

Teleological functional explanations: a new naturalist synthesis

Mihnea Capraru

Abstract. The etiological account of teleological function is beset by several difficulties, which I propose to solve by grafting onto the etiological theory a subordinated goal-contribution clause. This approach enables us to ascribe neither too many teleofunctions nor too few; to give a unitary, one-clause analysis that works just as well for teleological functions derived from Darwinian evolution, as for those derived from human intention; and finally, to save the etiological theory from falsification, by explaining how, in spite of appearances, the theory can allow for evolutionary function loss.

Competing interests: There are no conflicts of interest to declare.

1 Introduction

Teleology occurs when means are explained by goals. This has put off many a naturalist. For one thing, it seems to involve backwards causation, i. e., causation of the past by the future. Worse, it seems to involve causation by things that don't exist, because means can be explained even by goals that are failed. This is why, after Darwin, many naturalists have denied any teleology outside of the mental purposes of humans and like-minded animals, and why they have maintained, of the apparent teleology in the living world, that it is nothing but the mistaken, anthropomorphic projection of mental purposiveness. Mental purposiveness, however, is all too easy to take for granted, and it is all too easy to forget that mental purposiveness *itself* must be naturalised, lest we are to concede that mind is not natural. As naturalists, therefore, we must account for mental teleology. It is attractive to derive mental teleology from a simpler, more general naturalist conception that subsumes teleology in mind and in biology. Toward this goal, philosophers and biologists have worked for more than seventy years. I aim to give such a general naturalization of teleological function, robust enough to weather the difficulties accumulated through the decades, yet intuitive enough to count as an explanation, as an answer to the metaphysical question, what a teleological function is.

When something has a teleofunction, that teleofunction consists in causing an effect of a certain type *E*. This doesn't mean, however, that every time some-

thing causes an *E*-type effect, that thing has the teleofunction to do so. This is why we are going to distinguish, with Karen Neander (1991), between teleofunctions and causal-role functions. When we ascribe to a thing a function, we give an implicit explanation in terms of the types of effects that the thing causes. What it is that we explain, though, varies with the kind of function we ascribe.

When we identify an effect-type as a *causal-role* function of a system component, we explain how that system comes to have the capacities that it has (Cummins, 1975).¹ On the other hand, when we identify an effect-type as a *teleofunction* of a thing *x*, we explain not the capacities of *x* or of systems containing it, but rather we explain *x* itself. What, though, is the precise sense in which we explain *x*? Common sense does not seem to be precise enough in this regard; instead, we must answer this question by developing a satisfactory theory of teleofunction. Since Larry Wright's essay of 1973, the prevailing view has been that when we ascribe a teleofunction to a thing *x*, we explain *x* in the sense that we causally explain *x*'s existence. When researchers suggest, for instance, that the cecal appendix has the function to shelter beneficial, symbiotic bacteria during gut infections, they purport to explain causally why the cecal appendix exists (Smith et al., 2013).

This view is known as the etiological theory of teleofunction. Aside from

¹Cummins-style views are defended more recently by Paul Sheldon Davies (2001), Marcel Weber (2017), or Brandon Conley (2023).

Wright's own, further versions of the theory stem from Karen Neander (1991; 1991), Philip Kitcher (1990; 1993), and most sophisticatedly, Ruth Millikan (1984; 1989; 1993; 1999; 2002). Each of these theories aspires to explain both the teleofunctions of biological traits and those of technological artefacts. In the sphere of biology exclusively, further etiological theories come from Peter Godfrey-Smith (1994), Paul Griffiths (1993; 2006), and Justin Garson (2016; 2019); in that of technology, from Beth Preston (1998; 2009) and from Pieter Vermaas and Wybo Houkes (2006).

Arguably a further etiological approach is the recently popular organisational account of functions (e. g., Moreno and Mossio, 2015:63–87). To be sure, its proponents do not always understand it as etiological; their respective versions are reviewed by Garson in 2016, pp. 100–105. What all the versions have in common is the claim that if a trait t has a teleofunction f , that is because t does f , and by so doing t contributes to the persistence of a wider system s such that s contains t and the persistence of s ensures the persistence of t .

As Marc Artiga and Manolo Martínez argue in 2015, the organisational account cannot explain cross-generational functions, e. g., those of gametes, unless it collapses into the etiological theory. To this let me add that the organisational account does not seem to be able to explain the functions of traits pertaining to the extended phenotype.² If the parasite *Toxoplasma gondii* produces

²Dawkins, 1982.

a certain neuro-behavioural trait in its mouse host because that trait propagates the parasite, then that trait fails to contribute to the persistence of a wider system that contains it: it does not help the mouse survive, nor does it help the particular mouse–parasite system that contains it. Without *ad-hoc* modifications, the organisational account entails that such a trait has no function. But in all appearance it does: its function is plainly to propagate the parasite.

There has been, to be sure, a certain level of dissent over the etiological theory. Christopher Boorse (1976; 2002) has defended an older view, the goal-contribution view, on which teleofunctions do not explain the existence of their bearers, but merely the bearers' contributions to so-called 'goals' (on whose nature more later). John Bigelow and Robert Pargetter have proposed what is arguably a version of the goal-contribution view, namely, that we should assimilate the teleofunctions of biological traits to these traits' current contributions to the fitness of their organisms (1987).³ More recently, Bence Nanay has proposed that a thing has the teleofunction to cause *E*-type effects iff, were it to cause such effects, it would thus contribute to the fitness of its organism—if biological—or to the goals of its users—if an artefact (2010).

Such goal-contribution theories have been addressed by an extensive body

³Bigelow and Pargetter do acknowledge a certain derivative sense in which teleofunctions sometimes explain existence: namely, present biological-trait tokens often exist because *other*, prior tokens have performed their fitness-enhancing teleofunctions. However, according to Bigelow and Pargetter, there is no sense in which a trait token *x* exists because *it* (*x*) has the teleofunctions that it has (1987:194–95).

of literature, as old as Richard Taylor's 1950, and often containing formidable objections. Here is not the place to rehearse or to discuss these objections. Instead, I will make a conscious theoretical choice in favour of the currently prevailing view, to wit, the etiological theory. This is partly because I am persuaded by the extant objections to goal-contribution theories, but first of all, it is because I think we stand a better chance to naturalise teleofunctional explanations if we work from the etiological platform.

Let me clarify the sense in which we explain the existence of x when we ascribe teleofunctions to x . Take, for instance, the stomach, which has certain teleofunctions in digestion. All that this entails is that there is *some* causal history that explains the existence of present-day stomachs in terms of contributions to digestion. Aristotle or W. Paley would have disagreed with us—at least at first—about the nature of this history, but they would have agreed that *some* such history explains the stomachs' existence, and therefore they would have agreed that the stomach has digestive teleofunctions.

With the possible exception of Wright, all etiological theories share a crucial tenet underlined most forcefully by Millikan in 1984: teleofunctions are normative. When we ascribe to a thing x a teleofunction, we do not necessarily ascribe to x its actual effects; rather, we ascribe to x a more complex, normative relation between x and certain effects, effects that x may fail to cause actually. As Millikan illustrates (1984:34), we explain the existence of sperm by saying

that they are there to fertilise eggs, even though most never do. Likewise, and fortunately, most fire alarms never need to warn of dangerous fire.

