons for worrying about whether the account of teleological structure picks out a feature that should be thought of as a requirement for the possession of interests. On Basl's view, we saw good reasons for thinking it does not: behaviors that were selected (for) in the past might well turn out to be harmful to their bearer in the present. On Holm's view, we saw no reason for thinking that it does: a thing's ability to maintain the boundary conditions for its own existence bears no obvious relation to its ability to be harmed or benefited. In light of these doubts, we might start to wonder whether "teleological structure" was the right thing to be looking for in the first place.

The reason that most biocentrists have focused on the teleological structure of organisms seems to be methodological. Humans clearly have interests, and so if we try to develop an account of biological interests by looking for what plants have in common with humans (but not with rocks), teleological structure might well be where

[27] See Mossio and Bich (2017) for an alternative description of self-maintenance (which they call "self-determination") involving organizational closure. As they concede, their view might not rule out nonliving things as self-determining systems. Even if it did, however, the questions raised below about the relevance of self-determination to welfare will still apply to their view.

we end up.[28] Not only does teleological structure clearly distinguish plants from rocks, but it might seem analogous to certain welfare-related features of humans: plants have (nonconscious) goals or functions; people have (conscious) ends or intentions. If our possession of ends or intentions is somehow essential the way in which human individuals have a good of their own, then perhaps the analogous goals or functions of plants can be used to explain why plants have a good of their own. This extensionist strategy has been criticized in many ways (e.g., Plumwood 1993), but the rejection of (3) doesn't even require abandoning extensionism. It's possible that we've just extended interest attributions to plants using the wrong criterion. Perhaps some other similarity that we share with plants explains why they have a good of their own. Of course, determining what this other similarity is would require a much deeper analysis of what it is to have a good of one's own—i.e., what we are saying about a thing when we say that it has a good, and why that good is a good of its own. To my knowledge, however, no biocentrist has pursued this alternative (nonteleological) account of what it is to have a good of one's own.

4. The distinguishing mark of things that have biological interests is having a good of their own.

The earlier criticisms of Holm's view might lead us to question the legitimacy of (4). Why, after all, must a thing's good be a good *of its own* in order for the thing to count as having interests? Is this just a Western cultural preference for individualism, independence, and self-sufficiency at work in our theorizing?[29] Recall that this requirement was introduced to rule out artifacts as possible bearers of interests. The claim was that it might be good for the car that its oil pump continue to circulate oil around the engine, but this good is really something that is good for me, the person who wants to drive the car, not something that is good for the car in its own right. There might be some slipperiness in this line of reasoning.[30] For our purposes, however, the important questions are (a) What condition is being imposed when we say that something's good must be a good of its own in order for that thing to have interests? and (b) What justifies the imposition of that condition?

The discussion of Holm's view suggested that there are problems with thinking of this condition as requiring that one's good is not dependent upon the goods of others. This would rule out things that depend on one another to get by, and there seems to be no independent justification for thinking that such things cannot be harmed. The condition also should not be that a thing's good not be identical to or coextensive with the goods of others. That our interests are shared should not rule them out as interests. Nor should the condition be that a thing comes to have this as a good on its own—i.e., that it is not caused to have this good by others. The causal story of the history of

[28] See Basl (2019) for both a discussion and an example of this methodology. See also Taylor, Agar (2001), Johnson (1991), Rolston (1988), Taylor (1986) for exemplars.

[29] Feminist scholars have argued that Western ethical traditions do treat self-sufficiency and autonomy as criteria for full personhood, and that it is a problem (Meyers 1997; see also Plumwood 1993). It tempting to see something like this going on in the background of views that make *self*-maintenance required for the possession of interests or require that a good be a good *of one's own*.

[30] Some questions we might raise: Is satisfying my desire to drive the car really good for me? If it were bad for me, would we have to say that the functioning of the oil pump is bad for the car?

all organisms, for example, essentially involves their ancestors: neither our bodies nor our interests are self-created. Having ruled out content-independence, identity, coextension, and causal independence, what is left? In the case of the car, the worry seemed to be that the *only reason* that a functioning oil pump is good for the car is that it's good for me. That is to say, absent me and my good (or perhaps more accurately, my preferences), there would be no grounds for saying that a functioning oil pump is good for the car. More broadly, absent people and their preferences, there would be no grounds for saying that anything at all is good for cars. What we have here seems to be a kind of isolation test: If, absent everything else and its interests and preferences, we can still say that something is good for an entity, then that entity has a good of its own. However, with highly interactive organisms, it is often the case that one organism's interests are the sole reason for another organism's interests. For example, certain soil conditions are bad for trees only because they are bad for rhizobia, the symbiotic organisms that live among tree roots and fix nitrogen in the soil. It is difficult to find a construal of (4) that doesn't devalue high degrees of interactivity, such that highly interdependent organisms will end up having the same status as cars when it comes to welfare claims.

