12

Neo-Teleology

Robert Cummins

1. Two Species of Functional Explanation

There are two subpopulations of functional explanation roaming the earth: teleological explanation, and functional analysis. The two are in competition. In this chapter, I hope to help select the latter, and nudge the former to a well-deserved extinction.

1.1. Teleology

Teleology is the idea that some things can and should be explained by appeal to their purpose or goal or function. It is, for example, the idea that one can explain why rocks fall and fire rises by appeal to the fact that the goal of matter is to go to its natural place, and that this is down for rocks and up for fire. It is also the idea that one can explain why (though not how) an acorn grows into an oak (rather than a beech or a clam) by appealing to the fact that the goal or function of a growing acorn is to become an oak tree. More plausibly, teleological explanation seeks to account for the existence or presence of a biological trait, or structure or behavior *by appeal to its function.* It is said that animals that have hearts have them because of what hearts are for.[1] Hearts are for circulating the blood; they are not for generating a pulse. Therefore, circulating the blood is their function, and they are 'there' – animals have them – because they perform this function.

1.2. Functional Analysis

Teleological explanations and functional analyses have different explananda. The explanandum of a teleological explanation is the existence or presence of the object of the functional attribution: the eye has a lens because the lens has the function of focusing the image on the retina. Functional analysis instead seeks to explain the capacities of the system containing the object of functional attribution. Attribution of the function of focusing light is supposed to help us understand how the eye, and, ultimately, the visual system, works. In the context of functional analysis, a what-is-it-for question is construed as a question about the contribution 'it' makes to the capacities of some containing system.

While teleology seeks to answer a why-is-it-there question by answering a prior what-is-it-for

Robert Cummins, "Neo-Teleology," in André Ariew, Robert Cummins, and Mark Perlman (eds.) *Functions: New Essays in the Philosophy of Psychology and Biology.* Oxford: OUP (2002), pp. 157–172.

I received many hours of feedback and help with this chapter from the members of my weekly lab group: Pierre Poirier, Jim Blackman, David Byrd, and Martin Roth. I want also to thank Denise Cummins, André Ariew, and Mark Perlman for their extensive comments, as well as an anonymous reviewer for the Press.

question, functional analysis does not address a why-is-it-there question at all, but a how-does-it-work question. These last are answered by specifying the structure (design) of the system. Rube Goldberg devices are natural candidates for this sort of explanation. In my horse pasture, I have a device that opens, at a pre-set time, a gate dividing the pasture in two. Here is how it works. There is a wind-up alarm clock. When the alarm on the wind-up alarm clock goes off, a string wound on the key-stem unwinds, releasing a ratchet on a pully. A weight on one end of a rope over the pully falls, jerking open the gate latch attached to the other end of the rope. This is, if you like, a rather abstract mechanical description. It is also a functional analysis of the capacity to open the gate as this is realized in my Rube Goldberg device. The components are identified functionally, and their interactions are described in a way that, necessarily, abstracts away from the medium-dependent details.[2] When we understand how the thing works in the way provided by a functional analysis, we understand how others might be built – how other instantiations of the same design could do the same job, and, perhaps, do it better. This is possible, because the system and its components are specified functionally, and hence in a way that allows for multiple instantiations. By substituting functional equivalents at various points in the design, taking care to accommodate the need for adequate interfaces with other components, we can make incremental changes in the system while preserving its overall viability. This is precisely how we must understand a system to see how it could be incrementally improved, and hence how it could evolve.

Of the two forms of functional explanation, I suspect teleology is much the oldest. Teleology is a natural framework for thinking about tools, cooking and storage utensils, and shelters. These ideas extend quite naturally to the body: eyes are tools or instruments for seeing, ears for hearing, hands for grasping, teeth and jaws for chewing. Functional analysis, on the other hand, got a grip on the mind, I suspect, only with the invention of relatively complex artifacts. Carts and harnesses lend themselves to functional analysis. Machines such as catapults and water clocks are unthinkable without it. This kind of thinking extends naturally to social structures such as bureaucracies, and to complex anatomical systems: the digestive system, the circulatory system, the nervous system.

There is more to proto-teleology than attributing functions to tools and sense organs, however. What I am calling teleology is the idea that an appeal to something's function can explain 'why it is there': why there are hammers and hands. If having a function is to explain why a thing, or type of thing, exists, then there must be some background story about a mechanism or process that produces the items in question, and produces them because of their functions. It is this requirement for what I will call a *grounding process* that has proved to be the Achilles heel of teleology.

