Mechanistic Explanations and Teleological Functions

Andrew Rubner

This article defines and defends a notion of teleological function that is fit to figure in explanations concerning how organic systems, and the items that compose them, are able to perform certain activities, such as surviving and reproducing or pumping blood. According to this notion, a teleological function of an item (such as the heart) is a typical way that items of that type contribute to some containing system's ability to do some activity. An account of what it is for an item to contribute to a containing system's ability to perform an activity is provided. I argue that the view acquires its normative status in virtue of the fact that it obeys a function–accident distinction and obeys a function–dysfunction distinction; that the view is ahistorical; and, that its ahistoricity provides it with an advantage over one of its main competitors.

1. Introduction

Life scientists are often interested in how questions. For example, a vision scientist may ask how a perceiver is able to see what is where, or a physiologist may ask how an organism is able to maintain itself, or an ecologist may ask how a predator is able to capture its prey. These kinds of questions take a particular form:

Q: How is a system S able to ϕ?

Here, ϕ is some activity such as capturing prey or pumping blood.[1] In many cases, questions of this form are answered by appeal to the activities of the items that constitute the system in question. Of course, it cannot be any activity of such constituent items, it must be those activities that contribute to the realization of the relevant activity of the containing system. For example, when explaining how the circulatory system is able to transmit nutrients and other substances to different parts of the body, one would want to cite the heart's ability to pump blood and not its ability to make a thump-thump noise, since only the former—from a physiological perspective—contributes to the circulatory system's ability to transmit nutrients

---

[1] Greek letters will stand in for activities throughout.

and so on to other parts of the body. We may say that in answering Q we are interested in citing the functions of the constituents of the relevant system.

Putting everything together, we can say that life scientists are interested in a particular kind of explanation-how that I will refer to as a ‘mechanistic explanation’. They may be more carefully characterized as follows:

Mechanistic Explanation (ME):

> Explanandum: How is a system S able to ϕ?
>
> Explanans: Because there exists n items $X_1, ..., X_n$ such that each X_i is a part or process of S, and for each X_i there is some j such that $j \geq 1$ and X_i has functions $\psi^i_1, ..., \psi^i_j$, and if each X_i were to fulfil each of its functions, then S would ϕ.

We say that X_i is a part or process of S to highlight the fact that S may be able to ϕ by carrying out some process that is not necessarily a proper part of S (Cummins [1975], p. 762; Cummins [1983], pp. 29–30).

This raises the question of what notion of function is operative in the *explanans* of mechanistic explanations. One live option—as suggested and argued for by Neander ([2017])—is that the relevant notion of function is that of a teleological function. Teleological functions have a normative aspect—they are supposed to capture what an item is supposed to do. To make the characterization of this normative aspect tractable, following Neander ([2017]), we can say that a notion of function is that of a teleological function if, but only if, the notion allows for dysfunctional items, and it obeys a function–accident distinction.[2] If a notion of function allows for dysfunctional items, then, for example, it will make good on the claim that a token item can have the function to ϕ without being capable of ϕ-ing. And, if a notion of function allows for a function–accident distinction it will correctly distinguish between those activities of an item that are its functions and those activities that may be beneficial for the relevant system in a certain class of circumstances. For example, such a notion will correctly predict that a function of the heart is to pump blood, but not to diagnose illnesses even though an irregular heart beat can be used to diagnose an illness that can in turn save a patient’s life. The idea is that such a contribution is accidental whereas the contributions made when the heart pumps blood are systematic, and thus only pumping blood should count as something the heart is supposed to do.

This article explores the idea that the notion of function operative in mechanistic explanations is that of a teleological function (rather than some other kind of function). I do so by putting forward and developing a notion of function that is both fit to figure in the *explanans* of mechanistic explanations and is a teleological function. Subtleties and qualifications aside,

---

[2] Neander ([2017]) calls these functions ‘normal-proper functions’. See also (Rubner [2022]).

according to the account developed here, a function of an item is a typical way that items of that type contribute to some containing system's ability to do some activity. One advantage of the view, or so I will argue, is that it preserves a certain amount of ahistoricity. That is, unlike competitor approaches—such as the one found in Neander ([2017])—there is an important and straightforward sense in which the view is ahistorical.[3] Moreover, the view that emerges may be considered a variant of views that attempt to understand teleological function in terms of statistical normality (Boorse [1977], [1997]; Garson and Piccinini [2014]; Maley and Piccinini [2017]; Piccinini [2020], chap. 3) insofar as it makes appeal to a notion of typicality. It is argued that while the view developed here obeys the function–accident distinction, its main competitor in this camp, as represented by Garson, Maley, and Piccinini (Garson and Piccinini [2014]; Maley and Piccinini [2017]; Piccinini [2020], chap. 3), fails to do so. Moreover, the present article provides an account of contribution. This not only adds to the family of views that understand teleological function in terms of statistical normality, but also to other views of function that can be put to work in the context of mechanistic explanation, such as those of Cummins [1975] and Craver ([2001], [2013]).[4] This is because views in both families often make appeal to the notion of a contribution without providing a full account of that notion.

Section 2 engages in two background discussions, one concerning the type-individuation of items and another containing some further detail on mechanistic explanations. In section 3, I define and defend my preferred notion of teleological function. An account of what it is for an item to make a contribution to a containing system is provided. In section 4, I clarify and defend the view's claim to ahistoricity, and argue that it provides the view with an advantage over others that are historical.

2. Some Background

It is sometimes claimed that teleological functions are ascribed to item types, for example, the heart, and that item tokens, for example, Fido's heart, have them in virtue of being tokens of the relevant type (Boorse [1977]; Neander [2017]). This is primarily because such accounts want to maintain that whether a token item, X, is dysfunctional or functional depends on whether it can fulfil the function ascribed to the type of item that X is a token. If this is right, then theories of teleological function need to be accompanied by a story of how items are type individuated.

Following Amundson and Lauder ([1994]), I will maintain that items are type individuated by a morphological or structural criterion. Leaving gross anatomical structure aside, a mor-

---

[3] The view will thus be of interest to philosophers of mind and epistemologists who appeal to teleological functions but are dissatisfied with an aetiological account of those functions (see, for example, Schellenberg [2018], pp. 35–36). Moreover, I will not be able to compare the view developed here to the organizational accounts of Schlosser ([1998]), McLaughlin ([2001]), Mossio et al. ([2009]), and Saborido et al. ([2011]), which are sometimes understood as being neither aetiological nor statistical. I also won't address the 'analysis versus explication' problem raised by Griffiths and Matthewson ([2018]). I take it that a proper treatment of that problem would require a metaphilosophical discussion that would take us too afar afield.

[4] See Craver ([2001]) for helpful discussion of Cummins's view of function in relation to mechanistic explanation.

phological criterion for type individuating items may include examination of the microscopic structure of tissues (unveiled by histological methods), developmental origins, and the principle of connection—which states that items 'which have identical connections or position within an overall structure pattern' are identical (Amundson and Lauder [1994], p. 455). To be clear, it would be a mistake to understand this view as saying that only one of these criterion is relevant for any particular identification. As Neander ([2002], p. 410) rightly points out, there may be cases in which cardiac tissue appears in cancerous growths that occur in odd places in the body. Presumably, however, those cancerous tissues could be ruled out as being hearts by some other morphological criterion. So, the idea is rather that there is a cluster of morphological methods (those described above) for distinguishing Xs from non-Xs and this provides sufficient reason to hold on to the view that items can be type individuated by some structural or morphological criteria. Importantly, then, items are not type individuated by appeal to their functions on this account.[5]

When we begin a mechanistic explanation, as described in ME, we can substitute S with any type of organic system (that is, any species) that displays a certain kind of organizational complexity: humans, boney fish, birds and so on. Such systems are organizationally complex insofar as they are composed of interlocking subsystems that are themselves decomposable into parts and processes that may or may not themselves constitute subsystems—collectively, we can refer to all such subsystems, parts and processes as 'items' or 'sub-items' (compare Wimsatt [1997]; Craver [2001]). Based upon this process of decomposition, such systems display different levels of organization. For example, textbooks standardly describe humans as dividing into six distinct levels of organization: the organismal level—the highest level of organization—the organ system level, the organ level, the tissue level, the cellular level, and the chemical level (McKinley et al. [2016], pp. 6–7).

Mechanistic explanations feature both a system and a certain activity that that system is able to do—that is, a capacity, or ability, of that system—which the variable ϕ stands in for in (ME). Note further that each X_i featured in the *explanans* of a mechanistic explanation has some function ψ^i_j. By replacing any X_i with S and simultaneously replacing ϕ with a corresponding ψ^i_j another *explanandum* and a corresponding *explanans* is generated, and due to this, a sub-mechanistic explanation is generated. By repeating this process until no further decomposition is possible, a series of sub-mechanistic explanations are generated. It will be helpful to call non-sub-mechanistic explanations 'overarching mechanistic explanations'.

---

[5] Some read Amundson and Lauder ([1994]) differently on this point. See, for example, Griffiths ([2006], p. 17). However, Amundson and Lauder ([1994], p. 453) do write '[selected effect] functions are not the foundation for the classification of biological traits. To be sure, [causal role] function does not define basic categories either. The classifications come from a third, non-functional source'. The third, non-functional source is the structural, morphological criteria discussed above. However, it is possible that Lauder ([1999]) would consider some dispositional properties of an item as relevant to their classification. For example, a pattern of muscle movement (Lauder [1999], p. 186). Furthermore, it is sometimes argued that it is not possible to type individuate items by morphological or structural criteria alone (Neander, [1991]; Nanay, [2010]). But see Amundson and Lauder ([1994], p. 457) for a reply.

