

Internal perspectivalism: The solution to generality problems about proper function and natural norms

Jason Winning  • jason.winning@gmail.com
Department of Philosophy, University of California, Berkeley

Abstract

In this paper, I argue that what counts as the proper function of a trait is a matter of the *de facto* perspective that the biological system, itself, possesses on what counts as proper functioning for that trait. Unlike non-perspectival accounts, internal perspectivalism does not succumb to generality problems. But unlike external perspectivalism, internal perspectivalism can provide a fully naturalistic, mind-independent grounding of proper function and natural norms. The attribution of perspectives to biological systems is intended to be neither metaphorical nor anthropomorphic: I do not mean to imply that such systems thereby must possess agency, cognition, intentions, concepts, or mental or psychological states. Instead, such systems provide the grounding for norms of performance when they internally enforce their own standard of (i.e., their own perspective on) what constitutes proper functioning or malfunctioning. By operating with a fixed, determinate level of generality, such systems provide the basis for an account of proper function that is immune to generality problems.

Keywords: Generality Problem; Natural Norm; Perspectivalism; Perspectivism; Proper Function; Teleology

1. Introduction

In a wide variety of contexts, scientists and philosophers take for granted that it is meaningful to refer to or inquire into *the* biological function of a given trait, property, or part (from now on I will just say ‘trait’ for sake of brevity) of an organism or biological system. The idea here is that although a given trait may be capable of performing any number of functions—in the sense that there are various causal interactions it may enter into, and hence various causal roles it can play in those interactions—there might only be one function that it *has* (its *proper function*).¹ The

¹ This usage of the phrase ‘proper function’ follows Millikan (1984, p. 2). Often, ‘teleological function’ is used to mean what I intend by ‘proper function’ (e.g., Maley & Piccinini, 2017). Not all accounts of proper

distinction between properly functioning versus malfunctioning traits is extremely important to philosophy because it has been thought to provide resources for resolving a number of philosophical debates. For example, proper function has figured centrally in philosophical accounts of representation (Dretske, 1988), computation (Piccinini, 2015, p. 118), action and agency (Burge, 2009), mechanisms and mechanistic explanation in biology (Garson, 2013), health and disease (Boorse, 1997), epistemic norms (Sullivan-Bissett, 2017), and even value as such (Foot, 2001), among many other projects falling under the heading of “naturalizing normativity.” Proper function is an equally important notion for biologists. Their investigations are often inquiries into what *the* function of a given trait is, or into what *the* function of parts that contribute to produce a larger phenomenon or behavior are, or what *the* function of certain activities or developmental processes (I will also include these under the umbrella of ‘trait’) are in the life cycle of given organisms.

An account of proper function should, at minimum, tell us what it is in virtue of which proper function ascriptions are true or false. In this paper, I discuss a challenge for accounts of proper function that is tantamount to the “generality problem” that has been raised against process reliabilism in epistemology (Section 2). I argue that this should not be surprising: Both types of account explain a normative status (epistemic justification of beliefs in one case, proper function of traits in the other) by reference to the type a certain token falls under, but such tokens fall under many such types of varying generality, and such accounts must offer a way to tell us which is the relevant type in given cases (Section 3). Perspectival accounts of proper function² can potentially deal with this challenge, since the perspective that a functional attribution is relativized to may involve a univocal interpretation of the situation that disambiguates between the problematic multitude of types.

I further divide perspectival accounts into two categories: external perspectivalism and internal perspectivalism. Up to now, discussions of perspectivalism have tended only to focus on external perspectivalism (e.g., Massimi, 2018). On an external perspectivalist account, proper function within a given system depends on the perspective taken by an observer that is external to the system in question (such as a scientist who is investigating the system). Such accounts are sometimes not considered to be fully naturalistic and have other well-known problems (Section 3.2). On an internal perspectivalist account, by contrast, the proper function of a trait within a given system depends on the perspective had by the *system itself*. I argue that internal perspectivalism is superior because it (a) does not succumb to generality problems, and (b) can provide a fully naturalistic, mind-independent grounding of proper function and natural norms. (Sections 4–6).

function (as I am using the term) provide for a notion of malfunction. It should also be noted that, at least on some views, traits can have multiple proper functions; for example, it might be said that the hind legs of turtles have the proper functions both of locomotion and of excavation (Preston, 1998).

² As far as I know, no perspectival version of process reliabilism has been offered, perhaps because process reliabilists tend to be epistemic externalists. Sosa (1991, pp. 290–291) offers a perspectivalist solution to generality problems that arise for virtue epistemology.

My attribution of perspectives to biological systems is intended to be neither metaphorical nor anthropomorphic: I do not mean to imply that such systems thereby must possess agency, cognition, intentions, concepts, or mental or psychological states.³ The notion of a perspective to be developed here (in Sections 4 and 5) is much thinner, depending only on minimal recognition and response capabilities by means of which even “simple” organisms without nervous systems, and subsystems within such organisms, are selectively sensitive to details of their outer and inner environment at a fixed level of generality. Such systems provide the grounding for norms of performance when they internally enforce their own standard of (i.e., their own perspective on) what constitutes proper functioning or malfunctioning.⁴ Since they operate with a fixed, determinate level of generality, such systems provide the basis for an account of proper function that is immune to generality problems.

2. What is a Generality Problem?

I will use the case of process reliabilism in epistemology to explain what generality problems are since this is an existing context in which they are well-known and in which examples are easily generated and grasped. Process reliabilism is a theory of epistemic justification stating that

The justificational status of a belief is a function of the reliability of the process or processes that cause it, where (as a first approximation) reliability consists in the tendency of a process to produce beliefs that are true rather than false. (Goldman, 1979, p. 10)

Process reliabilism attempts to say what makes a justified belief different from an unjustified belief (a justified belief must have been formed by a *reliable* type of process), and it attempts to tell us what makes it the case that beliefs have any justificational status in the first place (the purported fact that any given belief will have been formed by a process that has some determinate level of reliability).

Conee and Feldman (1998) argue that belief-forming process tokens will always fall under a number of different types of varying generality and of varying reliability. For example, suppose I look out my window and form the belief that my friend Sharon is walking on the sidewalk outside. Suppose further that the process by which I formed my belief counts as a i) visual process at night, ii) visual process at night in a well-lit area (there is a street light nearby), and iii) visual process on a foggy night. We might characterize my belief as having been formed by a token process of any of these types, among others, depending on the level of generality at

³ My internal perspectivalism is therefore very different from that of Sinnott, who argued that “biological organization ... and psychological activity ... *are fundamentally the same thing*” (1961, p. 48). Pattee, whose work partly inspired the present view, might also be read as an internal perspectivist about biological function (e.g., 1970, p. 130; 1982).

⁴ This paper is concerned with *metaphysical* questions, i.e., questions about the nature of proper functions themselves, rather than epistemological questions about how humans can/should go about investigating them. In this paper, I do not address questions about how humans discover, test/confirm hypotheses about, or reason about proper functions.

which we choose to individuate such processes. But note that some of these belief-forming process types (BFPTs) are reliable, and some are not. Process reliabilism will therefore give wildly differing answers depending on the level of generality at which we individuate processes. Without resources for specifying the relevant level of BFPT generality, such an account fails even to give you the means to know when you are anywhere near the right answer. It therefore does not tell us which non-normative facts make it the case that beliefs are justified or unjustified.