Different kinds of phenomena subject to teleological explanation have required different grounding processes. Teleological mechanics appealed to the selective attractiveness of natural places. The intentions, plans, and actions of designers, creators, and manufacturers have been rung in to support teleological explanations of quite literally everything, and remain popular as underpinnings of teleological explanations of artifacts. In the hands of Aristotle, and Hans Dreisch (1867–1941), teleological developmental biology appealed to the regulating capacites of entelochies, a sort of inner goal-directed agent. Finally, natural selection has become a popular grounding process for the teleological explanation of biological traits, and sometimes for traits of artifacts as well.

Teleological explanation of motion failed because the grounding processes were transparently insensitive to function. Even if one could make sense of natural places, any force or mechanical constraint that would get something to its natural place would get it there whether or not it was the function of the thing to go there. To take an example from Ptolemaic astronomy, if a star has its apparent motion because it is attached to a rigid moving sphere, centered at the earth, it will trace a circular orbit around the center of the sphere regardless of what its function happens to be. The same point holds of Newtonian gravitational explanations of planetary orbits. Teleological appeal to functions in mechanics therefore appears idle and misleading. Indeed, it no longer seems plausible to suppose that celestial bodies have mechanical functions

at all. Like all teleology, teleological mechanics requires a grounding process. But the only grounding processes likely to satisfy render the appeal to functions utterly superfluous.[3]

Teleological explanation of growth and development fared even worse. The need for a grounding process spawned vitalism and the doctrine of entelochies. Entelochies could not be found. Moreover, appeals to them are regressive, since their own guiding and regulating behavior was itself teleologically explained, but without the hope of a corresponding grounding process. The whole misconceived enterprise rapidly became extinct when advances in cellular and molecular biology generated more adaptive theories competing for the same niche. Like post-teleological mechanics, those theories also appeal to factors that are transparently insensitive to the functions that were crucial to the teleological stories. So transparent is this, in fact, that it now seems silly to think it is the function of an acorn to develop into an oak rather than a birch, and that this explains why planting acorns never yields birch trees.

2. Neo-Teleology

Nobody much likes teleological mechanics or teleological developmental biology any more. It has been eliminated root and branch from mechanics and the other non-life sciences, recalled only by vestigial forms such as least energy principles that are without exception explained away as mere *façons de parler*. And it is similarly absent from developmental biology. But teleology survives in evolutionary biology, or anyway in the philosophy of it, as the idea that one can explain why an organism has a biological trait or structure by appeal to the function of that trait or structure. According to neo-teleologists, as I shall call them, we *have* hearts because of what hearts are *for*. Hearts are *for* blood circulation, not the production of a pulse. Hence, hearts are there – animals have them – because their function is to circulate the blood.

It is important to read the neo-teleologist claim transparently. Neo-teleologists hold that mammals have hearts because of something special that hearts do (or did). The heart, of course, does (did) lots of things. Among the things the heart does (did) is the thing we single out as its function, and it is *that* effect of heart presence – the one that counts as its function – that accounts for the presence of hearts. Neo-teleologists do not hold, as classical teleologists *did*, that circulation had its effect because it is or was the heart's function. They hold, rather, that the effect of heart presence that accounts for heart presence – call it e – accounts for heart presence because e has the property of being a circulating of the blood, not because e has the property of being the heart's function. Thus, neo-teleology is a mere shadow of the original (mistaken) idea, having only very limited aspirations. But, in spite of being widely influential, it is still mistaken in very much the same way that classical teleology was mistaken: the only plausible grounding processes render appeal to functions superfluous and misleading.

The idea behind neo-teleology is that evolutionary biology can provide the relevant grounding process and hence get you an answer to a why-is-it-there question from an answer to a what-is-it-for question. No doubt there is a sense of 'Why is that thing there?' that is just a way of asking what it is for. I point to a little rubber hemisphere on the carbureter of your lawnmower and ask you, 'Why is that thing there?' You reply by telling me its function – 'It is for priming the engine' – and this is an appropriate and satisfactory answer. But this only means that something that looks like teleology but is not can be had cheap.[4] What I am calling neo-teleology is more than this. It is the substantive thesis that, in some important sorts of cases at least, a thing's function – the effect we identify as its function – is a clue to its existence. If it is not to degenerate into the trivial thesis that 'why is it there?' can sometimes just mean 'what is it for?', neo-teleology must be the idea that, for example, there are eyes because they enable vision, wings because they enable flight, and opposable thumbs because they enable grasping.