Note that there is both a system and a capacity featured in an overarching mechanistic explanation. This overarching capacity is not chosen arbitrarily, but is rather reflected in some central explanatory project of some life science. For example, where we are concerned with whole, individual organisms, it is plausible that at least one overarching capacity is the capacity for individual organisms 'to maintain [themselves], to grow, to develop and to produce offspring'—what Wouters ([2005], p. 42) calls maintaining the life-state. This is because this reflects a central explanatory project of organismal biology, namely, to explain how individual organisms are able to maintain the life-state (compare Wouters [2005], p. 42).

I say that maintaining the life-state is at least one overarching capacity since, as Boorse ([1977], p. 556) points out:

> Most behavior of organisms contributes simultaneously to individual survival, individual reproductive competence, survival of the species, survival of the genes, ecological equilibrium *et cetera*. As a result, it appears that different subfields of biology (e.g. genetics and ecology) may use different [overarching capacities] as the focus of their function statements.

In other words, different capacities may be relevant for different explanatory projects and, because of this, determine various overarching mechanistic explanations that may feature diverging function ascriptions in the various *explanantia* generated.

Boorse takes this to show that the capacity at which a mechanistic explanation begins has to be, at least to some extent, determined by the interest of the theorist. It may be correct to say that the choice of *explanandum* is interest-relative—no doubt, different branches of the life sciences will be interested in explaining different phenomena.[6] But this does not imply that function ascriptions used in the explanation of such phenomena are interest-relative, or are somehow correct only relative to the perspective of the theorist providing such explanations.

To see this, note two things. First, it is possible that items can be multi-functional, and the notion of function defined below will allow for this. Second, the kinds of overarching capacities Boorse is concerned with in the above quote can all be considered biological goals. For our purposes, we can think of biological goals as capacities exercised at a particular level of organization (compare Weber [2017], p. 4747). For example, maintaining the life-state constitutes a biological goal because it is a capacity carried out by the whole organism, and so is a capacity that is exercised at the organismal level of organization. Likewise, the circulatory system's ability to transport oxygen, carbon dioxide, and nutrients to other parts of the body also counts as a biological goal since it is a capacity carried out by an organ system, and so is a capacity that is exercised at the organ system level.

There are a number of naturally occurring capacities that are goals with respect to the organismal level of organization that life scientists observe and any of those capacities may be

---

[6] As Reiss and Sprenger ([2020], sec. 3.1) write: 'It is almost universally accepted that the choice of a research problem is often influenced by interests of individual scientists, funding parties, and society as a whole'.

constitutive of some overarching mechanistic explanation. Given this, it is possible that an item may receive multiple function ascriptions given different overarching mechanistic explanations. However, as long as it is plausible (on empirical grounds) that the capacity that is featured in an overarching mechanistic explanation is a genuine capacity of the relevant system, and that we allow for the possibility of multi-functional items, there is no need to think that function ascriptions are made relative to the interests of the theorist just because there may be multiple places to start an overarching mechanistic explanation that may appear to generate conflicting function ascriptions. The point, in effect, is that although function ascriptions depend on which overarching capacity is used to start a mechanistic explanation (which must be chosen by a theorist), what counts as an overarching capacity is not relative to the interest of the theorist, nor are function ascriptions once that overarching capacity is chosen. So, according to the present account, function ascriptions are not interest-relative, unlike the views of Cummins ([1975]), Hardcastle ([1999]), and Craver ([2001], [2013]).[7]

One may want to opt in for a different, and perhaps more constrained, way of fixing the goals at the organismal level. For example, Piccinini ([2020], p. 68) argues that the goals of an organism should be those capacities that 'individuate living organisms as a kind'. That is, the exercise of such capacities is what differentiates living systems from non-living systems. These capacities, according to Piccinini ([2020], pp. 68–69), are the ability to survive, develop, reproduce, and help others. Goals at the organismal level can be individuated in such a way without disturbing the main tenets of the present account.

3. A Proposal

My proposed notion of function is defined as follows, where X ranges over item types and S over system types:

T: A function of X in S is to ϕ relative to a goal G and circumstance C if and only if:

(i) G is a goal of S.

(ii) A typical way that items of type X make a contribution to G in C is by ϕ-ing.

To clarify, T is stated as defining a notion of function for overarching mechanistic explanations,

---

[7] For a helpful discussion of why function ascriptions are interest-relative on such views, see Garson ([2016], chap. 5).

Weber ([2017]) makes a similar argument. According to him, statements of the form 'A function of X is to ϕ' only have truth values relative to a system that X stands in a suitable relation to and a goal of that system. And he takes goals to be capacities exercised at a particular level of organization, and takes it as uncontroversial that such goals exist. Weber then argues that once such parameters have been filled in, the truth or falsity of a function statement is an objective matter. He says 'the only thing that might depend on the investigator's interest may be the extent in which the truth of a function statement is worth knowing or *scientifically significant*' ([2017], p. 4747).

but naturally extends to sub-mechanistic explanations once they have been generated from an overarching mechanistic explanation.

Now, three central notions have been introduced in T: that of a circumstance, that of typicality, and that of a contribution. I address each in turn. (A note of caution: Since cashing out what a contribution amounts to will take the most space, I put it last in the order of presentation. However, some readers may find it helpful to jump to section 3.3 to familiarize themselves with that notion first.)

3.1. Circumstances

The present account relativizes function ascriptions to a circumstance in order to acknowledge that an item may have a typical way of making a contribution only in a restricted class of circumstances. For example, while it is true that a typical way that polar bear fur makes a contribution to maintaining the life-state of polar bears is by reducing heat loss, it is also true that it only contributes in this way in circumstances suitably like those found in an Arctic environment. The polar bear's fur would not contribute in such a way, for example, in a tropical environment (compare Garson [2016], p. 73). Further, it is possible that there will be a set of circumstances that are relevant when ascribing a particular function. For example, given the vital nature of hearts (for say vertebrates), the relevant circumstances for hearts will be any circumstance that a vertebrate could find themselves in. In such a case we can substitute C in T with a set of circumstances and adjust accordingly.

For our purposes, circumstances can be thought of as classes of spatio-temporal regions. To take our above example, regions that resemble those found in the Arctic would constitute one class of spatio-temporal regions while regions like those in the tropics would constitute another. Furthermore, classes of spatio-temporal regions may be individuated by a single feature rather than by a complex of features, as is likely the case in our previous examples. For example, when discussing the function of a gag reflex we will want to say that the spatio-temporal regions that contain asphyxiation constitute a class of their own (compare Garson [2016], p. 74). Furthermore, as the previous case suggests, a circumstance does not need to be restricted to spatio-temporal regions that are external to (that is, outside of), say, a whole organism. Nevertheless, a circumstance should be distinct from, that is, not individuated with respect to, the item to which the relevant function ascription is being made. That is, we should think of a circumstance as something the relevant item can occupy, and hence, as something distinct from the item itself. Sometimes I will describe circumstances as being events (for example, attending a doctors' appointment). I take this to be harmless since, for our purposes, we can think of events as classes of spatio-temporal regions extended over an interval of time.[8]

---

[8] Some authors (for example, Lewis [1986]) take events to be identical to classes of spatio-temporal regions, and thus take events to be identical to circumstances as the term is being used here. That is fine as long as we can recognize a distinction between events that take place over an interval of time and events that take place at a single point in time. Some circumstances will be more akin to the former while others more akin to the later.

Note that because the view relativizes function ascriptions to a circumstance, we cannot determine whether an item has a particular function until we have specified what that circumstance is. This raises the question of where our knowledge of these circumstances comes from. To be clear, what we need to know is whether a particular item has a function in some specified circumstance. Take the polar bear's fur. If we want to know whether polar bear fur has the function to reduce heat loss relative to a specific circumstance C, we can do so by observing whether reducing heat loss is a typical way that polar bear fur contributes to the relevant goal in C. More generally, if we want to know whether an item X has the function to ϕ in a circumstance C, then we can do so by observing whether ϕ-ing is a typical way that items of type X contribute to the relevant goal in C. Thus, in order to verify a function ascription, we need to observe the relevant item in the relevant circumstance. Nevertheless, according to the present account, an item may have a function with respect to a circumstance it has never encountered. That is, if an item, X, finds itself in a novel circumstance, C^n, it is perfectly compatible with the present account that X could have a function with respect to C^n at a time t despite not having been in C^n prior to t. Thus, the relevant circumstance used in evaluating function ascriptions is not necessarily a circumstance in which the relevant item has been in historically (compare Walsh [1996], p. 564; Garson [2016], p. 74). And so function ascriptions on the current view are not necessarily tied to any historical property of the relevant item.

3.2. Typicality

We can define typicality, following Wilhelm ([2022], p. 563), as follows: A property, P, is typical among members of a reference class, Γ, if and only if nearly all members of Γ exemplify P.[9] For our purposes, the relevant reference classes are items of type X that are making a contribution to some goal in some circumstance, and the relevant properties are 'ways that items could make a contribution to some goal in some circumstance'.[10]

According to T, for any item of type X and system of type S, a typical way that items of type X make a contribution (relative to some goal of S and circumstance) determines the function of items of type X. And, if X is constituted by items $x_1, ..., x_n$, then the function of each x_i is determined by the typical way that items of type x_i make a contribution to X's ability to carry out its function, which would be the relevant goal.

For example, consider the circulatory system in humans and one of its sub-items, namely, the heart—given the vital nature of the heart and circulatory system, presumably the relevant circumstances are any in which a human may find themselves. A typical way that the circulatory system makes a contribution to the maintenance of the human life-state is by transporting

---

Furthermore, an adequate metaphysics of events may want to complicate this picture, but such complications can be put aside for our purposes. For helpful discussion, see (Parsons [1990], chap. 8).