As Michael Bishop put it, “Without a principled solution to the generality problem, the reliabilist can always start with the intuitively correct judgment about a belief’s justificatory status and then cherry pick a BFPT that yields that judgment,” (2010, p. 287). Conee and Feldman argue that

Given the multiplicity of belief-forming process types and their variations in reliability, it is easy to make *ad hoc* case-by-case selections of types that match our intuitions. But case-by-case selections of relevant types does not constitute working out a reliabilist theory of justification. (1998, pp. 3–4)

In the intervening years, solutions have been proposed, but for my purposes what matters is that it continues to be recognized as a serious problem. It is important to realize that generality problems are not as innocuous as other types of indeterminacy⁵: an account that suffers from generality problems not only fails to give an exact answer; it fails even to give you the means to know when you are *anywhere near* the right answer.

3. How Generality Problems Arise for Existing Theories of Function

In this section I will briefly illustrate how generality problems arise in a very similar way for two prominent families of theories of biological function: those that base function on natural selection history (a.k.a. etiological theories), and those that base function on external perspectives (e.g., systemic causal role theories).

3.1. *Etiological Theories of Function*

According to the most prominent etiological approach, “a token trait’s function depends on what traits of the relevant type were selected for” (Neander & Rosenberg, 2012, p. 613).⁶ Whether or

⁵ See Neander (1995, p. 113) for a useful breakdown of several kinds of function indeterminacy that have been previously discussed.

⁶ I count the etiological approach as non-perspectival because, as many authors have pointed out (e.g., Darwin, 1876, p. 63; Polányi, 1958, p. 385; Nagel, 1977, pp. 286–287; Nyberg, 2009, p. 187; Okasha, 2009, p. 720), “natural selection” does not involve the presence of something in the background picking winners and losers according to its own normative standard or perspective. Instead, organisms succeed or fail at reproducing merely because of whatever particular traits happen to be helpful (or “adaptive”) in specific environments on specific occasions (Beatty, 1984, pp. 192–193). I will revisit this point in Section 4.1, where I consider the question of what it means for something to have a perspective. See Footnote 13.

not trait X is functioning properly is a matter of whether or not trait X can perform the task that it was selected for performing. Others have pointed out reasons why etiological theories may yield indeterminate answers to the question of what the proper functions of traits are. For example, as Walsh argues, natural selection is “relative to a selective regime” (1996, p. 553), defined as “the total set of abiological and biological (including social, developmental and physiological) factors in the environment of the trait which potentially affect the fitness of individuals with that trait” (1996, p. 564; cf. Bechtel, 1986). Etiological theories of function therefore “face indeterminacy problems, for there are more and less specific descriptions of selective regimes and functional outcomes” (Sterelny, 1995, p. 255). Goode and Griffiths (1995, p. 107) point out that indeterminacy is not necessarily problematic for theories of function. But as we saw in the previous section, generality problems are a particularly pernicious form of indeterminacy.

Consider the following example. Suppose the auditory system of a preyed-upon mouse population has become especially attuned to a certain frequency band B over a large number of generations, because its main predator (a certain hawk population) makes a distinctive sound in which frequency band B is dominant.⁷ Suppose further that over those many generations, due to the nature of their habitat, the mice gained almost no survival or reproductive benefit from detecting noises otherwise. Now suppose a mutation occurs in the hawk population, making it almost silent in frequency band B (producing what we may call “stealthy hawks”; call the non-mutated hawks “noisy hawks”). We may ask, is the auditory system of one of the mice malfunctioning when it fails to detect the stealthy hawk approaching? This depends on how we characterize the task that the auditory system was selected for, i.e., the task that conferred a benefit to its possessors. We may describe this as:

- Detecting predators that are noisy with respect to frequency band B
- Detecting noisy hawks
- Detecting hawks
- Detecting predators
- Detecting sounds in general

If “detecting sounds in general” is the task that the auditory system was selected for, then arguably, it is not malfunctioning. The mouse is generally able to detect sounds, just not those made by the stealthy hawks. Similarly, if “detecting predators that are noisy with respect to frequency band B” is the selected-for task, then it is not malfunctioning. But if the task is “detecting hawks” or “detecting predators,” then it is indeed malfunctioning. Our intuitions may pull us toward one answer or another, but as was the case with process reliabilism, it will be our intuitions doing the heavy lifting, not the account itself.⁸ The point here is that the bare historical

⁷ This example is loosely inspired by one given by Walsh (1996), but it is being used to make a different kind of argument.

⁸ Neander addresses this type of problem by arguing that when there are multiple ways to characterize functions, “we should give priority to that description of a trait’s function that is the lowest level in the

facts of “selection,” by themselves, are insufficient to pick out proper functions and malfunctions in given cases.⁹

3.2. *External Perspectivalist Theories*

External perspectivalist theories relativize function ascriptions to the perspectives of observers who are themselves external to the system in question, for example, scientists studying the system. External perspectivalist theories are largely a response to the fact that if functions are merely defined as the causal contribution a trait makes to a system, then almost anything can have any function (Davies, 2001, refers to this as a “promiscuity objection”) since there are so many ways to delineate systems and causal contributions. Roughly, the vast majority of systems and contributions to systems can be eliminated as being irrelevant to functional ascription because they are not of “explanatory interest” (Cummins, 1975, p. 764).

However, interest-relativity as a strategy for coping with promiscuity objections presents a dilemma. Either “explanatory interest” ineliminably reduces to subjective mental states or it does not. If it does (as on Craver’s 2001, p. 71; 2013, p. 155 view), then external perspectivalism has the same problem as process reliabilism: it treats the intuitions of scientists/philosophers (i.e., whoever is deciding what counts as “explanatorily interesting” in particular cases) as an inscrutable black box, and therefore doesn’t do the work that a theory of function is supposed to do.¹⁰

On the other hand, if “explanatory interest” reduces to some kind of criterion that is not merely subjective, then the external perspectivalist needs to specify what this criterion of “explanatory interest” consists in, and must do so in a way that avoids generality problems. Otherwise, as with process reliabilism, it will fail to provide a theory of proper function with definite content. The theory will need to make reference to the type a certain token (the trait token itself, or some token fact about it) falls under, but such tokens will usually fall under multiple such types of varying generality. For example, one of the most prominent approaches that takes this second option is the self-maintenance-based approach to proper function defended by Moreno among others. They avoid promiscuity objections by limiting the systems in question to those that are self-maintaining, and by limiting the relevant causal contributions to those that are contributions to the system’s self-maintenance. Moreno and Mossio (2015, p. 74) point out, however, that multiple conflicting function attributions are possible on this approach because whether (and

analysis (most mechanistic),” where mechanistic levels are connected by “asymmetrical by-relations” (1995, p. 137). But in the above case, unlike the cases Neander considers, the varying characterizations do not correspond to different mechanistic levels connected by asymmetrical by-relations.

⁹ Enç (2002) enumerates a series of arguments against etiological theories of proper function along similar lines as well.

¹⁰ Woodger went further, arguing that attempts to “explain internal teleology by means of external teleology” are no better than vitalism or Cartesian dualism if they leave the external teleology unaccounted for (1929, p. 441; cf. Pattee, 1977, pp. 260–261).

how) a system counts as self-maintaining depends on the relevant “regime of self-maintenance,” and these can be described at multiple levels of generality.¹¹

There are, however, additional problems for external perspectivalism. Bigelow and Pargetter argue that there is often a mismatch between functions and what we take an interest in:

it is assumed that biological structures *would* have had the functions they do have even if we had not been here to take an interest in them at all. And some of the effects of structures that we take an interest in have nothing to do with their function. And some functions are of no interest to us at all. (1987, pp. 183–184; cf. Maley & Piccinini, 2017, p. 240)¹²

Etiological theorists and external perspectivalists may yet be able to devise effective responses to such challenges. My intention in briefly considering the above theories and objections has not been to make an argument by elimination, but instead to motivate exploration for an alternative theory that does not face generality problems and that does not make proper functions relative to the subjective mental states or interests of external observers. It is to this exploration that I now turn.