3. Neo-Teleology and Natural Selection

Neo-teleology as just construed has no lack of able defenders (Millikan 1984; Neander 1991; Griffiths 1993; Kitcher 1993; Godfrey-Smith

1994; Allen and Beckoff 1995).[5] Generally, these authors are associated with selectionist etiological accounts of functions. Notice, however, that a defense of a selectionist etiological account of functions is, in effect, a defense of neo-teleology, since selectionist etiological accounts of functions equate functional attributions with what I am calling neo-teleological explanations: to say the function of the heart is to circulate the blood is, on these accounts, to offer a neo-teleological explanation of the presence of hearts. I prefer to attack the position by attacking the explanations in question rather than a thesis about what functions are, since evolutionary theory bears directly on the viability of the explanations, and only indirectly on the thesis that functional attributions are equivalent to such explanations.

Contemporary defenses of neo-teleology all share a basic selectionist strategy. The underlying idea is that traits are selected for because of the effects that count as their functions, hence exist in organisms because they have (or had) the functions they do (or did). Neo-teleology is thus packaged as what appears to be an uncontroversial part of the theory of natural (or artificial) selection. No natural places, entelechies, designer's intentions or other skyhooks (Dennett 1995) appear to taint neo-teleology; it is selectionist through and through. Since selectionist explanations are clearly legitimate scientific explanations, how could anyone object to neo-teleology? Surely we have here a subspecies of teleology that has found a legitimate grounding process.

And yet I am unpersuaded. Biological traits once explained by a teleology grounded in appeals to the intentions, plans, and actions of a creator have, in discerning minds, given way to appeals to evolution generally, and to natural selection in particular.[6] Neo-teleologists want to read this as the discovery of a legitimate grounding process for a teleological explanation of these traits. I am inclined to read the same intellectual development as analogous to what happened in mechanics and developmental biology: not a vindication but a replacement. The grounding processes of evolution, rightly understood, do not ground neo-teleology, because they are insensitive to function. Functions, I believe, have a legitimate place to play in science generally, and in biology in particular. But neo-teleology has the role of functions in selectionist explanations, and hence in biology, quite wrong: Biological traits, mechanisms, organs, etc., are not there because of their functions. They are there because of their developmental histories. Functions, I believe, enter into science legitimately as elements of functional analyses. Functional analysis is a powerful explanatory strategy that is widespread in all of the sciences. I have defended this view elsewhere (Cummins 1975, 1977, 1983), and will not comment on it further here. My focus in this chapter is rather to expose what I think are the vices of neo-teleology. An understanding of functional analysis will be relevant only because, as mentioned above, it appears to be precisely the framework we need to understand how complex systems could evolve.

4. Against Neo-Teleology

The basic idea of my argument is quickly conveyed. Traits are acquired in a variety of ways. Some are learned. Some, like sunburn, limb loss, and the effects of disease, are the direct result of environmental influences. None of these is of interest here because they are not heritable, hence are not subject to (non-cultural[7]) selection. The traits that *are* subject to selection *develop*. For convenience, in the rest of this chapter, 'traits' will be restricted to heritable traits the expression of which is the result of development and hence highly canalized.

Development is determined by a complex interaction between genes and environment. It is utterly insensitive to the function of the trait developed. Selection, on the other hand, is sensitive to the effects that are functions, but is, in the sense relevant to neo-teleology, utterly incapable of producing traits. It can preserve them only by preserving the mechanisms that produce them. Nor can selection, in the sense relevant to neo-teleology, produce the mechanisms that underwrite a trait's development; it can preserve only whatever mechanisms it finds already there.

I say selection cannot produce traits *in the sense relevant to neo-teleology*, for there is a sense in which selection can assemble complex traits or structures. This is what gives selection its awesome explanatory power. But the creative power

of selection is not the kind of process to which neo-teleologists appeal. I will return to this in a later section.