Finally, teleofunctions have another important aspect—they are (usually) performed non-accidentally. Teleofunctions are not merely benefits that things provide by chance, as when Bertrand Russell's smoking habit saved his life by making him move to the smoking section of the plane whose nonsmoking passengers died in the subsequent crash. Rather than accidental, teleofunctions are benefits that things are expected to provide, so to speak, 'by design.'

In sum: Teleofunctions are actual or 'supposed' effects in terms of which we can explain the existence of the things to which we attribute the teleofunctions; they are normative; and they are non-accidental. Finally, if—as assumed—teleofunctions explain the existence of their bearers, then they are, very plausibly, etiological.

Extant etiological theories are beset, however, by several problems that threaten to undermine our hope of naturalising purpose, as it is difficult to see how to analyse teleofunction so as to solve all these problems at once. For brevity, we will refer to these problems as 'Goldilocks,' 'Unity,' and 'Function Loss.' Here is what they consist in.

Goldilocks: Some etiological theories attribute spurious teleofunctions while others exclude genuine ones, and none seems to get all the cases right. In particular, for biological teleofunctions many theories require that the relevant

etiologies involve natural selection. This, as we will see, is done to avoid an absurd proliferation of natural teleology. We can show, however, that natural selection requires too much and thus forces us to ‘undergenerate’ teleofunctions. This means that we must find a way to soften the natural selection requirement, without, however, overcorrecting in the opposite direction.

Unity: Etiological theories are either disjunctive, or limited in scope; that is, either they have separate clauses for the teleofunctions of biology and for those of human technology, or they address only one of these realms. (To be sure, this problem may seem less pressing than the other two. But it is conceptually important; from a conceptual standpoint, both biological and intentional teleofunctions are functions in the same sense.)

Function Loss: Notice that biological trait types can lose their teleofunctions, as with ostrich wings, which no longer have the teleofunction to enable flight. Yet the existence of ostrich wings *is* caused, in part, by their remote precursors’ successful use in flight. Although etiological theorists are aware of this, no one has yet explained how it is possible—if the etiological theory is right—for a trait to *lack* a teleofunction whose past performance is a cause of that trait’s existence.

To solve our three problems I will build into the etiological theory a subordinated clause inspired by an idea from the field of cybernetics (or systems theory). Fifty years ago, this idea used to inspire most naturalist analyses of teleology,

both by philosophers and by biologists (Perry, 1917; Perry, 1921; Rosenblueth, Wiener and Julian Bigelow, 1943; Braithwaite, 1946; Rosenblueth and Wiener, 1950; Nagel, 1953; Nagel, 1961; Nagel, 1977a; Nagel, 1977b; Boorse, 1976; Boorse, 2002). Although the cybernetic idea seems insufficient in its original form, it does contain an insight that will prove crucial.

Here is the cybernetic idea in short: Consider a target-seeking missile. Once fired, the missile will tend to meet its target no matter from what initial position it was fired, and even if deflected from its course. Furthermore, the missile has these tendencies not merely because of external forces, but in part because the missile is disposed to adjust to disturbances so as to compensate for them. This, of course, goes within limits—some initial positions will be too tricky to start from, some deflections too strong to overcome. Within such limits, though, the missile will behave in a way that may prompt the casual onlooker to say that the missile ‘seeks’ the target, that it ‘tries’ to find it, that it ‘pursues’ the ‘goal’ of finding it. Many cyberneticists and like-minded philosophers have understood this last clause literally; they have purported to identify teleology with such ‘seeking’ behaviour, and they have called the end-states of such processes ‘goals.’ Finally, philosophers of science have identified a thing’s teleofunctions with that thing’s contributions to such ‘goals’—most notably Ernest Nagel (1953; 1961; 1977; 1977) and Christopher Boorse (1976; 2002). On this view, the eye has the teleofunction to aid sight *because* that is

how the eye contributes to the 'goals' of its organism.

As I explained earlier, I do not subscribe to the goal-contribution theory of teleofunction. Nevertheless, I will adopt for modified purposes the cyberneticists' concept of a goal. Since I do not believe that such goals are teleological by themselves, I will rename them to 'homestates.' Furthermore, when processes are directed toward homestates, I will call such processes not 'goal-directed,' but as they are often called in biological contexts, plastic. (Compare: phenotypic, developmental, neural, behavioural plasticity.⁴)

Let me say a few more words about homestates. First, their name is a neologism that I intend as an allusion to pigeons' homing behaviour, a behaviour that illustrates the phenomenon. Travelling pigeons can be released at many remote places; they can also be thrown off course by strong winds or by electromagnetic fields interfering with magnetoreception; however, to many of the remote places, and to many of the perturbations, they will properly adjust so as to find home.

Second, despite the name, homestates are not always single states. In full mathematical generality, homestates are a species of what dynamic systems theory calls attractors. An attractor is a *set* of states such that, from numerous initial conditions, the system tends to wind up within that set (although the

⁴To prevent misunderstanding let me note that Nagel uses the same expression, 'plasticity,' in a subtly different sense, namely to denote just one aspect of what I call plastic or homestate-directed behaviour—the remaining aspect he calls 'persistence' (1977). In this paper I never use Nagel's terminology.

system doesn't need to settle on any particular state inside). Although every homestate is an attractor, not every attractor is a homestate. An attractor is only a homestate when the system tends toward it, partly, because the system is disposed to *adjust* plastically to the circumstances. More on this will follow on page 21.

Third, homestates are not restricted to the plastic system's internal states, but often consist in relations between the system and its environment. A traveling pigeon, for instance, has the homestate of being located at its geographical home, and this is not an internal state of the pigeon, but rather, a relation between the pigeon and a certain geographical location.

Fourth, there is not generally such a thing as *the* homestate of a plastic system. A dynamic system can have multiple attractors, and a plastic system, in particular, can have multiple homestates. For an antelope both safety and satiety are homestates. Homestates may be in conflict; food, after all, is not always safe to reach, and watering holes are dangerous. Evidently in such cases not all homestates can be reached fully and simultaneously.

Fifth, and finally: what kind of an entity can have homestates? Any entity that exhibits plastic behaviour. This can be an organism, or a group of organisms, or an institution, or a sophisticated artefact, or an ecosystem, etc. The list is open, as long as plasticity is present.

Here, then, is how I propose that we analyse teleofunction:

When an entity x has the teleofunction to cause effects of type E , this is in virtue of the fact that

- x 's origin has a causal explanation,
- and that explanation places x under a norm
 - to cause E -type effects,
 - and thereby to contribute to the homestates of x 's originators.

Allow me to clarify this formulation.

First, I am employing the term 'norm.' This needs to be explained, as we are giving a naturalistic theory. I am using the term as explicated by Millikan, in a broad sense which is unproblematically naturalistic: "Normative terms ... can indicate *any kind of measure from which actual departures are possible*. For example, a numerical average is also a 'norm'" (1999:192–93, emphasis mine). Most importantly, not every norm is an 'ought,' i. e., not every norm is prescriptive. Often I will use the phrase 'etioloical norm' to refer to the kind of norm at issue for us, i. e., to the norms under which things fall in virtue of the causal explanations of their origins. For instance, an animal's heart is under an etioloical norm to pump blood.⁵ This is because that heart's origin is causally explained, in part, by the fact that the animal has ancestors whose hearts have

⁵Of course some animals don't have hearts.

pumped blood and thus contributed to the ancestors' homestates, which, in turn, has eventually led to the animal's birth. Such etiological norms are not prescriptive.