That said, (4) might also reflect an important idea in Western moral philosophy: that individuals matter morally precisely because each individual has a good of their own. Animal rights theorists have cited this as a reason that moral importance must be attributed to individual animals, not merely to the species, ecosystems, or human communities to which they belong. The fact that *this particular* dog has interests, interests that are independent of the interests of dogs generally, of its species, of its family or owners, of its municipality, of its local ecosystem, etc. means that its interests must be considered independently; they cannot be subsumed under these other interests (Regan 1983). In political philosophy, the fact that each citizen has an independent good, interests that are independent of the interests of other citizens or of the state generally, means that laws must be justifiable to each citizen in order to be just (Rawls 2005). To give up (4), then, might be to undermine the idea that the uniqueness of a thing's own point of view on the world matters. In other parts of moral philosophy, this idea is one that has been considered very important.

5. Biological interests do exist.

The final possibility, of course, is that we might give up (5): the idea that nonsentient, nonconscious organisms have interests at all. Perhaps we are simply mistaken when we say that that injecting poison into its roots harms the oak tree. Perhaps plants are just like rocks or cars after all. Rejecting (5) requires us to view quite a lot of ordinary discourse as mistaken, or at best as merely metaphorical, and it privileges consciousness in our understanding of interests in a way that some have argued is parochial (Goodpaster 1978). Nonetheless, the widespread view that living things are importantly different from nonliving things might be explained as merely a quirk of human psychology. Just as we often read intentions into nonintentional phenomena, read purposes into random chance events, and overestimate in-group similarities and out-group differences, perhaps we have a tendency to empathize with other living things, thinking of them as, like us (and unlike nonliving things), putting forth effort, *trying* to achieve things. Unlike rocks, plants seem to be doing the same sort of thing

that we are: trying to get enough to eat, to grow and develop, to stay safe from danger, to repair damage, to fight disease, and so on. But perhaps this way of perceiving plants is just a kind of biophilic cognitive bias: our psychologies might be set up to make life particularly cognitively and emotionally salient in our perceptions of the world and to interpret the behavior of living things as effortful and intentional in ways that it is not.[31]

How we should respond to the failure of etiological function accounts of biological interests is a difficult question, and I have not argued for any particular answer to it here. What the above analysis shows, however, is that there are deep philosophical and conceptual questions involved in the decision to accept or reject claims (1)-(5). These are questions about the very nature of harm and benefit; answering them requires not simply looking at what living things have in common with one another or with us, or identifying some sort of vaguely teleological aspect of the nonhuman world, but rather thinking about what we must be presupposing about a thing when we say that it hasn't just been damaged, but also harmed by that damage.

The fact that some interests might not matter morally, as Basl and Holm argue they do not, does not let us off the hook in getting our ascriptions of welfare right. To attribute interests to flames and pencil sharpeners but not to organisms that are maintained by others is to make substantive philosophical claims, whether or not we decide that the relevant interests are morally important. That is to say, theories of welfare matter in their own right, not merely because of the role they play in moral theories.

6 Conclusion

While etiological function accounts of interests solve many theoretical problems for biocentrists, they are ultimately bad accounts of the interests of organisms. They make evolutionary history, design history, and/or reproductive success the determinants of welfare, and facts about these things are simply not good guides to what will benefit an organism. Alternative accounts face their own problems; their prospects depend on whether they can justify the conceptual relationships that they posit between possessing certain characteristics and having a welfare. What this paper has shown is that such a justification will require not merely a different account of function with fewer problematic implications for plant interests, but a deeper philosophical understanding of the nature of harm and benefit itself.

Acknowledgements The author thanks Robin Attfield, Antoine Dussault, Jeff Kasser, Jeffrey Pannekoek, and participants in the Fall 2014 Environmental Philosophy graduate seminar at Colorado State University for helpful discussions of some of the ideas presented here. An earlier version of this paper was presented at Texas A&M University, and a related paper was presented at the annual conference of the Groupe de Recherche en Éthique Environmentale et Animale in Montreal; thanks to both audiences for their helpful thoughts and suggestions. Finally, particular thanks to two anonymous reviewers from Synthese. This paper was enormously improved by the careful scrutiny it received in the review process. My deepest gratitude for the attention, care, and rigor they devoted to their reviews of the manuscript.

[31] For a criticism of this line of reasoning, see Plumwood (2009).

References

Agar, N. (2001). *Life's intrinsic value: Science, ethics and nature*. New York: Columbia University Press.

Attfield, R. (1981). The good of trees. *Journal of Value Inquiry, 15,* 35–54.