4. Avoiding Generality Problems: Non-Minded Systems That Have Perspectives

4.1. Pattee, Dretske, and Selective Loss of Detail

In order to solve generality problems, any theory that analyzes proper function in terms of some other factors must tell us what it is about any given case that *fixes* the appropriate level of generality at which such factors should be considered. What kind of thing can “fix” or “pick out” a level of generality in the first place? Clearly, human beings can pick out a level of generality; this is just what Conee and Feldman and Bishop argue is happening when philosophers “apply” process reliabilism to given problems. Arguably, this is because human beings are capable of *abstraction*; they are capable of taking different perspectives on a situation, and characterizing it at varying levels of description, some of which are more fine-grained than others. The question then is: what sort of thing is it to *have a perspective* at all, i.e., the sort of perspective that allows us to “fix” *some* level of generality at all? What is special about the human capacity for *abstraction* that is relevant to generality problems? Howard Pattee’s answer to this is that *selective loss of detail* is the key:

All forms of discrimination, measurement, classification, selection, pattern recognition, etc., are accomplished by selective loss of details. Quite generally, it is this selective loss of details from the immense array of sensory inputs and memory stores that makes symbol

¹¹ Moreno and Mossio’s solution to this problem is to relativize function attributions to self-maintenance regime descriptions (2015, p. 74). While this allows them to circumvent generality problems, it arguably does so at the cost of rendering the notion of proper function less interesting and useful.

¹² Mark Bedau (1992, p. 37) makes similar arguments against what I am calling external perspectivalism about goal-directedness.

systems of all types, including all languages, functionally useful. Such discriminations are certainly one of the most primitive and fundamental functions of the brain. (1992, p. 190)

For instance, the ability to discriminate houses from other types of buildings requires a selective loss of detail: it requires sensitivity (or attention) to the details that distinguish houses from non-houses, but insensitivity (or lack of attention) to details that are irrelevant to making such discriminations (e.g., whether the building has four windows or five). By switching to a different perspective, different details are lost or treated as insignificant. The *generality* at which discriminations are made is given by the level of detail that is treated as significant, and the level of detail that is ignored, i.e., *selectively lost*.

But Pattee does not believe that a mind or even a nervous system is necessary for this. Relatively simple dynamical systems and devices can exhibit a selective loss of detail in the way that they are sensitive in their interactions with other objects. Similarly, Dretske argues that simple analog/digital devices, by systematically ignoring information at a certain level of detail, are capable of processes that involve assigning tokens to types, such as classification, categorization, or recognition:

Until information has been lost, or discarded, an information-processing system has failed to treat *different* things as essentially the *same*. It has failed to classify or categorize, failed to generalize, failed to “recognize” the input as being an instance (token) of a more general type. (1981, p. 141)

Dretske argues that it makes sense to speak of such a device as having its own “point of view,” according to which multiple tokens are considered to be of “the same” type. Similarly, Popper argued that since no two physical events are ever exactly the same, “for logical reasons, there must always be a point of view ... before there can be any repetition; which point of view, consequently, cannot be merely the result of repetition” (2002, p. 59). Dennett (1991) and Ross, Ladyman, and Collier (2007, pp. 220–227) use the word ‘perspective’ in the same sense as Dretske’s ‘point of view’. Devices that perform pattern recognition operate with a determinate scheme of sorting tokens into types—a *de facto* perspective on which tokens should get assigned to which types.¹³ In his book-length treatment of the notion of a “point of view,” Hautamäki argued that while this notion gets used in many different ways, the commonality is that when a given thing is being considered from a point of view, “certain aspects of [that thing] are considered while others are ignored” (1986, p. 7), i.e., a selective loss of details. On Hautamäki’s account of points of view, out of all of the possible “determinables” (i.e., types) that

¹³ I can now say with greater precision why (as noted earlier in Footnote 6) natural selection and the accounts of function that appeal to it are non-perspectival: because natural selection does not operate with a determinate scheme of sorting tokens (phenotypes, tasks, environmental conditions, outcomes, outcome histories, etc.) into types that systematically ignores information at a certain level of detail, i.e., into types at a fixed level of generality.

a determinate (token) object or situation may fall under, a point of view represents a selection of certain of these determinables as relevant and others as irrelevant (ibid., p. 65).¹⁴

Pattee argues that selective loss of details is ubiquitous within biological systems, as it is essential to the working of enzymes; in fact, Pattee argues that ultimately, all

pattern recognition and selective action [in biological systems] is mediated by enzymes or enzyme-like molecules. This is the case for the cell's sensing of the external environment, for sensing between the cells, and for intracellular recognition of patterns. (1982, p. 172)

Enzymes act as detection devices because of the selectivity with which they bind to other molecules. They have binding sites at which they form electrochemical bonds to specific types of molecules (often other proteins or carbohydrates). The formation of an electrochemical bond is much more like a discrete process than a continuous one; normally a protein either binds to a molecule or it doesn't. Further, enzymes are insensitive to differences between molecules of a given chemical species; either the catalytic reaction is carried out or it is not. Subtle differences in conformation between substrate molecules that bond to the enzyme do not translate to qualitative differences in catalytic reaction; such differences are effectively "ignored" by the enzyme:

out of all the microscopic collisions with ... an enzyme, only very special ones are capable of triggering their catalytic function, and furthermore when this function is triggered only a very limited or simple result takes place. This is in contrast to ordinary dynamical systems where almost any collision results in a complex perturbation spreading through the entire system with no coherent result whatsoever. (Pattee, 1973, p. 43)

In this sense, protein binding and catalysis, which is the basis for enzymatic reactions and intracellular signaling, counts as a fixed practice of recognition and response that consistently sorts substrate molecules, as well as operates on them, at a fixed level of generality. For these reasons, Pattee (e.g., 1982, p. 172) argues that enzymes can be said to *recognize* or *classify* molecules; enzymes have their own perspectives on which differences between molecules are significant and which are not (which determinables are relevant and which irrelevant, in Hautamäki's terminology)—perspectives on how to sort molecule tokens into types (viz., substrate and nonsubstrate).

¹⁴ My main reason for adopting Dennett's and Ladyman, Ross, and Collier's use of 'perspective' rather than Dretske's, Popper's, and Hautamäki's term 'point of view' here is to highlight that the key advantage that external perspectivalism has over forms of non-perspectivalism can be retained while dropping the dependence on mentality or cognition. Generally, the two terms are treated as synonymous. I would define 'perspective' as an isolated set of pattern-types which the system uses to implicitly carve its own world of interaction (cf. Varela, 1997). My use of 'perspective' is also similar to Devlin's (1991, p. 151) notion of a 'scheme of individuation'. Dennett's (and my) use of 'perspective' should not be confused with his use of 'stance'; a perspective counts as a "stance" for Dennett when it is adopted by minded beings as a "predictive strategy" (1981, p. 15).