Second, notice that we are talking about *an* entity, which means an individual object, or an individual token of a trait, or an individual event or state of affairs, and not a kind thereof. In particular, in biological contexts, we will speak of the origin of particular trait tokens, not of trait types. (E. g., Porky's having a snout is a trait token, namely, a token of the trait-type of having a snout.⁶) This matters, in particular, because trait types are often associated with teleofunctions that change with time. For instance, penguins' flippers do not have the teleofunction to aid flight, though some of their remote precursors did,⁷ while they do have the teleofunction to aid diving, which many of their precursors did not.

(I have encountered the worry that an etiological theory should not ascribe teleofunctions directly to trait tokens, but rather to trait types. This worry is motivated by Elliott Sober's well-known argument that natural selection explains

⁶In ordinary English usage, having a snout is called a trait, but the snout itself is not called a trait. However, I will often refer to individual organs as trait tokens. This will not affect our reasoning. We might as well speak of 'tokens of organ types,' but that would be awkward.

⁷These precursors, of course, were not flippers, nor were their organisms penguins. The long-term lineage that connects a present day token flipper to its remote non-flipper, non-penguin precursors is determined by the chain of local homologies

between one generation and the next: between any two generations, a token trait and its precursor are rendered homologous by the simple fact that they have the same non-functional, developmental explanation (Capraru, 2018). Hence these token-to-token homologies are independent of trait types, their functions, and their species.

statistical, population-level phenomena, rather than the properties of particular tokens.⁸ Sober's argumentation seems correct, but it does not affect our theory. As we will see in Section 3, not every etiological theory needs to appeal to natural selection as such. It is also possible to appeal to the particular history of fitness-enhancing events on the particular lineage of a particular token. It is this latter approach that we will take; hence it is safe for us to bypass the question what natural selection explains.)

Third, as is characteristic of etiological views, I maintain that Swampman's traits have no teleofunctions.⁹ For why this is less counterintuitive than it seems, I defer to Neander's argumentation from 1991a (pp. 179–80), which unfortunately there is no space to summarise here.

Finally, our analysis talks about x 's *originators*. x 's originators are, plainly enough, the entities from whose plastic activity x originated. For instance, many biological trait tokens are generated by their organisms during development, while many artefacts are generated by people acting intentionally. Notice that for x to be under an etiological norm to contribute to the homestates of x 's originators, it must be that x 's precursors have contributed to the homestates of *their own*—the precursors'—originators, and it must be that x 's origin is explained, in part, by the fact that the precursors' originators have thus reached

⁸Sober, 1984, pp. 135–155; Sober, 1995.

⁹Swampman is a hypothetical entity that resembles Donald Davidson molecule for molecule, yet came into being through an astonishing accident in a swamp (Davidson, 1987:443–44).

their homestates.¹⁰¹¹

By now we have introduced a certain amount of novel terminology. This is not necessarily desirable, but it is necessary. As I already explained, we need ‘homestate’ to avoid talking in terms of ‘goals,’ because the latter would suggest that our theory is a goal-contribution theory, when in fact it is etiological. On the other hand, we need ‘originators’ and ‘etiological norms’ as chunking abbreviations for the much longer expressions that define them. The need for ‘plasticity’ is similar, though this term is not a new coinage, but a loanword from biology.

In Sections 2, 3, and 4 respectively, I will show how we can now tackle the problems we labelled Goldilocks, Unity, and Function Loss. Thus we will be able better to grasp teleofunctional explanation, through a synthesis of our best naturalist ideas, a synthesis that subsumes plasticity and goal contribution under the normative aspect of the etiological theory.

¹⁰To be sure, the causal explanation needn’t involve the fact that these homestates *were homestates*; all that matters is that the homestates were reached.

¹¹I have encountered the assumption that my view is similar to Paul Sheldon Davies’s view from 2001. This is not the case at all. Davies denies the existence of teleofunctions, while I affirm it; he is a causal-role function ‘exclusivist,’ in the manner of Cummins, while I am a teleofunctional/causal-role pluralist, in the manner of Neander and Millikan.

2 Goldilocks

When thinking about teleofunctions in biology, it is natural to relate them to the cycles of biological reproduction. This is because, on the etiological view, teleofunctions explain existence. Contemporary eyes, for instance, exist because their possessors' ancestors used their own eyes to see, and thus survived and produced offspring. It seems natural to conclude that biological traits' teleofunctions coincide with these traits' past contributions to the realised fitness of their organisms, i. e., roughly speaking, to their reproduction.

Yet as we will see, things are not that simple. First, traits sometimes contribute to reproduction in accidental, irrelevant ways (Millikan, 1993:37–38). Second, there are many self-reproduction cycles that clearly do not involve any teleofunctions (more on this in Subsection 2.1). To solve these two problems, Millikan and others base their approach not on mere past contribution to reproduction, but on past natural selection (Millikan, 1993:33–39).

Despite the virtues of the natural selection approach, it has been shown to be too restrictive, perhaps for the first time by Philip Kitcher in 1993. If we require natural selection, then we ask too much and we exclude genuine teleofunctions. This leaves us in a bind: if we require only past contribution to reproduction, we overgenerate teleofunctions; if we require natural selection, we undergenerate them. Hence we must find a third way between natural selection and mere reproduction, a way to get teleofunctions just right.

2.1 Overgenerating teleofunctions: mere reproduction

Pointless cycles of self-reproduction. It is well-known that if we ascribed a teleofunction to every trait that has been contributing to the reproduction of its bearers, we would come to very implausible conclusions. Consider, for example, patients who harm themselves compulsively and thus cause endorphins to be released. Endorphins are natural painkillers, and they are rewarding much in the same way as their artificial analogues, the opiates. As they reward the patient, the endorphins lead to renewed self-harm; yet even though the endorphins thus contribute to the reproduction of self-harm, we do not regard it as their teleofunction to do so.

We can now see one reason why most etiological theorists follow Millikan (1993:35) and demand that teleofunctions stem from histories of natural selection. Self-harm-rewarding endorphins were never *selected* for their role in perpetuating self-harm, hence if we adopted the natural-selection requirement, we could avoid ascribing such spurious teleofunctions.

Accidental contributions to fitness. Porky's ancestor Eerie Pig was once attacked—let us imagine—by a wolf. In the moonlight Eerie's ears arose quite eerily and thus frightened the wolf, who scrambled to escape. This saved Eerie's life, enabled him to reproduce, and indirectly caused Porky's origin. Should we now say that Porky's ears have the teleofunction to scare the wolves away?

Ideally not. “How many [people] have to have been luckily saved from auto accidents by their love of *fast* speeds (surely that can happen too) before love of speed acquires a function?” (Millikan, 1993:38). According to Millikan, we can only answer such questions in a principled way if we require natural selection. Then, we could answer that love of fast speeds has the teleofunction to save drivers’ lives if and only if it has been selected for that. As for pigs’ ears, they were never selected for scaring away the wolves, and thus we don’t need to attribute to them this spurious teleofunction.

2.2 Undergenerating teleofunctions: natural selection

As we have seen, with a natural-selection requirement we could rule out two categories of spurious teleofunctions. Regrettably, however, we cannot adopt this strategy, because it excludes not only spurious teleofunctions but also genuine ones.