[9] For a detailed defence of this notion of typicality and its role in scientific explanation, see (Wilhelm [2022]) .

[10] Note that according to T, functions are not being directly attributed to the reference classes by which the relevant typicality statements are evaluated. Rather, functions are being attributed to item types.

oxygen, carbon dioxide, and nutrients. And, a typical way that the heart makes a contribution to the circulatory system's ability to carry out this task is by pumping blood. So, a function of the circulatory system is to transport oxygen, carbon dioxide, and nutrients while a function of the heart is to pump blood. The heart also constitutes a system that can be broken up into component items, namely, its four chambers and valves, and the functions of those items are determined in the same way that the heart's and the circulatory system's functions are determined, and so on until no further decomposition is possible.

It is important to note that the view does not claim that the function of an item is what it typically does full stop. If that were the view then it would clearly be inadequate. For example, sperm typically do not fertilize ova, even though the function of sperm is to fertilize ova. However, given that the function of sperm is to fertilize ova, the view here predicts that a typical way that sperm make a contribution to a goal of the system of which they are a part, relative to some circumstance, is by fertilizing the ova.

3.3. Contribution

The main goal of this subsection is to provide an account of what it means for an item to contribute to a goal of a containing system. Before doing so however, it is important to note that an item X with a function to ϕ may ϕ without contributing to the relevant goal in the relevant circumstance. In such cases, the item would be ϕ-ing below or above an appropriate rate so to speak (Schwartz [2007]). In many cases there will be many rates at which an item would be contributing to a central activity of the system in which it is contained. For example, all else equal, if a human heart at rest has a heart rate that falls outside the 65 to 74 beats per minute range then it would not be contributing to the circulatory system's ability to transmit nutrients and so on to other parts of the body by pumping blood (Cook et al. [2006]; compare Garson and Piccinini [2014]). What this means is that falling outside of this range would (among other things) increase the 'likelihood of heart failure or sudden cardiac arrest'(Garson and Piccinini [2014], p. 11). Furthermore, I say 'all else equal' to allow for certain anomalies. For example, the relevant interval of rates for adult human athletes (or active human adults more generally) may have a lower bound than for non-active adult humans (McKinley et al. [2016], p. 766). Moreover, it may be catastrophic for an athlete to have a resting heart rate at 90 beats per minute whereas for a non-athlete it would potentially be fine. More generally, the set of rates relevant for determining whether a token heart at a time t is contributing to its containing system will be relative to the properties of the bearer of that heart. For human hearts, for example, the interval of rates that is appropriate for a specific heart h will depend (among other things) on the age, sex and fitness level of the bearer of h (compare Garson and Piccinini [2014], p. 11).

We should also be careful in choosing which interval(s) of rates are relevant for determining whether an item is contributing to the relevant containing system. There is a sense in which the fact that my heart is not making a thump-thump noise at 250 decibels is contributing to my circulatory system's ability to transmit nutrients and so on to other parts of the body, in virtue of the fact that if it were doing so, something catastrophic would instantly occur (for example, my heart may explode). But, we do not want to say that the heart's making a thump-thump noise contributes to the circulatory system's ability to transmit nutrients and so on to other parts of the body.

The problem can be solved by noting that relativizing a contribution to a single interval of rates is insufficient. Rather, we must relativize to two intervals of rates. One interval is an appropriate set of rates R and the other is an inappropriate set of rates R'.

Consider a concrete case: say that we have a heart that contributes to the circulatory system's ability to fulfil its function when it pumps blood within the 65-74 beats per minute range. Thus, the values within this range constitute the appropriate interval R_{beat}. Now, in addition to R_{beat}, there must be a contrast interval of rates R'_{beat}—the inappropriate set of rates—such that if the heart were ϕ-ing at a rate r in R_{beat}, then the likelihood that the circulatory system would fulfil its function would be greater than if it were ϕ-ing at a rate r in R'_{beat}. For example, R'_{beat} may contain rates greater than 74 beats per minute and those less than 65 beats per minute. However, for any inappropriate set of rates R', R' must have a lower bound l such that the likelihood that the relevant item will ϕ at l is non-negligible. That is, when determining whether an item is contributing to some goal of a containing system, we are not just concerned with any arbitrary set, or sets, of rates, we are only concerned with those at which the relevant item has a non-negligible chance of performing the relevant activity.

Now, let us return to the thump-thump noise that the heart makes. Presumably the appropriate set of rates would contain those roughly at or between 45 and 75 decibels. However, the probability that the heart makes a thump-thump noise loud enough such that the likelihood of the circulatory system not fulfiling its function decreases in virtue of how loud that thump-thump noise is, is negligible (perhaps even nomologically impossible). Thus, there is no set of inappropriate rates that satisfies our criteria in this case.

Let us move forward to specifying the conditions under which an item X contributes to a goal G of a system S in a circumstance C by ϕ-ing. In light of our above discussion, a first condition can be formulated as follows: an item of type X contributes to a goal G of a system S in a circumstance C by ϕ-ing (relative to rate-intervals R and R') only if, the probability that S fulfils G in C would be greater if X were to ϕ at an appropriate rate, that is, at some r in R, than at an inappropriate rate, that is, at some r in R'.[11] For example, if a human heart were pumping blood at an inappropriate rate, then the likelihood of heart failure, or perhaps a less life-threatening condition, would be higher than if that heart were pumping blood at an appropriate rate. And because of this, the likelihood of the circulatory system fulfiling its function (the relevant goal) would be higher in the second case.

---

[11] One may want to relativize contributions to a time as well. This detail is left out for readability; likewise for C, FC, D, and M, below.

To state our first condition more rigorously and compactly, we can regiment it as follows: Let $\phi^{X,C}_{r\in R}$ mean 'X ϕs at an appropriate rate in C', $\phi^{X,C}_{r\in R'}$ mean 'X ϕs at an inappropriate rate in C', $G^{S,C}$ mean 'S fulfils G in C', '$p \mathrel{\Box\!\!\rightarrow} q$' means 'if it were the case that p, then it would be the case that q', and $p(\cdot)$ be a probability function that is defined in the normal way, then:

Pr: An item, X, contributes to a goal, G, of a system, S, in a circumstance, C, by ϕ-ing (relative to rate-intervals R and R') only if $(\phi^{X,C}_{r\in R} \mathrel{\Box\!\!\rightarrow} p(G^{S,C}) = v) \wedge (\phi^{X,C}_{r\in R'} \mathrel{\Box\!\!\rightarrow} p(G^{S,C}) = v') \wedge (v > v')$.[12]

In English, the right hand side of the conditional in Pr says that 'if X were to ϕ at an appropriate rate, then the probability that S fulfils G in C would be v, and if X were to ϕ at an inappropriate rate, then the probability that S fulfils G in C would be v', and v is greater than v''.

Careful readers will want to subscript $G^{S,C}$ with an appropriate rate, but this detail is left out for readability. Importantly, $p(\cdot)$ should receive an objectivist interpretation, rather than a subjectivist or evidential interpretation. So, to say that $(\phi^{X,C}_{r\in R} \mathrel{\Box\!\!\rightarrow} p(G^{S,C}) = v) \wedge (\phi^{X,C}_{r\in R'} \mathrel{\Box\!\!\rightarrow} p(G^{S,C}) = v') \wedge (v > v')$ for arbitrary, X, ϕ, S, G, C, R and R' is to say that the objective chance of, or frequency at which, S fulfils G in C would be greater if X were ϕ-ing at an appropriate rate than if it were ϕ-ing at an inappropriate rate.

Pr is too weak on its own to capture the relevant notion of contribution we are after. To see this, note that the right hand side of Pr is satisfied by both the heart's ability to pump-blood and its ability to make a thump-thump noise. This is because whenever the heart pumps-blood, given the morphological properties of the heart, it makes a thump-thump noise and vice versa. And, whenever it is not pumping it is not making a thump-thump noise and vice versa. So, in addition to Pr we need another necessary condition that can rule out the thump-thump noise of the heart as being a legitimate contributor to the circulatory system's ability to fulfil its function.

We can begin to do so by noting that there is no causal connection between the heart's ability to make a thump-thump noise and the circulatory system's ability to transmit nutrients, and so on, to other parts of the body. To see this, first imagine three worlds w_1, w_2, and w_3 that specify the behaviour of the heart with respect to making a thump-thump noise and pumping blood.

(w_1) The heart makes a thump-thump noise and pumps blood (the actual world is a world of this type).

(w_2) The heart makes a thump-thump noise and does not pump-blood.

(w_3) The heart pumps blood and does not make a thump-thump noise.

Both w_2 and w_3 represent a manipulation on the set of activities that the heart is able to perform in our world. Because of this, we can infer what would be the case if w_2 or w_3 were actual. Now, if w_2 were to suddenly become actual, all else equal and given what we know about physiology,

---

[12] Here I am assuming that v and v' are real numbers between zero and one.

the likelihood that the circulatory system would fulfil its function would decrease for those who possess hearts. However, the same would not be true if w_3 were to suddenly become actual. That is, all else equal, given what we know about physiology, and the fact that silent pumps are possible, the likelihood that the circulatory system would fulfil its function would remain the same. To put it differently, in w_2 we may say that there is an intervention on the heart's ability to pump-blood and, because of this, the likelihood that the circulatory system would fulfil its function would decrease. And, that in w_3 there is an intervention on the heart's ability to make a thump-thump noise and, in spite of this, the likelihood that the circulatory system would fulfil its function would remain the same.