Attfield, R. (2012). Biocentrism and artificial life. *Environmental Values, 21*(1), 83–94.

Basl, J. (2012). Nothing good will come from giving up on aetiological accounts of teleology. *Philosophy and Technology, 25,* 543–546.

Basl, J. (2017). A trilemma for teleological individualism. *Synthese, 194,* 1057–1074.

Basl, J. (2019). *The death of the ethic of life*. Oxford: Oxford University Press.

Basl, J., & Sandler, R. (2013a). The good of non-sentient entities: Organisms, artifacts, and synthetic biology. *Studies in History and Philosophy of Biological and Biomedical Sciences, 44,* 697–705.

Basl, J., & Sandler, R. (2013b). Three puzzles regarding the moral status of synthetic organisms. In G. E. Kaebnick & T. H. Murray (Eds.), *Synthetic biology and morality: Artificial life and the bounds of nature*. Cambridge, MA: MIT Press.

Bigelow, J., & Pargetter, R. (1987). Functions. *The Journal of Philosophy, 84*(4), 181–196.

Cahen, H. (1988). Against the moral considerability of ecosystems. *Environmental Ethics, 10,* 195–216.

Cummins, R. (1975). Functional analysis. *Journal of Philosophy, 72,* 741–765.

Dussault, A. C., & Bouchard, F. (2017). A persistence enhancing propensity account of ecological function to explain ecosystem evolution. *Synthese, 194*(4), 1115–1145.

Foot, P. (2001). *Natural goodness*. Cambridge: Cambridge University Press.

Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs, 28*(3), 344–362.

Goode, R., & Griffiths, P. E. (1995). The misuse of Sober's selection for/selection of distinction. *Biology and Philosophy, 10,* 99–108.

Goodpaster, K. E. (1978). On being morally considerable. *Journal of Philosophy, 75,* 308–325.

Holm, S. (2012). Biological interests, normative functions, and synthetic biology. *Philosophy and Technology, 25*(4), 525–541.

Holm, S. (2017). Teleology and biocentrism. *Synthese, 194,* 1075–1087.

Hursthouse, R. (1999). *On virtue ethics*. Oxford: Oxford University Press.

Johnson, L. E. (1991). *A morally deep world*. Cambridge: Cambridge University Press.

McShane, K. (2014). Individualist biocentrism vs. holism revisited. *Les ateliers de l'éthique/The Ethics Forum, 9*(2), 130–148.

Meyers, D. T. (1997). *Feminists rethink the self*. Boulder, CO: Westview Press.

Mossio, M., & Bich, L. (2017). What makes biological organisation teleological? *Synthese, 194,* 1089–1114.

Mossio, M., Saborido, C., & Moreno, A. (2009). An organizational account of biological functions. *British Journal for the Philosophy of Science, 60*(4), 813–841.

Nolt, J. (2009). The move from is to good in environmental ethics. *Environmental Ethics, 31*(2), 135–154.

Odenbaugh, J. (2017). Nothing in ethics makes sense except in the light of evolution? Natural goodness, normativity, and naturalism. *Synthese, 194*(4), 1031–1055.

Plumwood, V. (1993). *Feminism and the mastery of nature*. New York: Routledge.

Plumwood, V. (2009). Nature in the active voice. *Australian Humanities Review, 46,* 113–129.

Preston, C. (2018). *The synthetic age: Outdesigning evolution, resurrecting species, and reengineering our world*. Cambridge, MA: MIT Press.

Rawls, J. (2005). *Political liberalism*. New York: Columbia University Press.

Regan, T. (1983). *The case for animal rights*. Berkeley: University of California Press.

Rolston, H., III. (1975). Is there an ecological ethic? *Ethics, 85,* 93–109.

Rolston, H., III. (1988). *Environmental ethics*. Philadelphia: Temple University Press.

Samuelsson, L. (2010). On the demarcation problem and the possibility of environmental ethics: A refutation of "a refutation of environmental ethics." *Environmental Ethics, 32*(3), 247–265.

Taylor, P. (1986). *Respect for nature: A theory of environmental ethics*. Princeton: Princeton University Press.

Thompson, A. (2007). Reconciling themes in neo-Aristotelian meta-ethics. *The Journal of Value Inquiry, 41,* 245–263.

Thompson, J. (1990). A refutation of environmental ethics. *Environmental Ethics, 12,* 147–160.

Varner, G. E. (1990). Biological functions and biological interests. *The Journal of Southern Philosophy, 28*(2), 251–270.

Varner, G. E. (1998). *In nature's interests? Interests, animal rights, and environmental ethics*. Oxford: Oxford University Press.

von Wright, G. H. (1963). *The varieties of goodness*. London: Routledge & Kegan Paul.
Wright, L. (1973). Functions. *Philosophical Review, 82,* 139–168.