Recall that biological teleofunctions are determined not merely by history in general, but only by sufficiently recent history. Traits can lose teleofunctions they once had; ostriches’ and penguins’ wings do not have the teleofunction to aid flight, though long ago their precursors did.

The main objection to the natural-selection requirement plays on this: a teleofunctional trait can be maintained in its population by factors other than natural selection.

The oldest version of the objection stems from Philip Kitcher (1993:490–92): Suppose that moths in a certain population exhibit a wing pattern that helps them camouflage. This pattern was selected for in the past, but that was very long ago. Recently there have been a few moths with alternative patterns, patterns better at camouflaging. These moths, however, have died for unrelated reasons—perhaps someone sprayed insecticide in their area. Thus the original wing pattern has been maintained not by selection, but by accident. Yet although the pattern was not recently selected for, it still has the teleofunction to camouflage the moths.¹²

The non-selectional maintenance objection also has two other interesting versions in Peter Godfrey-Smith (1994:356–57) and in Peter Schwartz (2002:249). Godfrey-Smith has chosen to retain the natural selection requirement and to bite the bullet:

Perhaps traits are, as a matter of biological fact, retained largely through various kinds of inertia. Perhaps there is not constant phenotypic variation in many characters, or new variants are eliminated primarily for non-selective reasons. ... Then many ... function statements will be false. If functions are to be understood as

¹²One might have hoped that Millikan could avoid undergenerating functions by appealing to analogues of natural selection, as she often does for other reasons: e. g., cultural selection, or classical and operant conditioning (2004, chap. 1). But these alternatives cannot be deployed to tackle the non-selectional maintenance counterexamples, because what is missing in them is not merely natural selection, but any kind of analogous selection mechanisms whatsoever.

explanatory, in Wright's sense, there is no avoiding risks of this sort (Godfrey-Smith, 1994:356–57) .

I believe, however, that there is hope.

2.3 Getting teleofunctions just right: fitness and plasticity

We have seen in the previous subsection why we cannot use the natural-selection requirement, even though it would solve the problem of accidental contributions to reproduction and the problem of pointless self-reproduction cycles. We will therefore need to solve the two problems another way.

Let us start with the problem of pointless self-reproduction cycles. Recall how we are analysing teleofunction:

When an entity x has the teleofunction to cause effects of type E , this is in virtue of the fact that:

- x 's origin has a causal explanation
- and this explanation places x under a norm
 - to cause E -type effects,
 - and thereby to contribute to the homestates of x 's originators.

In the sphere of biology, reproduction is a homestate of living organisms. *This* is why, in biology, reproduction matters to teleology. On the other hand, in

the case of pointless self-reproduction cycles, reproduction is *not* a homestate of the relevant originators, hence it is not teleologically relevant. Therefore reproduction will give us biological teleofunctions, but it won't give us spurious pointless-cycle teleology.

Let us take these claims in turn.

Remember what we call a homestate. A homestate is a (set of) state(s) that a system tends to reach from diverse starting points and in spite of a range of perturbations, because the system tends to adjust appropriately to these various conditions. When a system is directed toward homestates, we call the system's behaviour 'plastic.' It is essential to plasticity that it involves *adjustment* in reaction to disturbances or even in anticipation. This was emphasised in 1917 by Ralph Barton Perry, who anticipated quite well the cybernetic, goal-contribution approach to teleofunctions. As Perry illustrates, when a candle flame is bent by the wind it tends to return to its upright shape, yet the flame does not behave plastically; the flame tends to regain its shape not because it adjusts to the wind, but rather because of external buoyancy forces. (This, indeed, is why not every attractor is a homestate. The point is particularly important because it resolves a commonplace objection to goal-contribution views, the ball rolling at the bottom of a bowl.)

Let us now look at biological trait tokens and see what originators they have and how they normally contribute to these originators' homestates.

First, biological trait tokens are usually generated by their own organisms, either during development or later. Organisms behave plastically, directed toward sundry homestates, most prominently toward development, survival, and reproduction. To these overarching homestates organisms subordinate many others, including mating, rearing the young, avoiding predation, foraging, cooperation, or capturing sunlight. When a trait token x is under an etiological norm¹³ to contribute to its organism's homestates by causing E , then x has the teleofunction to cause E .

Rather more remarkably, some trait tokens are generated not by their own but by other organisms (Dawkins, 1982; Griffiths, 1993). Consider the behaviour of rodents parasitised by *Toxoplasma gondii*. Such rodents seek out cats, get eaten by them, and thus enable the parasite to reproduce (Berdoy, Webster and Macdonald, 2000). This suicidal behaviour has the teleofunction to communicate the parasite to the cats. Notice, however, that this teleofunction relates not to the rodents' homestates, but rather to the parasite's, who counts as the relevant originator of this behaviour.

Some very simple organisms—perhaps retroviruses—might turn out to be so primitive that they don't pursue homestates, or at least not enough to ground all their teleofunctions. To anticipate this possibility let me argue that biological traits have one further homestate-directed originator: Darwinian evolution

¹³'Etiological norm,' recall, is shorthand for 'norm under which an entity falls in virtue of the causal explanation of its origin'; more discussion follows in Section 3.

itself. Recall that a system counts as having homestates if its behaviour is plastic. Evolution plastically renders organisms fitter and better adapted to their environments. (Note that this is *not* to say that evolution is teleological, not even in a naturalised sense. Recall that homestate-directedness is not teleological by itself.)

Take, for instance, any of a large number of plant species and transfer them to dry areas. Their leaves will adapt to lose less water, for instance by reducing their surface area, by thickening their skin, or by falling off during the dry season. Suppose, moreover, that we interfered and removed from these areas the best-adapted individuals. If so then we would slow down adaptation, but after a while it would get back on track. (Daniel McShea gives a related argument in 2012:677–78.)

One may object that there is no fact of the matter as to *whose* fitness evolution is directed toward. For instance, wolves and coyotes occupy similar ecological niches; both wolves and coyotes feed on small prey and on carrion (while wolves also feed on larger prey). When humans migrate into an ecosystem and reduce the wolves' fitness by hunting them down, coyotes are left with most of the wolves' food sources. Coyotes thus become fitter and take over the ecosystem. As we see, coyote fitness and wolf fitness are conflicting homestates, in the sense that if one goes up, the other is likely to go down. It may seem, now, that something is inconsistent here; but notice that I never said that evolution's

homestates are coherent. Evolution—which is not rational—is directed toward conflicting homestates, of which, of course, it is only likely to realise some.

In the living world, as we see, reproduction is not *mere* reproduction; it is a homestate of the originators of biological traits.¹⁴ Thus we can subsume biological cases under our general analysis of teleofunction.

Take, now, the pointless cycle that we gave as a counterexample to the *mere* reproduction requirement, namely, the cycle of self-harm releasing endorphins causing renewed self-harm. Although we may say that the self-harm has the teleofunction to make endorphins, few would say that *the endorphins* have the teleofunction to cause self-harm. Let us now observe a few things.

First, releasing endorphins is often a homestate for human organisms. All else being equal, humans seek high-endorphin states (but they do not only seek these, hence why they do not always have high endorphin levels). When patients become conditioned to seek endorphin releases, this is because endorphins can cause desirable mental states, such as pain relief, stress relief, or improved mood.