Given an interventionist account of causation along the lines of (Meek and Glymour [1994]; Pearl [2000]; Sprites et al. [2000]; Woodward [2003], [2012]), among others, we can define a causal connection between two token events $\mathcal{C}$ and $\mathcal{E}$ as follows:

#: There is a causal connection between $\mathcal{C}$ and $\mathcal{E}$ if and only if an intervention on $\mathcal{C}$ would change the probability of $\mathcal{E}$.[13]

We may say that for any heart h contained in a circulatory system c, there is a causal connection between h's pumping blood (at an appropriate rate) and the fulfilment of c's function (in the relevant circumstances), but not between h's making a thump-thump noise and the fulfilment of c's function. Of course, the intervention that gives us reason for accepting this claim is not nomologically possible, given what we know about physiology and anatomy. However, many supporters of # (see, for example, Woodward [2003]) find nothing in principle wrong with citing merely logically or conceptually possible interventions when determining whether there is a causal connection between a relevant $\mathcal{C}$ and $\mathcal{E}$. If one finds an intervention objectionable it cannot solely be because it is not nomologically possible—some alternative reason must be given.

Putting this altogether we may provide the relevant analysis of contribution as follows:

C: An item X contributes to a goal G of a system S in a circumstance C by ϕ-ing (relative to rate-intervals R and R') if and only if

(i) $(\phi^{X,C}_{r\in R} \,\Box\!\!\rightarrow p(G^{S,C}) = v) \wedge (\phi^{X,C}_{r\in R'} \,\Box\!\!\rightarrow p(G^{S,C}) = v') \wedge (v > v')$, and

(ii) there is a causal connection between $\phi^{X,C}_{r\in R}$ and $G^{S,C}$.

So, our notion of contribution is a causal one. Because of this, one may wonder whether there is any work left for our typicality clause to do in our official definition of function once we have this causal notion of contribution in the mix. The answer is 'yes'. The typicality clause in our definition will be put to work in making good on the function–accident distinction.

---

[13] Compare (Sprenger [2018], p. 373).

More specifically, given our notion of contribution, as will be shown in more detail below, the following are both true:

(1) It is possible that some heart will contribute to a circulatory system's ability to transmit nutrients to other parts of the body by pumping blood.

(2) It is possible that some heart will contribute to a human's ability to maintain itself by aiding in the diagnosis of various diseases.

In spite of this, however, what we want our view to predict is that a function of the heart is to pump blood and not to aid in diagnosing various diseases while maintaining that both 1 and 2 are true. I take it that the view will only be able to handle the function–accident distinction if it can do so. The typicality clause in our official definition of function allows the view to make these predictions, and hence, allows the view to handle the function–accident distinction.

Before seeing how the notion of function characterized by T satisfies our criteria for being a teleological function, observe that, according to it, an item may be multi-functional in (at least) two ways. Namely, given two goals G and G', an item X may typically make a contribution to G by ϕ-ing, but to G' by ϕ'-ing, and hence, in such a case, X would have a function to ϕ and to ϕ'. For example, the penile urethra gland has the function to conduct sperm relative to the reproductive system and has the function to evacuate urine relative to the urinary system.

Or, an item X may typically contribute to a goal by ϕ-ing in circumstance C, but when in circumstance C', it might additionally be typical of Xs to contribute by ϕ'-ing. Again, in this case, X would have a function to ϕ and to ϕ'. For example, consider the medulla oblongata in humans.[14] In circumstances containing no obtrusive objects in a subjects mouth, a typical way that the medulla oblongata makes a contribution to the respiratory system's ability to enable breathing is by regulating respiratory rates. However, in circumstances in which there is an obtrusive object in the mouth of a subject, another typical way that the medulla oblongata makes a contribution to that goal is by initiating a gag reflex. So, the view predicts that the medulla oblongata has both a function to initiate a gag reflex and to regulate respiratory rates, which is a desired result (see, for example, Garson and Piccinini [2014], p. 7).

3.4. Evaluating the proposal

As stated at the outset, a notion of function is fit to be identified with a notion of teleological function if and only if it allows for dysfunctional items and it abides by a function–accident distinction. I now show that the notion of function characterized by T can satisfy both conditions. Let us start with the former.

To a first approximation a token item X in a token system S is dysfunctional relative to some ϕ, a goal G and circumstance C if and only if G is a goal of systems of type S, and X is unable to

---

[14] This example is due to Garson and Piccinini ([2014]). They handle the case in a substantially different way, but I think their way of handling it generates issues elsewhere, which are discussed below.

contribute by ϕ-ing in C, where ϕ-ing is a typical way that items of type X make a contribution to G in C, where we define what it is for an item to be unable to contribute by performing some activity as follows:

M: A token item X is unable to contribute to a goal G by ϕ-ing in a circumstance C if and only if X would not contribute to G by ϕ-ing if it were in C.

Our definition of dysfunction will need to be modified, but note that, as stated, it is capable of distinguishing between cases in which an item is not performing its function because the item is not in a situation in which it would be appropriate to do so, and cases in which the item is not performing its function because it is dysfunctional (Kingma [2010]). For example, if the digestive system of some human subject were paralysed, then if that subject were to ingest food (that is, attempt to break a fast), the subject's digestive system would not break down food (one of its functions) in virtue of being paralysed. In this case such a digestive system ought to be considered dysfunctional. Contrast this with the case in which the digestive system is not paralysed but is not breaking down food in virtue of the fact that the possessor of the digestive system is fasting. In that case such a digestive system ought not to be considered dysfunctional.

To see how it draws this distinction, first note that, according to the current account, the digestive system's function to break down food is relativized to a certain goal—in this case, maintaining the life-state of humans—and some circumstance. The relevant circumstance is that in which the bearer of the digestive system is not fasting (a spatio-temporal region in which the bearer of the digestive system has ingested sustenance). Let us call this circumstance 'C^d'. Now, according to the definition we have on the table, it is not enough for a token digestive system to not be contributing by breaking down food in order to be considered dysfunctional. Rather, it would have to not contribute by breaking down food if it were in C^d. That is, only if a token digestive system were unable to contribute by breaking down food in C^d would it be considered dysfunctional. Given that paralysing the digestive tract of a digestive system would rule out the possibility of that digestive system breaking down food in C^d, a digestive system with a paralysed digestive tract would be considered dysfunctional by the present account.

However, the account as stated is inadequate. This is because it will count too many items as dysfunctional. To see this, note that if a subject S were to break fast after a vigorous exercise their digestive system may digest food at a rate that is similar to that of a subject S' who has paralysed their digestive system from taking too much paracetamol. Thus, whenever S's digestive system does digest food at a rate similar to that of someone who has taken too much paracetamol after exercising vigorously in C^d, the view here will predict that S's digestive system is dysfunctional. This result is counterintuitive (Piccinini [2020], p. 79).[15] Before moving forward, it is important to note that there may be a way that the digestive system is

---

[15] It is worth noting that one may find this consequence more acceptable than not. As Hausman ([2011], p. 666) notes 'exercise can cause a slew of pathologies from blisters to dehydration, and there is no reason why it cannot regularly cause temporary pathologies within the digestive system'; similarly, then, we may say that vigorous

functioning in such a circumstance, since it has a function, in addition to its function to digest food, to slow itself down after exercise in order to permit blood flow to the muscles (Hausman [2011], p. 666). The important thing to note is that the digestive system, in such circumstances, would only be dysfunctional relative to its function to digest food.

A principled way of resolving this issue is by adopting Piccinini's notion of a functional trade-off (Piccinini [2020], pp. 81–86). According to Piccinini, in a case where an item, such as the digestive system, has limited resources in order to fulfil all of its functions, some functions will take priority over others. As Piccinini ([2020], p. 82) puts it:

> Organisms perform many functions at once. Performing them at adequate rates requires expending energy and other resources. Since organisms possess bounded resources, performing multiple functions above a certain threshold may interfere with performing other functions above certain thresholds.

For example, in the case where a subject S breaks fast after vigorous exercise, their digestive system will prioritize its function to slow itself down after exercise in order to permit blood flow to the muscles over its function to digest food. In such a case, the digestive system will be digesting food (that is, ϕ_d-ing) at a rate similar to the rate it would be digesting food if S had taken too much paracetamol. However, this low rate of ϕ_d-ing—which leads to the judgement that the digestive system is not contributing to the relevant goal in this case—is caused by the fact that the digestive system is prioritizing its function to slow itself down after exercise in order to permit blood flow to the muscles. So, in such a case, the fact that S's stomach cannot digest food after vigorous exercise is due to a trade-off, which is not true in the case where S's stomach cannot digest food when S has taken too much paracetamol. So, only in the latter case should we consider S's stomach dysfunctional.

We are now in a position to provide our official notion of dysfunction:

D: A token item, X, in a token system, S, is dysfunctional relative to some ϕ, goal G, and circumstance C if and only if:

(i) G is a goal of systems of type S;

(ii) X is unable to contribute by ϕ-ing in C, where ϕ-ing is a typical way that items of type X make a contribution to G in C; and

(iii) X's inability to ϕ in C is not due to a trade-off.

Note that conditions i and ii jointly imply, by T, that a function of X in S is to ϕ relative to G and C.

exercise causes a temporary dysfunction, one that occurs as an adverse side-effect of vigorous exercise. However, as Kingma ([2016], p. 397) argues, going this route may open the door up to even more counterintuitive dysfunction ascriptions.