An external perspective can be a way of sorting tokens into types, and so indexing the sorting of tokens into types to such an external perspective screens off alternative interpretations. But if Dretske and Pattee are right, then a theory of proper function need not make essential reference to human minds or sophisticated cognitive systems; it could instead ground the normativity of function in the classificatory or recognitional capacities of certain kinds of much simpler systems. But even given Pattee's claim that biological systems contain vast numbers of classification and recognition devices, it remains to be seen how these resources can provide a basis for norms of biological performance, i.e., proper function. In the next section, I discuss work by Haugeland that attempts to demonstrate how basic classification and recognition capacities can provide the basis for norms of performance.

4.2. *Selective Loss of Performance Detail: Haugeland and Censoriousness*

In his 1990 paper "The Intentionality All-Stars," Haugeland's goal is to provide an account of the normativity of intentionality. Though he is not providing an account of biological proper function per se, his theory is worth examining in this context because of the more general insights it can offer about what is needed to avoid generality problems about norms of performance.¹⁵

According to Haugeland, normative standards of performance can spontaneously emerge within a community of "versatile and interactive creatures, not otherwise specified except that they are conformists" (1990, p. 147). Haugeland discusses two forms such conformism might take, i.e., two mechanisms of interaction by which such "creatures" might participate in and sustain a community of shared performance norms. The first of these is *imitativeness*, i.e., the tendency to imitate the behavior of the other creatures. Turner (1994) argues that imitativeness alone cannot provide a grounding for the sharing of norms. This is because imitating is subject to interpretation: there are many ways to interpret another's behavior, and therefore many different ways to imitate it. Again, a generality problem arises: whether or not A's behavior is an imitation of B's behavior depends on the level of generality or grain at which the behaviors are typed, i.e., at which behaviors are considered as being *the same* for purposes of imitation:

What is needed to preserve the concept of sameness in the case of open-textured rules is ... some way of distinguishing variation of habits, or the acquisition of a variant rule, from the acquisition of the same open-textured rule.¹⁶ (Turner, 1994, p. 76)

In other words, what is needed is a shared scheme for sorting and discriminating the performances of others. Turner and Haugeland both argue that what is needed to fix a norm is not merely a tendency of imitativeness, but a form of what Haugeland refers to as

¹⁵ Haugeland (1990, p. 147) credits Heidegger, Sellars, and Brandom as forerunners of the position he is advancing.

¹⁶ This is similar to Popper's point I mentioned earlier, that "for logical reasons, there must always be a point of view ... before there can be any repetition; which point of view, consequently, cannot be merely the result of repetition" (2002, p. 59).

ensoriousness that fixes and categorizes the relevant behavior *types*. Turner notes that, for this reason, norms “are identifiable only by observing what happens if they are breached” (1994, p. 28). Haugeland says that by *ensoriousness*, he means “a positive tendency to see that one’s neighbors do likewise, and to suppress variation” (1990, p. 147). He does not only have negative kinds of *ensoriousness* in mind:

What in general counts as censorious acceptance and rejection? Whatever community members do that promotes conformism: if they smile at conforming performances, and smiles promote repetition, if they fire electric shocks at those who err, and these shocks discourage aberration, then smiles and shocks are devices of censorship. And who are the members of this community? The members are basically whoever is brought into conformity with the rest of the group, and thereupon participates in the censoring of others. (1990, p. 148)

Haugeland explains that such *ensoriousness* can form the grounding of a socially-based norm if it is based on a fixed practice of recognition and sorting of behavior.

Imagine, for instance, that the rules of chess were not explicitly codified, but observed only as a body of conformist norms—“how one acts” when one plays chess. So, it is proper (socially acceptable) to move the king in any direction, but only one square at a time. For this to be a norm, players and teacher/censors must be able to “tell” (that is, respond differentially to) which piece is the king, what are the squares, what counts as a move, and so on; thus, the presupposed sorting of circumstances is effectively a sorting of items, features, and events within those circumstances. Meanwhile, according to other norms, the king must be protected when threatened, cannot be exposed to capture, can castle under certain conditions, and so on; and, crucially, for all of these norms, it’s the same instituted sort (“king”) that’s involved. Hence, the norms themselves are interdependent via depending on the same sorting of circumstances (items, features, ...). (1990, pp. 151–152)

As long as the sorting and categorizing is done in the same way by participants, there is no generality problem; which piece is counted as a “king” is not left open to interpretation based on differing levels of grain. In order to be playing chess at all, in this community, one must first of all attend to the same level of specification or detail as all the other participants when counting something as a “king” or not. But Haugeland stresses that the categorization must not only happen at the level of recognizing and categorizing items, features, and events; it must also happen at the behavioral level, so that input stimuli are recognized and categorized and then mapped to a categorical response: “norms have a kind of ‘if-then’ character, connecting sorts of circumstance to sorts of behavior” (1990, p. 151).¹⁷ Haugeland argues that because conformism

¹⁷ Schroeder has argued that

Regulating involves creating a rule of some sort, as the word ‘regulate’ suggests. This is the sort of activity that can be expected to create norms, for a rule is just one sort of norm. Hence it should be no

(in his sense) is the type of thing that serves not only to determine the behavior of individuals in the community but also that sets the framework for *typing* circumstances, behaviors, and conditions that hold between them, as well as for discriminating deviations, conformism can provide the necessary grounding for performance norms:

When behavioral dispositions aggregate under the force of conformism, it isn't herds that coalesce, but *norms* ... distinct, enduring clusters of dispositions in behavioral feasibility space, separated in that space by clear gaps where there are no dispositions The community-wide classes of similar dispositions that coalesce under the force of conformism can be called "norms"—and not just collections or kinds—precisely because they themselves set the standard for that very censoriousness by which they are generated and maintained. The censure attendant on deviation automatically gives these standards (the extant classes themselves) a *de facto* normative force. (1990, p. 149)

Haugeland stresses that conformism, on his view, does not require the capacity for following explicit rules, but neither does it consist in mere causal regularity. Here, Haugeland trades on a distinction that was earlier mapped out by Sellars, who distinguished between "merely conforming to rules," "obeying rules," and a third category situated in between the first two that Sellars refers to as "pattern-governed behavior" (1954, pp. 326–327). Haugeland also argues that conformist norms should not be confused with conventional norms; on his view, the latter require mental states:

The difference between norms and conventions lies in this explanatory appeal: conformism does not presuppose any prior beliefs or preferences on the part of individual conformists, and hence the persistence of norms cannot be explained in terms of agents' interest maximization or rational choice. Indeed, norms need not be, or even seem to be, in any way beneficial either to the individuals or to the group; the mechanism of conformism is completely blind to the character or merits of the norms engendered.¹⁸ (1990, p. 150)

For these reasons, the information processing capacity of "creatures" constituting the conformist community can be extremely modest. They do not have to be human beings: they could be very simple organisms or machines—anything that has the capability to perform recognition and enact a rule-like (pattern-governed) mapping between recognized circumstances and behavioral response categories.

surprise that a regulated object is subject to a norm of performance: that it has a function. (2004, p. 118)

However, not all forms of regulation rely on fixed typing schemes that define *sorts* in Haugeland's sense. Simple forms of negative feedback, for example, are insufficient for normativity because it is a matter of interpretation how far the variable being controlled has to be from the set point until malfunction should be ascribed.

¹⁸ Stated another way: Conformist norms need not be *categorically binding* (in the sense of Copp, 2015), and need not be constrained by categorically binding norms.