Second, self-harm is usually not a homestate for human organisms. However, it *can* become a homestate for a patient who does it compulsively. Such patients will often work around attempts to prevent self-harm. When self-harm

¹⁴A similar point was raised by David Buller in 1998, though Buller does not see things in terms of homestates, and does not attempt to analyse all teleological explanations, but only the evolutionary ones.

thus becomes a homestate, this is because self-harm leads to high endorphin levels, which in turn lead to desirable mental states, as explained above. Hence when a patient seeks self-harm as a homestate, this is because the patient seeks endorphins; but when the patient seeks endorphins as a homestate, this is not because the patient seeks self-harm, but rather because the patient seeks the ensuing mental states. Thus we find the function of endorphins in causing certain mental states, and we find the (non-Darwinian, conditioning-induced) function of self-harm in releasing endorphins. Present self-harming behaviour exists because its behavioural precursors have contributed to high endorphin levels; present high endorphin levels exist because their precursors have contributed to desirable mental states. Present high endorphin levels do *not* exist because their precursors have contributed to reinforced self-harm. Not that they haven't, but that is not why they exist; it is not self-harm that causes the patient to seek high endorphin states.

Having seen how to distinguish genuine from spurious teleofunctions in reproductive cycles, let us now address the problem of accidental contributions to reproduction.

As we remember, Porky Pig's ancestor Eerie had big ears that once scared a wolf away. Why don't Porky's ears now have the teleofunction to do the same? Recall that x falls under an etiological norm to cause E -type effects iff x falls under such a norm in virtue of the *causal explanation* of x 's origin. Hence it is

not enough for Porky's ears to be *caused* to originate by Eerie's lucky survival; their origin must thus be causally *explained*. This distinction makes the difference that we need, because not every causal factor is relevant to every causal explanation. Archduke Franz Karl married Princess Sophie and this led causally to the origin of their son, Emperor Franz Joseph. In turn, and along with the leaders of other nations, Franz Joseph played a (reluctant) role in causing the outbreak of World War I. Yet the marriage does not *explain* World War I. There is simply more to causal explanation than causation.

Back, now, to Eerie Pig. His unlikely survival participated in the causal process that led to Porky's origin. Yet his survival is not explanatorily relevant to Porky's origin, not any more than Sophie and Franz Karl's marriage is relevant to World War I. Hence it is not etiologically normal for Porky's ears to intimidate wolves. In general, when a trait token has contributed to reproduction in a way that we deem accidental, this is because the contribution is not explanatorily relevant to the origins of contemporary tokens of that trait. We can thus solve the problem of accidental contributions to reproduction.

On a final note, explanatory relevance comes in degrees, and so therefore does teleofunction. My smartphone has the teleofunction to play music, but not as much as it has the teleofunction to aid communication. If we should think that Eerie's adventure does explain the origin of Porky's ears a little, that would mean that Porky's ears do have *a little* of a teleofunction to scare the wolves. If

the wolves, furthermore, were scared away often enough, then relevance would build up and pigs' ears would acquire a full-blown second teleofunction.

3 Unity

Among the extant etiological theories of teleofunction, some fail to cover all the cases and others are disjunctive. It seems desirable, however, to give a unitary account that explains Darwinian and intentional teleofunctions equally well.

To this end, Neander identifies both Darwinian and intentional teleofunctions with what she calls 'selected effects.' Behind this unitary phrase, however, we have two very different things: Darwinian teleofunctions are selected in that they come from natural selection, while intentional teleofunctions are 'selected' in that they come from "a very different sort of selection process—the intentional design of an agent with forethought" (Neander, 1991a:174; a similar move is made by Daniel Dennett in 1995:407). Not only is the latter distinct from natural selection, but it often involves no selection at all, because agents can design artefacts without ever trying out alternatives or even contemplating them. Thus we seem to be gerrymandering the concept of teleofunction to cover natural selection and intentional human design.

In the philosophy of technology, Beth Preston, Pieter Vermaas, and Wybo Houkes avoid gerrymandering by limiting their objectives. Preston derives all

teleofunctions of artefacts from the cultural analogue of natural selection, thus rendering them all Darwinian (1998; 2003). Houkes and Vermaas, on the other side, derive all artefact teleofunctions from human intention (2003; 2006). Neither approach can do what we need. Preston cannot explain how *novel* artefacts can have teleofunctions, given that they have no history of cultural selection (Vermaas and Houkes, 2003:285). Vermaas and Houkes, in turn, cannot account for teleofunctions in biology;¹⁵ nor can they explain or make room for artefact teleofunctions derived from Darwin-like cultural evolution.¹⁶

Among proposals to unify Darwinian and intentional teleofunctions, the most complex is Millikan's. Millikan distinguishes three kinds of teleofunction (or in her terms, proper function). Let me illustrate the three with one of her favourite examples (Millikan 1984:39–45; 1999:198–204). A chameleon has neuro-physiological mechanisms that it uses to hide from predators by changing its colour to match the environment. These mechanisms have the *direct* proper function to hide the chameleon in its environment *in general* (no colour specified). When the chameleon enters a particular greenish environment, the mechanisms acquire the *adapted* proper function to hide the chameleon in this greenish environment. To this end the mechanisms produce a particular configuration of pigments. This pigment configuration, finally, has the *derived* proper

¹⁵To be fair, they aren't trying to.

¹⁶This manner of unintended artefact evolution was likely pointed out for the first time by David Hume in 1779.

function to hide the chameleon in the greenish environment. Millikan explains Darwinian teleofunctions as direct proper functions, and intentional teleofunctions as derived proper functions.

According to Millikan, direct proper functions are grounded in histories of natural selection. In brief, a trait's direct proper function is whatever it was selected for doing. That, however, is not how adapted proper functions are grounded. Assume that a chameleon originates from a population that never underwent natural selection in environments of a particular shade of green. Once the chameleon comes to such a place, however, its hiding mechanisms will still have the adapted proper function to hide it against that particular greenish background.¹⁷ The same goes for the greenish pigment configuration; therefore derived proper functions, too, are grounded differently from direct ones. This point is best made by Preston, who concludes that "selection in the case of derived proper functions does not necessarily have *any* of the characteristic features of selection in the case of direct proper function" (1998:233).

The crux of the objection is that Millikan's account does not explain what renders direct and derived proper functions the same kind of function. Direct proper functions result from a distal explanation by natural selection. Derived proper functions result from the proximal activity of a mechanism with another proper function of its own. What though, is it that renders such a derived proper

¹⁷We can see this even more vividly in the case of the octopus or the cuttlefish, who can camouflage themselves in a very large number of novel environments.

function – a proper function? If this is to be understood simply as a separate, independent clause in Millikan’s account, then the account is disunified. But if the derived function of the chameleon’s pattern is a logical consequence of the direct function of the chameleon’s brain structures – together with what Millikan calls the normal explanation for the performance of the brain structures’ function – then it ought to be contradictory to assert that a) the brain structures have the function to camouflage the chameleon, b) this function has normally been performed by producing patterns that matched the local backgrounds, but c) the present pattern itself lacks a function. Now I freely admit that c) is counterintuitive, because we already agree pre-theoretically that the pattern does have a function. But there is no apparent logical contradiction in asserting c) after a) and b). Thus it seems that derived proper functions are an independent clause in Millikan’s theory, hence there is strictly speaking no such thing as ‘proper function’ as such, but only direct function and derived function, with the latter being defined ultimately in terms of the former. (Evidently it is possible to add that a proper function is, by definition, *either* a direct function *or* a derived function, but this would only underscore the disunity.)