One may point out that the present account of dysfunction ultimately falls back on there being some circumstance(s) to which we relativize dysfunction ascriptions, and go on to say that once these circumstances have been determined, it is unclear what real work the account is doing in capturing a function–dysfunction distinction. Admittedly, our account of dysfunction does make crucial reference to the circumstances to which function ascriptions are relativized. This is because whether an item X is dysfunctional relative to its function to ϕ at some time t should be a matter of whether X, in its condition at t, would ϕ in a particular circumstance, and not a matter of whether X ϕ's at t *simplicter*, nor a matter of whether X, in its condition at t, would ϕ in an arbitrarily chosen circumstance (Kingma [2010]; Hausman [2011]; Garson and Piccinini [2014]). Moreover, the circumstances that determine whether an item is dysfunctional with respect to its function to ϕ on some occasion are the same as those that determine whether it has the function to ϕ. So, we are not appealing to an *ad hoc* set of circumstances when judging whether an item is dysfunctional, but to those that the relevant function ascription has already been relativized. And, as discussed earlier, these circumstances are not tied to any historical property of the relevant item. Importantly, however, the appeal to circumstances is not doing all the work in capturing a function–dysfunction distinction. This is for at least two reasons. First, recall that an account of function can only handle the function–dysfunction distinction if it can maintain that an item can have the function to ϕ without being capable of ϕ-ing. The account makes good on this desideratum by requiring it to be the case that if an item X is dysfunctional with respect to its function to ϕ, then X must be a token of an item type with the function to ϕ and also be incapable of ϕ-ing (in the relevant circumstance). The present account is able to do so because it ascribes functions to item types. So, the account's ability to ascribe functions to item types is doing a good amount of work in the account's ability to capture a function–dysfunction distinction. Second, the view only takes it as a necessary condition that the relevant item, which has the function to ϕ, cannot ϕ in the relevant circumstances. An additional necessary condition in D specifies that this inability to ϕ cannot be due to a trade-off.

Let us move forward to the function–accident distinction. An account of function needs to distinguish an item's functions from beneficial contributions an item may make (seemingly) by accident. For example, it is true to say that

(3) A function of the heart is to pump blood.

While it is false to say that

(4) A function of the heart is to aid in diagnosing various diseases.

Even though the heart can act as a useful diagnostic aid, since abnormal heart sounds may be taken as a symptom of some disease (compare Wright [1973], p. 148).

Plausibly, a notion of function abides by the function–accident distinction if and only if it can predict 4 false and 3 true and other cases like it. For the sake of space, I will only consider 3 and 4, but I believe the view extends naturally to other cases.

Before obtaining a clear prediction from the current analysis, a few more details need to be filled out. First, let us restrict ourselves to humans and their organs. Further, we need to specify the goals that are being contributed to. With respect to 3, the relevant goal is the circulatory system's ability to fulfil its function, which in turn, by performing that function, contributes to maintaining the life-state of humans. And, most plausibly, the relevant goal with respect to 4 just is maintaining the life-state of humans. And last, the circumstance(s) to which the function ascriptions are relativized to need to be decided. With regards to 3, as was done earlier, I will assume that any circumstance in which a human may find themselves is a circumstance in which the function ascription stated in 3 is relativized. With regards to 4, the relevant circumstances are physical examinations or doctors' appointments more generally.

We have already seen that the account predicts 3 is true. The current account does not predict, however, that 4 is true. This is because it is not true that nearly all hearts make a contribution to maintaining the life-state of a human by aiding in diagnosing various diseases, even in the given circumstances. To see this, first note that we want to maintain that there will be some hearts that do make a contribution to a human's ability to maintain itself in the relevant circumstances. In such cases, however, it has to be the case that the heart is actually diseased. Only in that case will it be true that

(5) If the heart were not aiding in diagnosing a disease (at some appropriate rate in a doctor's appointment), then the likelihood of the human maintaining it's life state would be lower than if it were not diagnosing a disease.

This is because if the heart weren't actually ailed, the likelihood of the human maintaining itself would be the same whether or not the heart acted as a diagnostic aid. This point is relevant because, while there are many instances of hearts acting as diagnostic aids, only some of the time will they be making a contribution to a human's ability to maintain itself. Thus, it is not typical for hearts to make a contribution in such a way.

To put the point concretely, suppose that we have four hearts, h_1, h_2, h_3, and h_4, where h_1, h_2, and h_3 are not ailed, but h_4 is and produces an abnormal heart sound. Let us further suppose that each heart is used as a diagnostic aid while in a doctor's appointment. In such a case, it may be true that h_4 contributes by aiding in the diagnosis of some disease, but h_1, h_2, and h_3 will not contribute in such a way for the reasons stated above. Thus, it will not be the case that nearly all such hearts contribute to maintaining the life-state of a human by aiding in diagnosing various diseases even in doctors' appointments. And hence, by our definition of typicality given in section 3.2, aiding in the diagnosis of various diseases will not be a typical way that such hearts contribute to maintaining the life state of humans in doctors' appointments. The claim I am making here is that this will be the case even when a more realistic reference class is chosen.[16]

---

[16] Note that switching the activity from 'diagnosing various diseases' to 'making a thump-thump noise' will not make this case any more difficult to handle. This is because it is not the case that nearly all hearts will contribute

Furthermore, this case highlights the work that the typicality clause is doing in the account. Without the typicality clause we would not be able to capture why the heart should not be assigned the function to aid in diagnosing various diseases. This is because, according to the present account, even if some token hearts contribute by acting as diagnostic aids in the relevant circumstances—as the above discussion has been assuming—this is not sufficient for the heart to have the function to aid in diagnosing various diseases. According to the current account, what more is needed is that nearly all token hearts contribute to their corresponding organism's ability to maintain itself under those circumstances by acting as a diagnostic aid. This final condition is of course captured by the typicality clause in the view's official definition of function.

One may object that if we understand the above case as we have done, then by analogous reasoning the view would not predict that a function of sperm is to fertilize ova—contrary to what was claimed earlier. To see why this objection does not hold it will be informative to go over the sperm case in some detail, and then contrast it with the case concerning hearts and diagnoses.

What we first want to do is show that the view predicts that a function of sperm is to fertilize ova. To do so, we need to specify the relevant circumstance, system, and goal of that system. Let us say the relevant system is the human reproductive system and the relevant goal is producing offspring. Furthermore, it is plausible that the relevant circumstance is copulation while ovulation is occurring.[17] Let us call this circumstance 'C^c'. Importantly, I take this circumstance to be the broadest possible circumstance in which a sperm could possibly contribute to the human reproductive system's ability to produce offspring by fertilizing an ovum. Any choice of circumstance narrower than this would run the risk of being *ad hoc*, or stacking the deck in favour of the present view. In sum then, we want to verify the following:

(6) A function of sperm is to fertilize ova in the human reproductive system relative to the goal of producing offspring, and circumstance C^c.

If 6 is true according to the present account then the following two statements are true:

(7) A goal of the human reproductive system is to produce offspring.

(8) A typical way that sperm contribute to the human reproductive system's ability to produce offspring is by fertilizing ova in C^c.

I will take it for granted that 7 is true. To see why 8 is true, first note that the relevant reference class for verifying the typicality claim is sperm that are contributing to the human reproductive system's ability to produce offspring in C^c, as the relevant reference class for any typicality

---

to maintaining the life state by making a thump-thump noise in doctors' appointments. Only hearts ailed in the right way will do so.

[17] I take this to include cases of *in vitro* fertilization where copulation technically never occurs.

claim that appears on the right hand side of T is an item of type *X* that is making a contribution to some goal in some circumstance. Further, note that 8 implies the following:

(9) Nearly all sperm contribute to the human reproductive system's ability to produce offspring by fertilizing ova in C^c.

To see why 9 is true say that we have a reference class containing 100 sperm, $s_1, ..., s_{100}$, and that 90 of these sperm, $s_{11}, ..., s_{100}$, do not fertilize an ovum in $C^c_{11}, ..., C^c_{100}$, respectively, but that sperm $s_1, ..., s_8$ fertilize eight ova, $o_1, ..., o_8$, respectively, in $C^c_1, ..., C^c_8$, respectively, and that s_9 and s_{10} both fertilize the same ovum o_9 in C^c_9. Further note that if two sperm fertilize the same ovum, the chance of reproductive success is negligible (McKinley et al. [2016], p. 1141). Among this class of sperm, it will be true that sperm typically do not fertilize any ovum, since only 10 out of 100 have done so; that $s_1, ..., s_8$ do contribute to the relevant reproductive system's ability to generate offspring by fertilizing an ovum in their respective C^c_is; and that s_9 and s_{10} do not contribute by fertilizing an ovum in C^c_9 (since two sperm fertilizing the same ovum curbs reproductive success). Furthermore, since eight out of the ten sperm that do fertilize an ovum contribute to the relevant reproductive system's ability to generate offspring in the relevant circumstance, it will be true that nearly all such sperm contribute by fertilizing an ovum in that circumstance. And hence, by our definition of typicality, it will be true that a typical way that sperm contribute to the human reproductive system's ability to produce offspring in C^c is by fertilizing an ovum. Again, the claim I am making here is that this will be the case even when a more realistic reference class is chosen.

Finally, let us contrast this with the case concerning hearts and diagnoses. As stated earlier, the view does predict that 4 is false, and I take this to be because the following is false:

(10) A typical way that hearts contribute to maintaining the life state of humans is by aiding in the diagnosis of various diseases in a doctor's appointment.

Note that the relevant reference class for this claim is hearts that are contributing to the human ability to maintain itself in a doctor's appointment. However, even in the relevant circumstance, it will not be typical for hearts to contribute to the maintenance of the life state of humans by aiding in the diagnoses of some disease. More specifically, the only time a heart *h* will contribute to maintaining the life state of some human by aiding in the diagnosis of some disease in a doctor's appointment is when *h* is ailed in the right way. And, given that nearly all hearts are not ailed in the right way even in the relevant circumstance, it follows that aiding in the diagnosis of some disease is not a typical way that hearts contribute to the maintenance of the life state of humans in the relevant circumstance. In sum then, the view is able to deny that a function of the heart is to aid in diagnosing various diseases, while maintaining that a function of sperm is to fertilize the ova.