4.3. *Division of Labor*

Whereas Haugeland's account depended on a community of homogenous "creatures" that are all monitoring and censuring each others' behavior, the entities doing the censuring (call these the "alpha creatures") might be distinct from those which are censured ("beta creatures"). Consider an example in which all of the creatures are robots in a fully automated factory (perhaps they were all designed by intelligent beings, or suppose that their existence and placement within the factory is a chance occurrence, like the parts of Dretske's "Twin Tercel," 1995, pp. 141ff). The beta robots are the workers. They perform the work in the factory, assembling and transporting components, and they each have a task that they usually perform successfully. But they sometimes develop problems causing them to fail at their tasks. For example, their internal components may become corroded or damaged, there may be lurking bugs in their software, or their parts may need recalibrating. Fortunately, the alpha robots are there to roam the factory and monitor the performance of the beta robots. When the beta robots are doing their jobs properly, the alphas simply continue vigilantly roaming through the factory. When they detect that a beta robot is not performing its job properly, the beta robot is censured. This may take many forms: perhaps taking it offline and sending it to a repair bay, or taking it offline and sending it to a scrapping or recycling area, or even sending a signal to the beta robot that stimulates a self-repair activity. Regardless, even in this simple situation, the automated factory may be said to realize a communal regime of "conformist norms" in Haugeland's sense, even if the information processing capabilities of the robots are very limited. Again, Haugeland stresses that in the minimal case,

conformism does not presuppose thought, reasoning, language, or any other cognitive faculty; the creatures do not in the first instance conform or censure wittingly or because they want to (except "tacitly")—they are simply built that way. (1990, p. 148)

Nonetheless, by enacting a specific framework for typing circumstances, behaviors, and conditions, as well as what count as deviations, and censuring these deviations, the censuring patterns of the alphas set a normative standard for the factory as a whole. Given the background of such norms of performance that are put into practice by the alphas, the beta robots in the factory may be said to have proper functions; the standard for whether a given beta robot is functioning properly is set by the censuring patterns of the alphas.

Note that by putting a certain set of standards into practice in the factory, the alphas may be said to collectively *have a perspective*: their pattern of censuring constitutes a tacit perspective on which differences in performance are significant and which ones insignificant. An external observer may take a different perspective and decide that the way the alphas work is not *objectively* best for the factory. The perspective possessed by the alphas may not be "objective" in this sense¹⁹ (as we saw, due to generality problems, there probably cannot be a theory of

¹⁹ Or, again in Copp's (2015) terminology, the performance norms grounded by the censuring patterns of the alphas are not categorically binding.

what counts as proper function in this non-perspectival sense), but it is objective in the sense that it is independent of any *external* observer. The system in question, by itself, realizes a normative standard. The collective activity of alpha robots also provides an objective criterion for delineating what is part of the system and what is not.

Such norms of proper function are not etiological or history-dependent,²⁰ and they are not dependent on the perspective(s) taken by external observers. They depend on the level of generality at which certain circumstances, behaviors, and conditions are typed, but instead of “starting with the intuitively correct judgment about a [trait’s proper function] and then cherry picking an account of [the organism’s selective regime, selection history, possible worlds, regime of self-maintenance, the functional characteristics of the larger system, hierarchical levels, etc.] that yields that judgment” (again paraphrasing Bishop), it is the alpha robots themselves that *fix* the relevant level of generality. For this reason, by grounding proper function in the censuring activities of alpha robots, the generality problems can be avoided. In the next section, I discuss how the theory applies to real biological systems.

5. Performance-Monitoring and Censuring Mechanisms in Subcellular Biology

The cells of biological organisms, including even single-celled organisms like bacteria, contain highly sophisticated molecular machinery. The everyday functioning of any given cell involves many different kinds of chemical reactions that are tightly orchestrated and coordinated in ways that biologists are only beginning to understand. There are many ways that such orchestration and coordination can (and does) fail, and like the automated factory, cells are highly dependent on certain components (proteins) that are capable of discriminating whether certain other components are doing their jobs properly. As discussed in Section 4.1, the enzymatic binding and catalysis by which proteins perform their work also counts as a fixed practice of recognition and categorical response. Just as chess players consistently categorize chess pieces at a fixed level of generality, enzymes sort substrate molecules at a similarly fixed level of generality. These quality control proteins can trigger various types of responses, ranging from throttling down upstream processes to reduce the malfunctioning component’s workload, to triggering a process called *programmed cell death*, in which the cell itself commits suicide in an orderly fashion to prevent the malfunctioning cell from causing further problems for the larger organism.

²⁰ Garson (2019) has argued that because they tend to rely on statistical norms, even accounts that do not define function explicitly in terms of selection history should yet be classified as history-dependent. Performance norms such as those described in this section, however, are not history-dependent because such norms derive from the recognition and response behaviors of the alphas, which are in turn realized by their own mechanistic constitution. The norms fall out of the fact that, as Haugeland put it in the above quotation, the alphas “are simply built that way.” Performance norms in Haugeland’s sense are therefore not statistical norms. They are recognition and response norms that are realized by the internal dynamical properties of the alphas. There is no difficulty for “swamp man” scenarios here; if alpha creatures spontaneously appear in a system and begin enforcing performance norms, those norms will apply to the system immediately.

One example of this is the *unfolded protein response* that is triggered by several different types of “performance monitoring” proteins in the endoplasmic reticulum (ER). An important task that many cells perform is to synthesize proteins that will be excreted from the cell to perform functions elsewhere. These proteins (among others) are synthesized in ribosomes that are embedded within the ER. Protein synthesis involves stringing together component amino acid molecules, and requires the string of amino acids to be folded into the right shape (important because the capacity of proteins to do work is largely a matter of how they are shaped). The folding process can occur spontaneously during synthesis, but is sometimes also assisted by helper proteins called *chaperones*. Problems may occur during the folding process, whereby a non-folded or misfolded protein results, or a properly folded protein can also become unfolded or misfolded as a result of various conditions in the cell, like excessive heat or pH levels (the unfolding of a protein is sometimes called *denaturing*). There are several types of proteins in the ER that are able to detect unfolded or misfolded proteins, including Ire1, ATF6, and PERK (Galluzzi, Bravo-San Pedro, & Kroemer, 2014).

When too many unfolded proteins are detected by such “sensor” proteins, several types of responses may be triggered. A signal may be sent to the nucleus of the cell causing it to produce more of the chaperone molecules or to enlarge the ER to increase its processing capacity. A signal may also be sent that causes the upstream processes of protein synthesis to be slowed down (to give the ER a chance to play catch-up, and hopefully reduce the number of unfolded or misfolded proteins). These signals may also halt the progression of the cell cycle, preventing the cell from making further progress towards mitosis, or cell division (since if this cell is malfunctioning, then the resulting cells may not be properly formed). When conditions are bad enough, these sensors can even send a signal that triggers *mitochondrial membrane permeabilization* (MMP), in which “executioner” enzymes are released from the outer membrane of the mitochondria. These executioner enzymes, called caspases, are like a controlled demolition team that systematically degrades the components of the cell, causing it to self-terminate without provoking an immune response or causing problems in the surrounding tissues (Hotchkiss, Strasser, McDunn, & Swanson, 2009).

Other examples of malfunction detection and response have been found in other organelles, including the mitochondria, nucleus, Golgi apparatus, and lysosomes (Galluzzi, Bravo-San Pedro, & Kroemer, 2014). For example, the SOK1 sensor protein in the Golgi apparatus detects malfunction by responding to elevated levels of *reactive oxygen species* (ROS), which are a potentially dangerous byproduct of certain intracellular reactions (Nogueira et al., 2008). Instead of triggering programmed cell death, another way that cells respond to malfunction is by causing the orderly disposal of malfunctioning organelles in a process called *autophagy* (Okamoto, Kondo-Okamoto, & Ohsumi, 2009; Youle & Narendra, 2011).