Perhaps in the chameleon’s case one may try, somewhat counterintuitively, to attribute the function not to the skin pattern itself, but to the broader matching configuration between skin and background. Similar matches, though with different colours, have occurred in the past, and because these matches have

saved the respective chameleons, they have been reproduced. It would seem, then, that unity might be thus saved, at the cost of quite a bit of revisionism: it would no longer be the device itself that has a function, but only the broader configuration that is reproduced.

This revisionary move, however, fails as soon as we consider something as simple as a pen. A pen is disanalogous from a chameleon colour, in that a chameleon colour is produced directly in the environment that it matches, whereas a pen is produced at one time, by one factory, at one place, only to be used at other times, often decades later, by unknown persons, on other continents, and perhaps in the future on other planets. Hence in the pen's case there is no particular device–environment match that is reproduced at all: when the pen is produced, it matches nothing in particular.

Given this disunity, even though our solution perhaps owes more to Millikan than to any other philosopher, we need to analyse teleofunction in a slightly different way.

3.1 Millikan norms

Consider once again Porky, whose ears originate from a reproduction cycle during which the ears' precursors contributed to their originators' homestates by helping these originators hear. Notice that this cycle determines a norm in the broad sense of a “measure from which actual departures are possible”

(Millikan, 1999:192–93). When seen in light of their origin, it is normal for Porky’s ears to help Porky hear.

Let us call this kind of norm a Millikan norm. The concept of a Millikan norm is inspired by Millikan’s own “normal explanations and normal conditions for direct proper functions” (1984:33–34).¹⁸

One may wonder just why we should hold Porky’s ears as bound by the Millikan norm to help Porky hear. What, in other words, is the naturalistic explanation for why Porky’s ears ought to fulfil this particular norm? This worry, though, is unmotivated; the worry stems from the assumption that all norms are prescriptive, and that assumption is inaccurate. For instance, one way to be normal is to be average, yet this norm is not prescriptive, and people often think it better to be above average and thus abnormal. Likewise, we are not *demanding*, prescriptively, that Porky’s ears obey the norm to help him hear. What we are saying is this: that it is normal, among other things, for Porky’s ears to help him hear, and that this is normal in virtue of how the ears’ origin is causally explained.

How would we ideally find out the Millikan norm for a trait token *t*? We would begin by reconstructing its entire precursor set; for instance, we would specify the set containing the ears of Porky’s mother, Porky’s father, and of

¹⁸In the interest of clarity, however, note that Millikan does not define her norms for teleofunctions as such, but a) for explanations of the past performance of teleofunctions, and b) for the environmental conditions involved in such explanations.

their own progenitors, etc. This precursor set is finite, but usually large. We would now ascertain what effects these ears and proto-ears have caused that have contributed to the survival of their organisms, and thus to their own proliferation. Often we will find, of course, that they have aided hearing. They have also aided balance, if we are including the inner ear. We may well discover further uses in the distant past. We would now take the set of all these past beneficial effects, and weigh it according to their causal proximity to Porky's present-day ears. That is, we would weigh the ears of Porky's proximal ancestors more heavily than those of his remote ancestors. The underlying intuition is that we are asking what is normal at the time at which Porky's ears exist, as opposed to what may have been normal at other times. (Likewise, what is normal at a spatial location depends more on its proximal vicinity than on its distal surroundings.)

We would now count as Millikan-normal any effects whose frequency, thus weighted, is non-negligible.¹⁹

3.2 Intentional teleofunctions

Unlike Darwinian teleofunctions, intentional teleofunctions stem not from natural history but from the intentions of rational agents. Most naturally, inten-

¹⁹What counts as non-negligible is, of course, somewhat vague around the edges, as with many other things.

tional teleofunctions are ascribed to intentionally designed artefacts. We can also ascribe intentional teleofunctions to intentional actions. (That is, to particular, token intentional actions.) Finally, we can ascribe intentional teleofunctions to other products of intentional action, including social groups and institutions that were set up intentionally, e. g., football teams or colleges.

When an entity x has an intentional teleofunction F , x 's origin is causally explained by the intention of an originating agent, to wit, the intention that x perform the function F and thereby contribute to the originating agent's homestates. The agent, thus, is the relevant originator of x , and the relevant etiological norm is to fulfil the agent's intention. (Notice that it is not artefacts, but agents, who need to have homestates—since it is the agents who originate the artefacts; sophisticated artefacts do have homestates of their own, but that is inessential.)

Our job, however, is not done yet. It is easy to agree that intentions set etiological norms on their products, hence that intentional functions satisfy our analysis. Yet we aren't aiming merely to analyse teleofunction, but also to show that teleofunctional explanations are naturalistically acceptable. This means, in turn, that we must show that intentional etiological norms are naturalistically acceptable.

I maintain that intentional norms are indeed naturalistically acceptable; this is because it is Millikan-normal for intentional norms to be satisfied. The origin of my intentions is causally explained by those intentions of my ances-

tors that got themselves fulfilled, not by those that failed. This is because my ancestors have normally survived by fulfilling their intentions and not by failing. (At times, of course, they may also have survived accidentally by failing to fulfil unwise intentions.) Then, furthermore, just as it is Millikan-normal for Porky's ears to help him hear, it is also Millikan-normal for my intentions to get themselves fulfilled. An intention is fulfilled, however, iff the intentional norm it generates is satisfied. Moreover, the intentional norm is satisfied *by virtue of* the intention's being fulfilled. This, then, is why it is Millikan-normal for intentional norms to be satisfied.

We encounter an important special case with commercial artefacts, whose originators are not their intended users. Take cars, which are generated by factory workers, designers, and engineers, but used by their owners, by taxi drivers and passengers, etc. In order to vindicate our analysis, we must explain how cars are under an etiological norm to contribute to the homestates of their *originators* by transporting their *users*.

That, though, is not hard to see. Cars are intended to contribute to their originators' homestates by transporting their users and thus: helping the originators acquire a good name and sell to more customers; helping them sell more cars to their extant customers; helping them to keep receiving wages or instalment payments; or not to pay damage; or simply to be honest and helpful. (Let us hope that most car makers do pursue these latter homestates too.)

Finally, let me underline an important advantage of our analysis. Notice that things can have more than one originator. A custom-made suit is originated by its tailor, but indirectly it is also originated by the customer who orders it. A bridge is originated by the workers, by the engineers who instruct the workers, by further engineers who design the bridge and by architects cooperating with them, and by the public authorities who commission the project. Answers to Socrates' questions are originated not only by the answerers, but also by Socrates himself.

This means that we can ascribe to things teleofunctions specific to particular originators. Take the bridge. For the public authorities, the bridge may have the teleofunction to connect two regions and thus to create economic activity; for the architects, it can have subtler, aesthetic teleofunctions; for some politicians, it can serve as an electoral argument. As we see, it is quite intuitive that things can have different teleofunctions for different originators. No previous accounts, however, seem to have addressed this aspect,²⁰ and it is an attractive feature of our analysis that it can do so.