According to the present account, then, what explains why it is a function of the heart to pump blood but not to diagnose various diseases is that typically hearts contribute in the former way but not in the latter. And so what is accidental is being explained by what is not typical. However, one may object that in other cases where we want to give an account of what is accidental and what is not, it is common to appeal to causal mechanisms or laws of nature. For example, we may say that a generalization is accidental because it does not pick out a law of nature. Further, it may be granted that statistical patterns will emerge once these causal mechanisms and/or laws are in place, but what is typical will be explained by the fact that it is not accidental and not vice versa. In other words, the view counterintuitively makes the claim that something is accidental because it is not typical, whereas we ordinarily think that something is not typical because it is accidental.

In essence, the objector is suggesting that the typicality clause in the view can be dropped in favour of causal considerations or by appeal to laws of nature in order to get rid of the unwanted reversal of explanation. With respect to the former move, note that the view already makes use of a causal notion, namely, the notion of contribution, in its official definition. But, as we saw above, that notion on its own cannot draw a function–accident distinction. And so it is unclear what other causal mechanism can be brought into play here. The second suggestion is interesting but would need to be developed in more detail. Presumably, such a view would have to make good on the claim that it is not a law of nature that hearts contribute to the maintenance of the life state by diagnosing various diseases while maintaining that it is a law of nature that hearts contribute to the circulatory system's ability to transmit nutrients and so on to other parts of the body by pumping blood. But, it is unclear why we should deny the former claim. For example, in the right conditions, as we saw above, it will be true that: if the heart were aiding in the diagnosis of some disease, then it would be contributing to the maintenance of the life state. So the law-like statement concerning hearts and diagnoses is counterfactual supporting, which is often taken as evidence that a law-like statement picks out a genuine law. For similar reasons, it may even be used to support predictions, another indication that a law-like statement picks out a genuine law.[18] In sum, it is not clear how laws of nature will be able to replace the typicality clause in our definition of function such that the resulting view will be able to capture a function–accident distinction. Because of this, I am content with maintaining that what explains why it is accidental for the heart to contribute to the maintenance of the life state by diagnosing various diseases is that it is not typical for hearts to contribute in such a way.

To see the advantages of the current view, I want to contrast it with the recent proposal of Garson and Piccinini ([2014]) that purports to be in the same family of views as the one being developed here. They claim that their view is a 'a novel and improved version of Boorse's

---

[18] Another option is to appeal to a particular account of lawhood, such as the best systems account (see, for example, Lewis [1983]), and argue that the law-like statement concerning hearts and diagnostics should not count as a law according to that conception. This would of course restrict the generality of the view. Further, it is unclear that such views are in the business of deciding what the laws of nature are (a descriptive question) rather than just providing an answer to the question of what laws of nature are, given we already have a pre-theoretic grip on what law-like statements should count as laws of nature. These are of course questions that outstrip the scope of this article.

biostatistical theory of functions' (Garson and Piccinini [2014], p. 1). With this I agree, since their view, among other improvements, provides a precise analysis of what Boorse ([1977], p. 558) calls 'statistically typical functions'—although Garson and Piccinini cash things out in terms of 'probability' instead of 'typicality'. I do want to note further, however, that the present view makes an additional contribution to such views by providing an account of what it is for an item to make a contribution to a containing system. But more importantly, unlike the current view, I want to argue that Garson and Piccinini's view predicts 4 to be true.

To see this, take into account their preferred definition of function (where some details—variables, and so on—are changed for consistency):

> A function of an item, X, in a reference class, Γ, is to ϕ if and only if:
>
> (i) $p(X$ contributes to survival or inclusive fitness in $\Gamma)$ is non-negligible.
>
> (ii) $p(X$ contributes to survival and inclusive fitness by ϕ-ing $|X$ contributes to survival or inclusive fitness in $\Gamma)$ is non-negligible.[19] (Garson and Piccinini [2014], p. 6)

With respect to our problematic case, no doubt i is satisfied with respect to humans—the probability that the heart contributes to the survival or inclusive fitness in humans is non-negligible. But, is it true that the probability that the heart, in humans, contributes to survival or inclusive fitness by aiding in diagnosing various diseases given the heart contributes to survival or inclusive fitness is non-negligible?

Prima facie, the answer is 'yes'. If there was only a negligible chance that the heart would contribute in such a way, then it is most unlikely that doctors would check the heart beats of patients during routine physical examinations. Perhaps Garson and Piccinini ([2014]) could respond by modifying or clarifying further what is meant by non-negligible, but, as their view stands, it looks to predict 4 as true.

It is worth emphasizing why the present account is able to handle the function–accident distinction, but Garson and Piccinini's is not. The main reason I take it is that Garson and Piccinini eschew appeal to typicality talk in favour of talk of non-negligible probabilities, whereas the present view cashes out typicality, following Wilhelm ([2022]), in terms of how many members of a reference class exemplify a particular property. What I have shown is that eschewing typicality talk in favour of non-negligible probabilities in the way that Garson and Piccinini do is inadequate to generate a version of the biostatistical theory that can handle the function–accident distinction. By contrast, understanding the notion of typicality in the way done here allows for a version of the biostatistical theory that can handle the function–accident distinction.

---

[19] Garson and Piccinini ([2014], p. 6) amend this definition 'by identifying the function of [an item] with its non-negligible contribution to survival or inclusive fitness on those occasions on which it does so contribute'. This amendment, however, does not help them here.

4. The Ahistorical Nature of the Account

I take the present view to be ahistorical insofar as it capable of ascribing functions to items without appealing to the selection history of those items. Some views of function, such as traditional aetiological views, are famously not ahistorical in this sense (see, among others, Milikan [1989]; Neander [1991], [2017]). such views have it that no function may be attributed to any item without a selection history, or without making claims about that item's selection history. This is because, according to those views, the function of an item, X, is that activity ϕ that has been selected for by some selection process over some historical time period because it provided some advantage to the bearer of X. That is, there is a built in historical aspect to such views insofar as function ascriptions depend on an item's selection history. Many theorists have found this consequence unattractive (see, among others, Boorse ([1976]), Amundson and Lauder ([1994]), Schlosser ([1998]), and Weber ([2005])). In what follows, I will argue that the present view is ahistorical in the sense just described and that this provides it with an advantage over the aetiological theory of function—at least in the context of mechanistic explanations. It should be noted that this is a substantive claim since Neander ([2017]) has recently argued that the relevant notion of teleological function at play in mechanistic explanations is an aetiological one.

My view is ahistorical in the sense that it ascribes functions to items without appealing to the selection history of those items because T does not make any reference to an item's selection history. And, as was noted earlier, the circumstances that function ascriptions are relativized to are not tied to any historical property of the relevant item and so are not tied to any property concerning that item's selection history.

The main reason why being ahistorical in this sense provides the view with an advantage over its aetiological competitor is because there is no independent reason for thinking that the notion of function that occurs in the *explanans* of a mechanistic explanation requires the kind of historical aspect that aetiological views have built into them. Before looking at why this is so, note that in contrast to views that evoke selection history, the present view allows items to have functions even if the relevant activity evoked in the function ascription does not explain why the item was selected for, or if the relevant item has no selection history at all. Aetiological theories, of course, cannot do this. According to such theories, 'we are not allowed to ascribe a function to the first beating heart that contributed to blood circulation' (Schlosser [1998], p. 321). Nor are we allowed to say that the first feathers that were co-opted for flight in birds had the function to enable flying (Schlosser [1998], p. 321). Although I will not argue for this here, many authors have taken the failure of aetiological theories to allow for such function attributions to be a cost (Bigelow and Paragetter [1987], pp. 194–95; Walsh [1996], pp. 555–58; Walsh and Ariew [1996], p. 498; Schlosser [1998], p. 321; Weber [2005], p. 36; Piccinini

[2020], pp. 74–75).[20]

We can demonstrate that the view has this latter advantage by appealing to an example from (Weber [2005], p. 36), which is ‘fictional but not unrealistic’. I take it that the case is realistic insofar as it is supposed to resemble the two kinds of cases discussed in the previous paragraph. Weber asks us to ‘imagine that a mutation occurs in a bacterial gene that allows the bacterium to utilize a previously indigestible sugar. This newly created gene is advantageous and spreads in the bacterial population’ (Weber [2005], p. 36). Presumably this means that the gene's corresponding enzyme—call it ‘$\mathrm{e}_{\mathrm{token}}$’—contributes to the bacterium's ability to reproduce in some circumstance C. For convenience, call the relevant bacterium ‘B’. The present view attributes a function to $\mathrm{e}_{\mathrm{token}}$ relative to the contribution it makes to B's reproduction. That is, the present view supports the following function ascription, where ‘e_{type}’ is the type of enzyme of which e_{token} is a token:

(11) A function of $\mathrm{e}_{\mathrm{type}}$ in B is to metabolize sugar (relative to B's ability to reproduce in C).

According to the present account, even if there is only one token of $\mathrm{e}_{\mathrm{type}}$ that contributes to some B's ability to reproduce in C by metabolizing sugar, that is sufficient for it have the function to metabolize sugar (if $\mathrm{e}_{\mathrm{token}}$ is the only enzyme of its type among the population of Bs that contributes by metabolizing sugar, then trivially nearly all tokens of $\mathrm{e}_{\mathrm{type}}$ will contribute in such a way). In contrast to the present view, however, an aetiological view is not able to attribute a function to the enzyme until it has been under some selection process.