Even when the damage repair mechanism is not known, its existence can be ascertained. For several decades, biologists have known that when filaments are sheared off of flagella, bacteria are able to replace them. Kelly T. Hughes and colleagues found that flagella can regenerate

from the same hook/basal body of a flagellum that was sheared off. Through clever experiments they determined that a separate mechanism for *repairing* malfunctioning flagella must exist, rather than simply retriggering the same mechanism that initially builds the flagellum, although this separate mechanism is not yet known (Rosu & Hughes, 2006).

These examples have several features in common. In each case, there is some situation indicative of performance which is either registered as present or not present (e.g., malfunctioning flagellum, too many unfolded/misfolded proteins, too many ROS molecules being produced/too much oxidative damage, malfunctioning or superfluous mitochondria), triggering an unequivocal all-or-nothing type of response (e.g., flagellar filament replacement, autophagy, halting the cell cycle, apoptosis). In other words, the recognition and typing of a performance token, an implicit if-then mapping from this recognized performance type to a categorical censuring response, and the context-dependent execution of that response, all the ingredients needed for a community of “creatures” to realize norms of performance based on censorious acceptance and rejection. Since what is being categorized is the performance of some subsystem or component, each of these processes can ground a distinction between functioning and malfunctioning for that subsystem/component.

6. Macro-Level Functions: Composite Recognition and Response

In the previous section, the focus was on individual protein molecules that act as “alpha creatures” at the subcellular level. But it might seem that there is no equivalent at the macroscopic level that can ground macro-level proper functions, such as the mammalian heart’s function of pumping blood. After all, there is no individual system in mammals we can point to that acts as an alpha creature, measuring overall performance of the heart and “censuring” the heart when it malfunctions.²¹

Recall from Section 4.1 that what is special about enzymes that allows them to act as alpha creatures is their selective sensitivity to *types* (chemical species) at a specific level of detail: differences in conformation between substrates of the right chemical species for that enzyme are ignored, so that enzymatic reactions implement the kind of fixed practice of recognition (in the Dretskean sense) and categorical response that can ground a normative perspective.²² While such enzymatic activity is ubiquitous in biological systems (Section 5), similar detection and censoring systems above the cellular level may not be nearly as common. However, I do

²¹ The hypothalamus plays a key role in regulating heart rate, but this is not quite the same thing: the hypothalamus only makes adjustments to normal heart functioning. It has no way to detect and “censure” the heart based on macro-level conditions of heart malfunction, like congestive heart failure or cardiac arrest.

²² Since it explains the normative status of proper functions in terms of such non-normative facts as the special dynamical properties of enzymes discussed by Pattee, as well as how they are situated within cells to realize such regimes of censorious recognition and response, the present view might be considered a form of reductionism.

not believe that this shows that most of the macro-level functions we attribute to biological traits do not exist. In this section I argue that instead of arising from the activity of individual, macro-scale alpha creatures, most macro-level proper functions emerge from the collective activity of micro-level alpha creatures.

To understand how macro-level proper functions (e.g., of an organ like the heart) can arise from the activity of the kinds of micro-level alpha creatures described in the previous section, begin again with a simple example. Suppose that there is a device that detects whether there is a '1' in a certain position, and another device that detects whether there is a '5' in the position next to it. Each device responds categorically in one of two ways depending on whether the condition in question is met. But when you consider both devices together, there are four conditions that the devices are *collectively*—not singly—responsive to: present 1 and absent 5, present 1 and present 5, absent 1 and present 5, and absent 1 and absent 5. There are similarly four *collective* response behaviors, one for each of these conditions. In this way, we can instead look at the two detection devices as comprising a single, more complex detection device.²³

Now suppose that the absence of a 1, in the case of the first detector, or the absence of a 5, in the case of the second detector, results in negative censuring behaviors (perhaps they are monitoring and censuring the mechanisms responsible for producing the 1s and 5s). The detectors are then each a type of alpha creature. If so, then it equally makes sense to consider the combination of the two detectors as a single, more complex alpha creature that monitors whether the sequence (1, 5) is present. It is monitoring the larger system that includes the subsystems that produce the 1s and 5s, and enforcing a standard in which the production of (1, 5) sequences constitutes proper functioning, and other conditions constitute varying degrees of malfunction.

The larger point here is that the activity of many alpha creatures can collectively realize a higher-level, and more complex, standard of functioning—not due to the fact that the higher-level censuring is causally dependent on the low-level censuring, but because it is *constituted* by it. The higher-level response is categorical and involves a selective loss of detail precisely because the lower-level responses are categorical and involve selective losses of detail. The detail that is lost at the higher level *just is* the sum of the detail that is lost at the lower level. As long as the lower-level alpha creatures operate with a fixed level of generality, any normative censuring standard that is additively built up from them will also operate with a fixed level of generality. Ultimately, activity that passes muster for many different types of alpha creatures at the cellular level *just is* activity at the macro scale when you add it all together. The alpha creatures collectively define a standard of what passes muster at the macro-level, and therefore define what proper functioning consists in (when no censuring activity is triggered) or what varying degrees of malfunction consist in (when varying amounts of negative censuring activity is triggered). Just as in the simple initial example, the function of the combined subsystems was

²³ In fact, this is the basis of parallel distributed (Selfridge & Neisser, 1960) and agent-based (Rosin & Rana, 2004) pattern detection techniques.

to produce sequences of (1, 5), the function of the heart is to pump blood because that is the collective activity that the detectors are collectively responsive to. The detectors, considered collectively, realize a perspective on what counts as proper functioning for the heart as a whole. This is the reason why we don't find a single, non-composite macro-level alpha creature that monitors the heart as a whole: such a macro-level alpha creature is not necessary and would be redundant, because that role is already fulfilled by the collective, finely orchestrated and emergent self-organizing activity of lower-level alpha creatures.

Go back to the automated factory example, and suppose it is a car factory. It might be the case that there isn't any particular alpha robot whose job it is to monitor the factory as a whole and detect whether it is producing entire cars properly or not. This might be because the operations of the individual alpha robots are so well orchestrated (without there needing to be a non-composite "orchestrator") that their activity collectively adds up to the implementation of a standard-verifying mechanism for the whole factory that works as well as a single, holistically functioning alpha robot would. In this way, an organ like the heart is like the car factory.²⁴

Space does not here allow a detailed critical analysis of this way of grounding macro-level proper functions. However, since on this account, proper function is always derived from the detection of individual enzymatic detectors (and other Dretske-style recognition devices that might exist in biological systems), an objection might run as follows: Even at the level of subcellular detection enzymes, it is a matter of interpretation whether (in the case of the enzyme that detects mitochondrial damage) increased ROS is being detected, or whether incomplete oxidative phosphorylation is being detected. This is similar to the argument that there is no determinate representational content in frog vision when a frog is detecting a fly, because there are equally good reasons to think that what is being registered by the frog's visual system is the presence of 'food' or the presence of 'small dark moving things' (Goode & Griffiths, 1995, p. 100).