²⁰The closest is Godfrey-Smith's 1994:349–50. That account, however, is only partial. For one thing, it addresses only Darwinian function, not intentional function. Accordingly, it only involves past fitness-contribution, not general homestate-contribution. For another thing, it requires the entities that bear teleofunctions to be parts of living organisms, or of other 'biologically real systems' subject to natural selection (Godfrey-Smith, 1994:350). This makes it impossible to apply the account to artefacts, etc.

4 Function loss

Biological trait types can lose teleofunctions. Even though our ancestors walked on all fours, our forelimbs no longer have the teleofunction to enable walking. This has been well appreciated by etiological theorists ever since Kitcher (1993), Griffiths (1993), and Godfrey–Smith (1994). What no etiological theorist has explained, though, is *why* a trait can fail to possess a teleofunction whose past performance is a cause of that trait’s existence. My arms exist, in part, because my four-legged ancestors were able to reproduce, and that in turn is because they were able to use their forelimbs to walk. Yet my arms don’t have a teleofunction to support walking. Some traits, therefore, do not have as a teleofunction *every* effect involved in their etiology.

This poses a significant yet under-appreciated challenge to the etiological theory. We must find a principled basis on which to say that some aspects of etiology count while others do not. Note that we cannot simply say that the etiology that counts is ‘recent,’ because what counts as recent varies from species to species and from trait to trait (Griffiths, 1993). Furthermore, recent or not, we must explain why the etiological theory is *not falsified* by the existence of traits that no longer possess every teleofunction involved in their etiology. It is not sufficient merely to exclude ancient history through *ad hoc* stipulation, because ancient events are also causes of the existence of present-day traits. If teleofunctions are determined by causal explanations of existence, and if certain

causes are ancient, then *prima facie* it seems that ancient history should matter.

Aside from the question how a teleofunction can be lost, there is also the question *when* a teleofunction is lost. As we will see, however, once we answer the question how we will also settle the question when.

Recall that we emphasised in Subsection 2.3 that etiological norms are not determined by causation as such, but rather by causal explanation. Thus even if x occurs in the causal history of y 's origin, if x is not explanatorily relevant to y 's origin, then it does not affect the teleofunctions of y .

Here, then, is how I propose to answer the question of function loss: Tokens of a biological trait still have the teleofunction F iff past performances of F are still relevant to the causal explanation of the origin of present tokens.

Let us look at ostrich wings, which no longer have the teleofunction to aid flight even though some of their remote precursors did. If we trace the causal history of ostrich wings, we will come upon flying birds. Yet these events are too ancient to explain their present-day consequences. As explanatory relevance wears off, teleofunctional relevance follows. This is why ostrich wings do not have the teleofunction to aid flight.

Having seen how a biological trait can lose its teleofunction, let us turn to the question *when* this happens.

Paul Griffiths has proposed what seems to be the only answer to date, the Evolutionarily Significant Time Period theory (ESTP):

Define an evolutionarily significant time period for a trait *T* as a period such that, given the mutation rate at the loci controlling *T* and the population size, we would expect sufficient variants for *T* to have occurred to allow significant regressive evolution if the trait was making no contribution to fitness. A trait is a vestige relative to some past [teleo]function *F* if it has not contributed to fitness by performing *F* for an evolutionarily significant period. (Griffiths, 1993:417)

(Regressive evolution occurs when traits fade away after becoming useless, as with the eyes of troglobites—animals that spend their entire lives in caves.)

Unfortunately, it doesn't seem that the ESTP approach will work. Here is why:

First, if ESTP were right, then we should be able (fallibly) to expect every trait to regress significantly by the time it has lost all its teleofunctions. Such traits should only fail to regress in the context of "a lucky absence of mutations" (Griffiths, 1993:417). Yet there are more reasons why such a trait may fail to regress: For instance, the trait may be controlled by a gene *A* linked to a further gene *B* such that *B* is maintained for its own unrelated teleofunction. Or the trait may be controlled by a gene that pleiotropically (simultaneously) controls an unrelated phenotypic trait that is useful and therefore maintained.

Second, ESTP seems to establish no well-motivated connection between

teleofunctions and the time it takes to lose them. For instance, if a trait is controlled by genes with a very low mutation rate, then it will take a very long time for the trait to regress. But why should the mutation rate matter to whether the trait has lost its function? Furthermore, if the trait is not significantly selected against, then it can only regress by random drift. But suppose that the effective population size is very large. Then it will take a very long time before random drift has any significant effects; and once again, why should population size matter to whether the trait still has its function?

ESTP cannot give a working and well-motivated answer to the question how long it takes for a trait to lose a teleofunction. From our own standpoint, however, the answer is straightforward: A trait T will retain its teleofunction F for as long as its past F -performing tokens are still explanatorily relevant with regard to its present tokens, and no longer.

Explanatory relevance, of course, can depend on the explanatory project. Evolutionary theory is interested in the function of the ostrich wing, but it is also interested in its remote origins. Within one explanatory project, it makes no difference that ancient wings were used to fly; within the other, it does. Why, now, are ancient wings not explanatorily relevant to ascertaining the present function of the ostrich wing? Recall that teleofunction depends on a) a causal explanation which determines b) an etiological norm. In the case of Darwinian teleofunctions, the relevant norm is what we have called a Millikan norm. But

for the ostrich wing, it is not Millikan normal to aid in flight; recall that what is Millikan normal depends on causal proximity, in the same way in which what is normal at a spatial location depends on what is spatially proximal. If a token wing has causally remote precursors that have aided flight, these precursors can still play a role in explaining the ultimate origin of ostrich wings, but they no longer determine what is nowadays normal.

5 Conclusion

When we give a teleofunctional explanation of x , we explain why x exists in terms of the effects that x is under a norm to cause. It is not a trivial effort to fit such explanations into a scientifically informed worldview. For the past seven decades, philosophers and biologists have pursued two prevailing strategies. At first, they tried to understand teleofunctions as contributions to cybernetic ‘goals,’ i.e., to what we have called homestates. Later, most abandoned the homestate-contribution view and replaced it with the etiological view, on which the teleofunctions of x are determined by the causes of x ’s existence.

The etiological theory faces several difficulties of its own. I have shown how to resolve these problems by grafting onto the etiological theory a subordinated homestate-contribution clause. This approach enables us to ascribe neither too many teleofunctions nor too few; to give a unitary analysis that

works just as well for teleofunctions derived from Darwinian evolution and for those derived from human intention; and finally, to save the etiological theory from falsification, by explaining how, despite the appearances, the theory can allow for evolutionary function loss.

As a welcome byproduct, we have been able to analyse the intuitive notion of a ‘function for,’ that is, the notion of things having different teleofunctions for different people, organisms, or social entities. This should allow us to clarify our understanding of the teleological features of actions, artefacts, and institutions stemming from complex societal histories.

References

- Artiga, Marc and Manolo Martínez (2015). “The Organizational Account of Function is an Etiological Account of Function”. In: *Acta Biotheoretica* 64.2, pp. 105–117.
- Berdoy, M., J. P. Webster and D. W. Macdonald (2000). “Fatal Attraction in Rats Infected With *Toxoplasma Gondii*”. In: *Proceedings of the Royal Society of London. Series B: Biological Sciences* 267.1452, pp. 1591–1594.
- Bigelow, John and Robert Pargetter (1987). “Functions”. In: *Journal of Philosophy* 84, pp. 181–196.