One may object that the fact that the present view is able to attribute a function to $\mathrm{e}_{\mathrm{token}}$ (prior to that type of enzyme spreading in the population) is in tension with the idea that the account captures the function–accident distinction (compare Garson [2019], pp. 29–30). That is, why should we not say that the original $\mathrm{e}_{\mathrm{token}}$ contribution is accidental? I think that our intuitions underdetermine whether the contribution made by $\mathrm{e}_{\mathrm{token}}$ should be considered accidental when only $\mathrm{e}_{\mathrm{token}}$ is present. In light of this, the view is designed to attribute the function to $\mathrm{e}_{\mathrm{token}}$, but is also designed such that the function ascription in 11 can be retracted if it turns out that $\mathrm{e}_{\mathrm{token}}$'s contribution in the single case, in retrospect, was better described as an accident. That is, the view is designed to err on the side of permissibility with respect to function ascriptions.

Recall that what allows the present view to claim that $\mathrm{e}_{\mathrm{token}}$'s contribution licenses the function ascription given in 11 is that the gene that gives rise to $\mathrm{e}_{\mathrm{token}}$ mutated in only a single member of B—and so the relevant typicality claim is trivially satisfied. Now, if the gene spreads throughout the population, there are two ways things could go with respect to that gene in the population of Bs. First, if the relevant genes $g_1, ..., g_n$ give rise to enzymes $e_1, ..., e_n$ that are of the same type as $\mathrm{e}_{\mathrm{token}}$, and nearly all e_is contribute to the reproduction of their corresponding bacterium by metabolizing sugar, then the function ascription in 11 will remain true. However,

---

[20] There are, of course, well-known attempts to amend this problem for aetiological theories; see, for example, (Griffiths [1993]; Godfrey-Smith [1994]). For some shortcomings of these proposals, see (Walsh [1996], pp. 556–88; Schlosser [1998], p. 321).

if it turns out that each (or most, and so on) e_i fails to contribute by metabolizing sugar, then 11 will turn out to be false and thus the function previously ascribed to $\mathrm{e}_{\mathrm{type}}$ will need to be retracted. All I take this to show is that a function ascription can be retracted as time moves forward. If one takes this to be costly, then it is counterbalanced by the fact that it is able to attribute a function to the first instance of the useful enzyme, a result that many have found desirable (see, again, Bigelow and Paragetter [1987], pp. 194–95; Walsh [1996], pp. 555–58; Walsh and Ariew [1996], p. 498; Schlosser [1998], p. 321; Weber [2005], p. 36; Piccinini [2020], pp. 74–75).[21]

Now, what I take to be the main advantage of the present view's ahistoricity is that it does not force life scientists to be making claims about an item's selection history that outstrips the evidence that a life scientist may have when attributing functions. To deepen this line of thought, take into account an instance of function discovery in the life sciences. In the early 1960s researchers were able to uncover a function of the thymus, an organ found in the immune system that was long thought to be non-functional. They did this by performing interventions on baby rats and rabbits—neonatal thymectomies—and observing the adverse effects this had on their growth and development. What researchers inferred from these studies was that the thymus was responsible for lymphocyte cell production, which is essential for the immune system to carry out its function—namely, to protect the body from illness and disease. And, because of this, researchers attributed the function of producing lymphocyte cells to the thymus in mammals (Burnet [1962]).[22] As Schaffner ([1993], p. 386) notes, 'no evolutionary argument is offered' by the researchers who make this function attribution. Moreover, although not stated explicitly, Schaffner suggests that no evolutionary evidence was available to such researchers. The evidence for the function ascription primarily concerns how the removal of the thymus affects the immune system's function. In other words, the function is inferred by seeing how the removal of thymus breaks its containing system. This signifies that we can explain how the immune system works by appeal to an activity of the thymus—and this activity, given our gloss of mechanistic explanations we began with, ought to count as a function of the thymus—even in the absence of any considerations concerning the selection history of the thymus.

Of course, the supporter of the aetiological view of function will say that even though no evidence concerning an item's selection history is required to make the function ascription, the

---

[21] Given that the view can attribute functions to mutations, it is plausible that it will also be able to attribute functions to a swamp system, for example, a spontaneously formed molecule for molecule duplicate of Donald Davidson. One may worry that this is also in tension with the view's ability to capture a function–accident distinction, since swamp systems, as they are usually described, arise as sheer accidents. I think a similar response to what we have given for the mutation case also suffices as a response to the objection when it is geared towards the swamp system case. Given that the swamp system is not typically conceived of as being something that spreads through a population, we can ask ourselves counterfactually: if a swamp system S were to reproduce, would its offspring have items that contribute in the same way that S's do? All though such counterfactuals are far-fetched, those pulled by the swamp system case in the first place will be inclined to say 'yes'.

[22] For philosophical discussion of this case, see (Canfield [1964]; Schaffner [1993], chaps 3, 4, and 8, especially pp. 385–88).

function ascription is implicitly historical in the way that aetiologists say it is. More specifically, the move for the aetiologist at this point is to say that we cannot make sense of the normativity of function ascriptions without bringing in the relevant item's selection history (Neander [1991], [2002]; Rosenberg and Neander [2009]). Importantly, this move requires the aetiologist to be committed to the claim that scientists making function ascriptions are committed to specific facts about the selection history of a particular item. Namely, that for any item X that is ascribed a function to ϕ, a selection process (for example, natural selection) acted on X because it conferred some advantage to the past bearers of X by ϕ-ing (compare Amundson and Lauder [1994], p. 459). For simplicity, we can say that making such a claim is tantamount to saying that all such attributed functions are historical. In sum, we can understand the supporter of the aetiological view as proposing a transcendental argument with roughly the following structure:

(P1) The only way that there could be teleological functions is if those functions were historical.

(P2) There are teleological functions.

(C) So, teleological functions are historical.

The argument is valid, but, in the absence of knock-down arguments to the view proposed here, the aetiologist cannot take the first premise for granted.

In sum, then, the view here is privileged over the aetiological view insofar as it does not require scientists to be implicitly making certain claims about the past when they are attributing functions in cases where such claims about the past either outstrip the evidence available to such scientists or are simply not relevant to their reasons for making such ascriptions.

An alternative motivation for thinking that life scientists are invoking history when they make function ascriptions is that they project their findings on to unexamined cases. That is, when ascribing a function to some item on the basis of experiments involving a particular sample of some specimen, the scientists conducting those experiments are only able to generalize their results to the whole species (or even to other species located on the same phylogenetic tree as the species under question) on the assumption that the specimen used in the experiment shares an evolutionary or developmental history with the other members of the relevant species. So, to make such generalizations, life scientists must be making implicit appeals to the evolutionary or developmental history of the relevant item.

My account can acknowledge that common descent and evolutionary or developmental history underwrite the relevant ampliative inferences. Indeed even those, such as Amundson and Lauder ([1994], pp. 454–55), who are most sceptical that life scientists are making claims about an item's selection history when ascribing functions in many contexts, are wholly committed to the idea that two organisms belong to the same species in virtue of descending from a common ancestor. What the view denies (and must deny) is that this amounts to making the claim that the relevant item X being ascribed the function to ϕ was selected for because it conferred some advantage to past bearers of X because it ϕ-ed. And thus, given that the account only claims to have an advantage with respect to other views insofar as it allows function ascriptions to be neutral with respect to such facts concerning X's selection history, the objector must provide further reason for why the fact that common descent and evolutionary or developmental history underwrite such ampliative inferences commits one to the claim that X did confer some advantage to its past bearers because it ϕ-ed.

At this point one may further argue that a statistical ahistorical view, such as the one put forward here, is untenable in any case because of the problem of pandemic disease. Neander ([1991], p. 182) states the objection as follows:

> A statistical definition of biological norms implies that when a [item] standardly fails to perform its function, its function ceases to be its function; so that if enough of us are stricken with disease (roughly, are dysfunctional) we cease to be diseased, which is nonsense.

Given the complexity of this issue, I will not be able to do full justice to it here. Rather, I just want to point out that Neander's original objection is made with respect to a very specific version of a statistical ahistorical view, namely, an early version of such a view provided by Boorse ([1977]).[23] Because of this, it is not obvious that the problem will plague the current view. That said, the objector would have to do more to show that the view faces such a problem. Although I don't have the space to show this here, I don't think it does.

5. Conclusion

This article has provided a novel statistical theory of teleological function according to which a function of an item X is to ϕ relative to a goal G and containing system S in a circumstance C just in case G is a goal of S, and a typical way that items of type X make a contribution to the realization of G is by ϕ-ing. Furthermore, the view provides an attractive alternative to aetiological views. I have also provided an account of what it is for an item to make a contribution to the realization of a goal of some system, and argued that the view has an advantage over one of its main competitors insofar as it can distinguish between those activities of an item, X, that are functions of X and those that are not functions of X but may be beneficial in a certain class of circumstances.

Acknowledgements

For extensive feedback on various versions of this article I would (especially) like to thank Seth Goldwasser, Alice van't Hoff, Paul Pietroski, Susanna Schellenberg, and two anonymous

---

[23] For a reply, see Boorse ([2002], pp. 95–103).

referees at this journal. I would also like to thank Diego Arana Segura, Sam Carter, Denise Dykstra, Justin Garson, Adam Gibbons, Brian McLaughlin, Lukas Myers, Ezra Rubenstein, Isabel Uriagereka Herburger, Bram Vaassen, Jonathan Schaffer, Daniel Stoljar, and Isaac Wilhelm.