However, this objection conflates two different types of problems. One pertains to the question of whether there exists a determinate category *for the frog* (i.e., by the frog's "lights," or according to its own perspective) of what is being detected. The other is the problem of whether there exists a univocal translation of the frog's categories into the ordinary language of humans.²⁵ The latter problem is an important one to solve if one wishes to ground the full

²⁴ Unfortunately it isn't possible to go into empirically detailed concrete biological examples of this, simply because we don't know in detail about how the activities of thousands of "alpha creature" individual proteins collectively sustain the organization of a system like the heart. It would require reverse-engineering the heart in *total* molecular detail, all the different genes that get expressed under what circumstances (and exactly what molecular pathways they are triggered by) in all the different types of cells of the heart and what their intercellular and intracellular effects are, and all of the non-linear interactions of those effects, etc., something we cannot yet approach. I don't, however, believe that such an exhaustive level of knowledge is necessary for scientists to competently discover and investigate macro-level proper functions, just as knowledge of molecular structure is not necessary for competently discovering and classifying crystals (Polányi, 1958, pp. 43–48).

²⁵ For a thorough discussion of this type of conflation, see Evans (1975).

richness of representational content of human cognition in basic detection capabilities, which is certainly not being attempted in this paper. There may very well be no fact of the matter whether we should characterize what the frog is detecting as ‘food’, ‘flies’, or ‘small dark moving things’. But it may be true at the same time that since the frog has no cognitive means of distinguishing between these situations, and only has a single category corresponding to its more basic recognition capacity, no corresponding ambiguity exists from the point of view of its own content-assignment scheme (cf. Dennett, 1996, pp. 42–43). The categorical discrimination necessary for such ambiguity to arise outstrips the frog’s own classificatory repertoire.

Similarly, even if we can conceptualize the performance-monitoring enzyme’s simple classificatory repertoire in different ways, this does not mean there will be a corresponding indeterminacy in the proper function grounded by the enzyme’s censuring activity. The level of generality that the enzyme works with is fixed, and is simply more coarse-grained than the conceptual discriminations that we ordinary language users have at our disposal.

We can now return to the mouse/hawk scenario discussed in Section 3.1. The proper function of the mouse’s auditory system would be determined by what kind of macro-level standard of mechanistic operation all of the micro-level enforcement ultimately adds up to. It may add up to a particular kind of auditory classification scheme that classifies auditory inputs into categories and generates responses. That scheme may or may not include types that directly map onto the situation types that were listed in Section 3.1, and may or may not have a univocal translation into the ordinary languages of human beings—but regardless, there will be some determinate norm of performance for the auditory system that is enforced internally.²⁶

7. Conclusion

Because performance-monitoring subsystems that recognize and respond to malfunctions are widespread in biological systems, all biological systems operate with their own implicit perspectives on what constitutes proper functioning. These implicit, operational perspectives can provide definite content to a theory of proper function because they amount to a fixed scheme for sorting performance tokens into types. The *proper function* of a given biological subsystem consists in the activities undertaken by that subsystem that are enforced by performance monitoring detection and response systems. This is an internal perspectivist account of proper function: what counts as the proper function of a biological subsystem or trait is a matter of the perspective that the biological system, itself, has on what counts as proper functioning for that subsystem.

One key problem faced by major existing positions in the biological function debate is that they assign functions to traits based on the type that some related token (system, selection history,

²⁶ This response was not available to the etiological theorist because again, natural selection does not involve a scheme of typing and sorting at a fixed level of generality. See Footnotes 6 and 13.

self-maintenance regime, etc.) falls under, but the token in question falls under many such types and this gives rise to conflicting function assignments. As we saw, this is very similar to a problem that has been faced by process reliabilism. What was needed was a theory of proper functions that includes sufficient resources to say what determines the relevant level of type generality in given cases. Internal perspectivalism is the theory of biological function that can meet this key desideratum. The question of when a mechanism is functioning or malfunctioning is one that can be answered by looking at that mechanism in the larger context of the organism and its own quality control “alpha creatures,” which are ubiquitous in biological systems. The larger system that includes the mechanism in question and the quality control “alpha creatures” can ground this type of normativity since their recognition and response patterns will operate with an implicit but definite categorization scheme that isolates a specific level of generality. The present account therefore provides the way to naturalize proper function without reference to natural selection, natural design, or fitness, and without resorting to external perspectivalism.

Acknowledgements

I am indebted to William Bechtel, Mark Bedau, Andrew Bollhagen, Dan Burnston, Rick Grush, Rebecca Hardesty, Jay Odenbaugh, Ben Sheredos, participants of the 2017 Active Matter Workshop at Georgetown University, and two anonymous reviewers for this journal for valuable discussions and comments on previous drafts.

References

- Beatty, J. (1984). Chance and Natural Selection. *Philosophy of Science*, 51, 183–211.
- Bechtel, W. (1986). Teleological Functional Analyses and the Hierarchical Organization of Nature. In N. Rescher (Ed.), *Current Issues in Teleology* (pp. 26–48). Landham, MD: University Press of America.
- Bedau, M. (1992). Goal-Directed Systems and the Good. *Monist*, 75(1), 34–51.
- Bigelow, J., & Pargetter, R. (1987). Functions. *Journal of Philosophy*, 84(4), 181–196.
- Bishop, M. A. (2010). Why the Generality Problem is Everybody's Problem. *Philosophical Studies*, 151, 285–298.
- Boorse, C. (1997). A Rebuttal on Health. In J. M. Humber & R. F. Almeder (Eds.), *What Is Disease?* (pp. 3–134). Totowa, NJ: Humana.
- Burge, T. (2009). Primitive Agency and Natural Norms. *Philosophy and Phenomenological Research*, 79(2), 251–278.
- Conee, E., & Feldman, R. (1998). The Generality Problem for Reliabilism. *Philosophical Studies*, 89, 1–29.
- Copp, D. (2015). Explaining Normativity. *Proceedings and Addresses of the American Philosophical Association*, 89, 48–73.
- Craver, C. F. (2001). Role Functions, Mechanisms, and Hierarchy. *Philosophy of Science*, 68(1), 53–74.
- Craver, C. F. (2013). Functions and Mechanisms: A Perspectivalist View. In P. Huneman (Ed.), *Functions: Selection and Mechanisms* (pp. 133–158). Dordrecht: Springer.
- Cummins, R. (1975). Functional Analysis. *Journal of Philosophy*, 72(20), 741–765.