- Boorse, Christopher (1976). "Wright on Functions". In: *Philosophical Review* 85, pp. 70–86.
- (2002). "A Rebuttal on Functions". In: *Functions: New Essays in the Philosophy of Psychology and Biology*. Ed. by André Ariew, Robert Cummins and Mark Perlman. Oxford University Press, pp. 63–112.
- Braithwaite, R. B. (1946). "Teleological Explanation: The Presidential Address". In: *Proceedings of the Aristotelian Society* 47, pp. 1–20.
- Buller, David J. (1998). "Etiological Theories of Function: A Geographical Survey". In: *Biology & Philosophy* 13, pp. 505–527.
- Capraru, Mihnea D. I. (2018). "Note on the Individuation of Biological Traits". In: *Journal of Philosophy* 115, pp. 215–221.
- Conley, Brandon A. (2023). "Dispositionalism and Dysfunction". In: *Philosophy of Science* 90, pp. 686–703.
- Cummins, Robert (1975). "Functional Analysis". In: *Journal of Philosophy* 72, pp. 741–765.
- Davidson, Donald (1987). "Knowing One's Own Mind". In: *Proceedings and Addresses of the American Philosophical Association* 60.3, pp. 441–458.
- Davies, Paul Sheldon (2001). *Norms of Nature: Naturalism and the Nature of Functions*. MIT Press.
- Dawkins, Richard (1982). *The Extended Phenotype*. Oxford University Press.

- Dennett, Daniel C. (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. Simon & Schuster Paperbacks.
- Garson, Justin (2016). *A Critical Overview of Biological Functions*. Springer.
- (2019). *What biological functions are and why they matter*. Cambridge University Press.
- Godfrey-Smith, Peter (1994). "A Modern History Theory of Functions". In: *Noûs* 28, pp. 344–362.
- Griffiths, Paul E. (1993). "Functional Analysis and Proper Functions". In: *British Journal for the Philosophy of Science* 44, pp. 409–422.
- (2006). "Function, Homology, and Character Individuation". In: *Philosophy of Science* 73, pp. 1–25.
- Hume, David (1779). *Dialogues concerning natural religion*. Reprinted in *Dialogues concerning natural religion and other writings*, ed. Dorothy Coleman, Cambridge University Press, 2007.
- Kitcher, Philip (1990). "Developmental Decomposition and the Future of Human Behavioral Ecology". In: *Philosophy of Science* 57.1, pp. 96–117.
- (1993). "Function and Design". In: *Midwest Studies in Philosophy* 18. Reprinted in Allen, C., Bekoff, M. & Lauder, G., *Nature's Purposes: Analyses of Function and Design in Biology*, MIT Press, 1998, 479–503, pp. 379–397.
- McShea, Daniel W. (2012). "Upper-Directed Systems: A New Approach to Teleology in Biology". In: *Biology & Philosophy* 27, pp. 663–684.

- Millikan, Ruth G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. MIT Press.
- Millikan, Ruth Garrett (1989). "In Defense of Proper Functions". In: *Philosophy of Science* 56, pp. 288–302.
- (1993). "White Queen Psychology and Other Essays for Alice". In: A Bradford Book. MIT Press. Chap. 2, pp. 31–50.
- (1999). "Wings, Spoons, Pills, and Quills: A Pluralist Theory of Function". In: *Journal of Philosophy* 96.4, pp. 191–206.
- (2002). "Biofunctions: Two Paradigms". In: *Functions: New Essays in the Philosophy of Psychology and Biology*. Ed. by André Ariew, Robert Cummins and Mark Perlman. Oxford University Press, pp. 113–143.
- (2004). *Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT Press.
- Moreno, Alvaro and Matteo Mossio (2015). *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Springer.
- Nagel, Ernest (1953). "Teleological Explanation and Teleological Systems". In: *Vision & Action*. Ed. by Sidney Ratner. New Brunswick, New Jersey: Rutgers University Press, pp. 192–222.
- (1961). "The Structure of Science: Problems in the Logic of Scientific Explanation". In: New York & Burlingame: Harcourt, Brace & World, pp. 398–428.

- Nagel, Ernest (1977a). "Functional Explanations in Biology". In: *Journal of Philosophy* 74, pp. 280–301.
- (1977b). "Goal-Directed Processes in Biology". In: *Journal of Philosophy* 74, pp. 261–279.
- Nanay, Bence (2010). "A Modal Theory of Function." In: *Journal of Philosophy* 107, pp. 412–431.
- Neander, Karen (1991a). "Functions as Selected Effects: The Conceptual Analyst's Defence". In: *Philosophy of Science* 58, pp. 168–184.
- (1991b). "The Teleological Notion of 'Function'". In: *Australasian Journal of Philosophy* 69, pp. 454–468.
- Perry, Ralph Barton (1917). "Purpose as Tendency and Adaptation". English. In: *The Philosophical Review* 26.5, pp. 477–495.
- (1921). "A Behavioristic View of Purpose". In: *Journal of Philosophy* 18, pp. 85–105.
- Preston, Beth (1998). "Why is a Wing Like a Spoon? A Pluralist Theory of Function". In: *Journal of Philosophy* 95.5, pp. 215–254.
- (2003). "Of Marigold Beer: A Reply to Vermaas and Houkes". In: *British Journal for the Philosophy of Science* 54, pp. 601–612.
- (2009). "Philosophical Theories of Artifact Function". In: *Philosophy of Technology and Engineering Sciences*. Ed. by Anthonie Meijers. Handbook of the Philosophy of Science. Elsevier.

- Rosenblueth, Arturo and Norbert Wiener (1950). "Purposeful and Non-Purposeful Behavior". In: *Philosophy of Science* 17.4, pp. 318–326.
- Rosenblueth, Arturo, Norbert Wiener and Julian Bigelow (1943). "Behavior, Purpose and Teleology". In: *Philosophy of Science* 10.1, pp. 18–24.
- Schwartz, Peter H. (2002). "The Continuing Usefulness Account of Proper Function". In: *Functions: New Essays in the Philosophy of Psychology and Biology*. Ed. by André Ariew, Robert Cummins and Mark Perlman. Oxford University Press, pp. 244–262.
- Smith, Heather F. et al. (2013). "Multiple independent appearances of the cecal appendix in mammalian evolution and an investigation of related ecological and anatomical factors". In: *Comptes Rendus Palevol* 12.6, pp. 339–354.
- Sober, Elliott (1984). *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. University of Chicago Press.
- (1995). "Natural Selection and Distributive Explanation: A Reply to Neander". In: *British Journal for the Philosophy of Science* 46, pp. 384–397.
- Taylor, Richard (1950). "Comments on a Mechanistic Conception of Purposefulness". In: *Philosophy of Science* 17.4, pp. 310–317.
- Vermaas, Pieter E. and Wybo Houkes (2003). "Ascribing Functions to Technical Artefacts: A Challenge to Etiological Accounts of Functions". In: *British Journal for the Philosophy of Science* 54, pp. 261–289.

- Vermaas, Pieter E. and Wybo Houkes (2006). “Technical Functions: A Draw-bridge Between the Intentional and Structural Natures of Technical Artefacts”. In: *Studies in History and Philosophy of Science* 37, pp. 5–18.
- Weber, Marcel (2017). “How Objective Are Biological Functions?” In: *Synthese* 194, pp. 4741–4755.
- Wright, Larry (1973). “Functions”. In: *Philosophical Review* 82, pp. 139–168.

DRAFT