Department of Philosophy
Rutgers University–New Brunswick
New Brunswick, NJ, USA
andrew.rubner@gmail.com

References

Amundson, R. and Lauder, G. V. [1994]: ‘Functions without Purpose: The Uses of Causal Role Functions in Evolutionary Biology’, *Biology and Philosophy*, **9**, pp. 443–69.

Bigelow, J. and Pargetter, R. [1987]: ‘Functions’, *Journal of Philosophy*, **84**, pp. 181–96.

Boorse, C. [1976]: ‘Wright on Functions’, *Philosophical Review*, **85**, pp. 70–86.

Boorse, C. [1977]: ‘Health as a Theoretical Concept’, *Philosophy of Science*, **44**, pp. 542–73.

Boorse, C. [1997]: ‘A Rebuttal on Health’, in J. M. Humber and R. F. Almeder (*eds*), *What Is Disease?* Totowa, NJ: Humana Press, pp. 3–134.

Boorse, C. [2002]: ‘A Rebuttal on Functions’, in A. Ariew, R. Cummins and M. Perlman (*eds*), *Functions: New Essays in the Philosophy of Psychology and Biology*, Oxford: Oxford University Press, pp. 63–112.

Burnet, M. [1962]: ‘The Thymus Gland’, *Scientific American*, **207**, pp. 50–57.

Canfield, J. [1964]: ‘Teleological Explanation in Biology’, *British Journal for the Philosophy of Science*, **14**, pp. 285–95.

Cook, S., Togni, M., Schaub, M. C., Wenaweser, P. and Hess, O. M. [2006]: ‘High Heart Rate: A Cardiovascular Risk Factor?’, *European Heart Journal*, **27**, pp. 2387–93.

Craver, C. [2001]: ‘Role Functions, Mechanisms, and Hierarchy’, *Philosophy of Science*, **68**, pp. 31–55.

Craver, C. [2013]: ‘Functions and Mechanisms: A Perspectivilist View’, in P. Huneman (*ed.*), *Function: Selection and Mechanisms*, Dordrecht: Springer, pp. 133–58.

Cummins, R. [1975]: ‘Functional Analysis’, *Journal of Philosophy*, **72**, pp. 741–65.

Cummins, R. [1983]: *The Nature of Psychological Explanation*, Cambridge, MA: MIT Press.

Garson, J. [2016]: *A Critical Overview of Biological Functions*, Switzerland: Springer.

Garson, J. [2019]: *What Biological Functions Are and Why They Matter*, Cambridge: Cambridge University Press.

Garson, J. and Piccinini, G. [2014]: ‘Functions Must Be Performed at Appropriate Rates in Appropriate Situations’, *British Journal for the Philosophy of Science*, **65**, pp. 1–20.

Godfrey-Smith, P. [1994]: ‘A Modern History Theory of Functions’, *Noûs*, **28**, pp. 344–62.

Griffiths, P. E. [1993]: ‘Functional Analysis and Proper Functions’, *British Journal for the Philosophy of Science*, **44**, pp. 409–22.

Griffiths, P. E. [2006]: ‘Function, Homology, and Character Individuation’, *Philosophy of Science*, **73**, pp. 1–25.

Griffiths, P. E. and Matthewson, J. [2018]: ‘Evolution, Dysfunction, and Disease: A Reappraisal’, *British Journal for the Philosophy of Science*, **69**, pp. 301–27.

Hardcastle, V. G. [1999]: ‘On the Normativity of Functions’, in V. G. Hardcastle (*ed.*), *Functions: New Essays in the Psychology and Biology*, Oxford: Oxford University Press, pp. 144–56.

Hausman, D. M. [2011]: ‘Is an Overdose of Paracetamol Bad for One’s Health?’, *British Journal for the Philosophy of Science*, **62**, pp. 657–68.

Kingma, E. [2010]: ‘Paracetamol, Poison, and Polio: Why Boorse’s Account of Function Fails to Distinguish Health and Disease’, *British Journal for the Philosophy of Science*, **61**, pp. 241–64.

Kingma, E. [2016]: ‘Situation-Specific Disease and Dispositional Function’, *British Journal for the Philosophy of Science*, **67**, pp. 391–404.

Lauder, G. V. [1999]: ‘Homology, Form, and Function’, in B. K. Hall (*ed.*), *Homology: The Hierarchical Basis of Comparative Biology*, New York: Academic Press, pp. 151–96.

Lewis, D. K. [1983]: ‘New Work for a Theory of Universals’, *Australasian Journal of Philosophy*, **61**, pp. 343–77.

Lewis, D. K. [1986]: *Philosophical Papers*, Volume 2, Oxford: Oxford University Press.

Maley, C. J. and Piccinini, G. [2017]: ‘Unified Mechanistic Account of Teleological Functions for Psychology and Neuroscience’, in D. M. Kaplan (*ed.*), *Explanation and Integration in Mind and Brain Science*, Oxford: Oxford University Press, pp. 236–55.

McKinley, M. P., O’Loughlin, V. D. and Bidle, T. S. [2016]: *Anatomy and Physiology: An Integrative Approach*, New York: McGraw-Hill.

McLaughlin, P. [2001]: *What Functions Explain: Functional Explanation and Self-Reproducing Systems*, Cambridge: Cambridge University Press.

Meek, C. and Glymour, C. [1994]: ‘Conditioning and Intervening’, *British Journal for the Philosophy of Science*, **45**, pp. 1001–21.

Millikan, R. G. [1989]: ‘In Defense of Proper Functions’, *Philosophy of Science*, **56**, pp. 288–302.

Mossio, M., Saborido, C. and Moreno, A. [2009]: ‘An Organizational Account of Biological Functions’, *British Journal for the Philosophy of Science*, **60**, pp. 813–41.

Nanay, B. [2010]: ‘A Modal Theory of Function’, *Journal of Philosophy*, **107**, pp. 412–31.

Neander, K. [1991]: ‘Functions as Selected Effects: The Conceptual Analyst’s Defense’, *Philosophy of Science*, **58**, pp. 168–84.

Neander, K. [2002]: ‘Types of Traits: The Importance of Functional Homologues’, in A. Ariew, R. Cummins and M. Perlman (*eds*), *Functions: New Essays in the Philosophy of Psychology* *and Biology*, Oxford: Oxford University Press.

Neander, K. [2017]: 'Functional Analysis and Species Design', *Synthese*, **194**, pp. 1147–68.

Neander, K. and Rosenberg, A. [2010]: 'Solving the Circularity Problem for Functions: A Response to Nanay', *Journal of Philosophy*, **109**, pp. 613–22.

Parsons, T. [1990]: *Events in the Semantics of English: A Study Subatomic in Semantics*, Cambridge, MA: MIT Press.

Pearl, J. [2000]: *Causality*, Cambridge: Cambridge University Press.

Piccinini, G. [2020]: *Neurocognitive Mechanisms: Explaining Biological Cognition*, Oxford: Oxford University Press.

Reiss, J. and Sprenger, J. [2020]: 'Scientific Objectivity', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/win2020/entries/scientific-objectivity/>.

Rosenberg, A. and Neander, K. [2009]: 'Are Homologies (Selected Effect or Causal Role) Function Free?', *Philosophy of Science*, **76**, pp. 307–34.

Rubner, A. [2022]: 'Normal-Proper Functions in the Philosophy of Mind', *Philosophy Compass*, **7**, available at <doi.org/10.1111/phc3.12852>.

Saborido, C., Mossio, M. and Moreno, A. [2011]: 'Biological Organization and CrossGenerational Functions', *British Journal for the Philosophy of Science*, **62**, pp. 583–606.

Schaffner, K. F. [1993]: *Discovery and Explanation in Biology and Medicine*, Chicago, IL: Chicago University Press.

Schellenberg, S. [2018]: *The Unity of Perception: Content, Consciousness, Evidence*, Oxford: Oxford University Press.

Schlosser, G. [1998]: 'Self-re-production and Functionality', *Synthese*, **116**, pp. 303–54.

Schwartz, P. H. [2007]: 'Defining Dysfunction: Natural Selection, Design, and Drawing a Line', *Philosophy of Science*, **74**, pp. 364–85.

Sprenger, J. [2018]: 'Foundations of a Probabilistic Theory of Causal Strength', *Philosophical Review*, **127**, pp. 371–98.

Sprites, P., Gylmour, C. and Scheines, R. [2000]: *Causation, Prediction, and Search*, New York: Springer.

Walsh, D. M. [1996]: 'Fitness and Function', *British Journal for the Philosophy of Science*, **47**, pp. 553–74.

Walsh, D. M. and Ariew, A. [1996]: 'A Taxonomy of Functions', *Canadian Journal of Philosophy*, **26**, pp. 493–514.

Weber, M. [2005]: *Philosophy of Experimental Biology*, Cambridge: Cambridge University Press.

Weber, M. [2017]: 'How Objective Are Biological Functions?', *Synthese*, **194**, pp. 4741–55.

Wilhelm, I. [2022]: 'Typical: A Theory of Typicality and Typicality Explanation', *British Journal for the Philosophy of Science*, **73**, pp. 561–81.

Wimsatt, W. C. [1997]: 'Aggregativity: Reductive Heuristics for Finding Emergence', *Philos-*

ophy of Science, **64**, pp. 372–84.

Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.

Woodward, J. [2012]: 'Causation and Manipulability', in E. N. Zalta (*ed.*), *Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/win2012/entries/causation-mani/>.

Wouters, A. [2005]: 'The Functional Perspective in Evolutionary Biology', in T. A. C. Reydon and L. Hemerik (*eds*), *Current Themes in Theoretical Biology*, Berlin: Springer, pp. 33–69.

Wright, L. [1973]: 'Functions', *Philosophical Review*, **82**, pp. 139–68.