- Darwin, C. (1876). *On the Origin of Species* (6th ed.). New York: Appleton.
- Davies, P. S. (2001). *Norms of Nature*. Cambridge, MA: Bradford.
- Dennett, D. C. (1981). True Believers: The Intentional Strategy and Why It Works. Reprinted in his *The Intentional Stance* (pp. 13–35). Cambridge, MA: Bradford, 1987.
- Dennett, D. C. (1991). Real Patterns. *Journal of Philosophy*, 88(1), 27–51.
- Dennett, D. C. (1996). *Kinds of Minds: Toward an Understanding of Consciousness*. New York: Basic.
- Devlin, K. (1991). *Logic and Information*. Cambridge: Cambridge University Press.
- Dretske, F. I. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: Bradford.
- Dretske, F. I. (1988). *Explaining Behavior*. Cambridge, MA: Bradford.
- Dretske, F. I. (1995). *Naturalizing the Mind*. Cambridge, MA: Bradford.
- Enç, B. (2002). Indeterminacy of Function Attributions. In A. Ariew, R. Cummins, & M. Perlman (Eds.), *Functions: New Essays in the Philosophy of Psychology and Biology* (pp. 291–313). Oxford: Oxford University Press.
- Evans, G. (1975). Identity and Predication. *Journal of Philosophy*, 72(13), 343–363.
- Foot, P. (2001). *Natural Goodness*. Oxford: Clarendon.
- Galluzzi, L., Bravo-San Pedro, J. M., & Kroemer, G. (2014). Organelle-Specific Initiation of Cell Death. *Nature Cell Biology*, 16(6), 728–736.
- Garson, J. (2013). The Functional Sense of Mechanism. *Philosophy of Science*, 80, 317–333.
- Garson, J. (2019). There Are No Ahistorical Theories of Function. *Philosophy of Science*, 86(5), 1146–1156.
- Goldman, A. I. (1979). What is Justified Belief? In G. S. Pappas (Ed.), *Justification and Knowledge: New Studies in Epistemology* (pp. 1–23). Dordrecht: Reidel.
- Goode, R., & Griffiths, P. E. (1995). The Misuse of Sober's Selection for/Selection of Distinction. *Biology and Philosophy*, 10, 99–108.
- Haugeland, J. (1990). The Intentionality All-Stars. Reprinted in his *Having Thought: Essays in the Metaphysics of Mind* (pp. 127–170). Cambridge, MA: Harvard University Press, 1998.
- Hautamäki, A. (1986). *Points of View and Their Logical Analysis*. Helsinki: Philosophical Society of Finland.
- Hotchkiss, R. S., Strasser, A., McDunn, J. E., & Swanson, P. E. (2009). Cell Death in Disease: Mechanisms and Emerging Therapeutic Concepts. *New England Journal of Medicine*, 361(16), 1570–1583.
- Maley, C. J., & Piccinini, G. (2017). A Unified Mechanistic Account of Teleological Functions for Psychology and Neuroscience. In D. M. Kaplan (Ed.), *Explanation and Integration in Mind and Brain Science* (pp. 236–256). Oxford: Oxford University Press.
- Massimi, M. (2018). Perspectivism. In J. Saatsi (Ed.), *The Routledge Handbook of Scientific Realism* (pp. 164–175). New York: Routledge.
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Moreno, A., & Mossio, M. (2015). *Biological Autonomy*. Dordrecht: Springer.
- Nagel, E. (1977). Functional Explanations in Biology. *Journal of Philosophy*, 74(5), 280–301.
- Neander, K. (1995). Misrepresenting & Malfunctioning. *Philosophical Studies*, 79, 109–141.
- Neander, K., & Rosenberg, A. (2012). Solving the Circularity Problem for Functions: A Response to Nanay. *Journal of Philosophy*, 109(10), 613–622.
- Nogueira, E., Fidalgo, M., Molnar, A., Kyriakis, J., Force, T., Zalvide, J., & Pombo, C. M. (2008). SOK1 Translocates from the Golgi to the Nucleus upon Chemical Anoxia and Induces Apoptotic Cell Death. *Journal of Biological Chemistry*, 283(23), 16248–16258.
- Nyberg, I. (2009). Can Moral Norms Be Derived from Nature? The Incompatibility of Natural Scientific Investigation and Moral Norm Generation. In M. J. Cherry (Ed.), *The*

- Normativity of the Natural: Human Goods, Human Virtues, and Human Flourishing* (pp. 175–196). Dordrecht: Springer.
- Okamoto, K., Kondo-Okamoto, N., & Ohsumi, Y. (2009). Mitochondria-Anchored Receptor Atg32 Mediates Degradation of Mitochondria via Selective Autophagy. *Developmental Cell*, 17, 87–97.
- Okasha, S. (2009). Causation in Biology. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *The Oxford Handbook of Causation* (pp. 707–725). Oxford: Oxford University Press.
- Pattee, H. H. (1970). The Problem of Biological Hierarchy. In C. H. Waddington (Ed.), *Towards a Theoretical Biology 3: Drafts* (pp. 117–136). Edinburgh: Edinburgh University Press.
- Pattee, H. H. (1973). Physical Problems of the Origin of Natural Controls. In A. Locker (Ed.), *Biogenesis, Evolution, Homeostasis* (pp. 41–49). Berlin: Springer-Verlag.
- Pattee, H. H. (1977). Dynamic and Linguistic Modes of Complex Systems. *International Journal of General Systems*, 3, 259–266.
- Pattee, H. H. (1982). Cell Psychology: An Evolutionary Approach to the Symbol-Matter Problem. Reprinted in H. H. Pattee & J. Rączaszek-Leonardi (Eds.), *Laws, Language and Life* (pp. 165–179). Dordrecht: Springer, 2012.
- Pattee, H. H. (1992). The Measurement Problem in Physics, Computation, and Brain Theories. In M. E. Carvallo (Ed.), *Nature, Cognition and System II* (pp. 179–192). Dordrecht: Kluwer.
- Piccinini, G. (2015). *Physical Computation*. Oxford: Oxford University Press.
- Polányi, M. (1958). *Personal Knowledge*. New York: Harper Torchbooks.
- Popper, K. (2002). *Conjectures and Refutations*. London: Routledge.
- Preston, B. (1998). Why is a Wing Like a Spoon? A Pluralist Theory of Function. *Journal of Philosophy*, 95(5), 215–254.
- Rosin, P. L., & Rana, O. F. (2004). Agent-Based Computer Vision. *Pattern Recognition*, 37, 627–629.
- Ross, D., Ladyman, J., & Collier, J. (2007). Rainforest Realism and the Unity of Science. In J. Ladyman & D. Ross, *Every Thing Must Go: Metaphysics Naturalized* (pp. 190–257). Oxford: Oxford University Press.
- Rosu, V., & Hughes, K. T. (2006). σ^{28} -Dependent Transcription in *Salmonella enterica* is Independent of Flagellar Shearing. *Journal of Bacteriology*, 188(14), 5196–5203.
- Schroeder, T. (2004). Functions from Regulation. *The Monist*, 87(1), 115–135.
- Selfridge, O. G., & Neisser, U. (1960). Pattern Recognition by Machine. *Scientific American*, 203(2), 60–68.
- Sellars, W. (1954). Some Reflections on Language Games. Reprinted in his *Science, Perception, and Reality* (pp. 321–358). Atascadero, CA: Ridgeview, 1963.
- Sinnott, E. W. (1961). *Cell and Psyche: The Biology of Purpose*. New York: Harper & Row.
- Sosa, E. (1991). Intellectual Virtue in Perspective. In his *Knowledge in Perspective: Selected Essays in Epistemology* (pp. 270–293). Cambridge: Cambridge University Press.
- Sterelny, K. (1995). Basic Minds. *Philosophical Perspectives*, 9, 251–270.
- Sullivan-Bissett, E. (2017). Biological Function and Epistemic Normativity. *Philosophical Explorations*, 20(sup1), 94–110.
- Turner, S. P. (1994). *The Social Theory of Practices: Tradition, Tacit Knowledge and Presuppositions*. Cambridge: Polity.
- Varela, F. J. (1997). Patterns of Life: Intertwining Identity and Cognition. *Brain and Cognition*, 34, 72–87.
- Walsh, D. M. (1996). Fitness and Function. *British Journal for the Philosophy of Science*, 47, 553–574.
- Woodger, J. H. (1929). *Biological Principles*. London: Routledge & Kegan Paul.
- Youle, R. J., & Narendra, D. P. (2011). Mechanisms of Mitophagy. *Nature Reviews Molecular Cell Biology*, 12, 9–14.