If the processes that produce traits are insensitive to their functions, how can functions account for why a trait is 'there' – that is, expressed in some specified population? The contemporary neo-teleologist answer is to concede that the processes that produce traits are insensitive to their functions, since, of course, traits do not have functions until *after* they are produced. But they argue that the processes that proliferate and preserve traits in a population are not insensitive to their functions. Certain traits spread through a population over time, and the mechanisms responsible are sensitive to function. Hence we can explain spread by appeal to function. Appeal to function thus gives us a handle on why a trait survived and proliferated, and hence a handle on why it is 'there'.

Imagine that crab grass invades a patch of Mendel's pea plants. The short ones will soon have trouble getting enough sunlight. The tall ones will do better. (They will all have trouble competing for root space underground.) The tall ones will reproduce more than the short ones, and will soon be far more common than the short ones, though they may be less common (per square foot) then either the tall or short ones previous to the crab grass invasion, and may eventually be crowded out altogether. In the meantime, tallness will, as we say, have spread through the population, and will be maintained.

This sort of story is supposed to explain why the pea plants in Mendel's crab-grass-infested garden are tall.[8] And it does. But how does neo-teleology get into the picture? Well, the idea is that the function of tallness in plants, or at least in these pea plants, is to achieve access to sunlight. Since gaining access to sunlight is what explains, via selection, why Mendel's pea plants are tall, we have explained why Mendel's pea plants are tall by appeal to the function of tallness in those pea plants.

5. Functions and Spread

The fundamental problem with neo-teleology is that traits do not spread because of (the effects that count as) their functions.

We can distinguish strong and weak variations of neo-teleology. The strong variation holds that any biological trait that has a function was selected for because it performed that function. The weak variation holds only that some traits were selected because of their functions.

A trait can be selected because of its function only if having that function counts as an adaptive variation in the population. For wings to be selected for because they enable flight, there must be a subpopulation in which wings enable flight, while wings in the rest of the population do not. For hearts to be selected for because they circulate the blood, there must be a subpopulation in which hearts circulate the blood, while hearts in the rest of the population do not. While it is plausible to suppose that there was a first flight-enabling wing somewhere among the ancestors of today's sparrows, those ancestors were not sparrows, nor was the wing in question anything like a contemporary sparrow wing. Similarly, somewhere in our ancestral line is to be found the first appearance of centralized blood circulation. But those ancestors were not even vertebrates, and the structures in question were nothing like our hearts.[9] It follows from these considerations that sparrow wings and human hearts were not selected because of their functions. Selection requires variation, and there was no variation in function in the structures in question, only variation in how well their functions were performed.[10]

Strong neo-teleology is refuted if there are legitimate targets of functional characterization that are not targets of selection. Strong neo-teleology must be rejected, since most, perhaps all, complex structures such as hearts, eyes, and wings patently have functions but were not selected because of (the effects that count as) their functions. And, since the selectionist etiological account of functions stands or falls with neo-teleology, it must be rejected as well, not because it is bad conceptual analysis (whatever that is), but because it equates functional attributions with bad evolutionary explanations.[11]

Weak neo-teleology survives this objection, but at a very considerable price. Weak neo-teleology comes out true only because of the rare though important cases in which the target of selection is also the bearer of a function that accounts for the selection of that trait. These will

be cases in which genuine functional novelty is introduced; a trait present in a subpopulation that is not just better at performing some function that is also performed in competing subpopulations (though not as well), but a trait that performs a function that is not performed at all by any counterpart mechanism in competing subpopulations. This unquestionably happens, and the importance of such seeding events should not be underestimated. But complex structures such as sparrow wings and human hearts were not introduced in this way. They were selected because they were better[12] at performing some function that was also performed by the competition. It follows from the equivalence of neo-teleology and selectionist accounts of function that these accounts will limit function attribution to those traits for which neo-teleology comes out true – namely, traits in which selection was triggered by the fact that the trait in question had a function that was entirely novel in the relevant population.

This is not merely a defense of gradualism. You do not have to be the village gradualist to be skeptical of the idea that there was variability in the presence or absence of T whenever T is rightly said to have a function. The point is rather that whether or not something has a function, and what that function happens to be, is quite independent of whether it was selected and spread. When we look for a place for selection to act on wings, say, we need to be looking for variations in wing design. All of the variant wings will have the same function – to enable flight. Thus, one cannot look to differences in the function of the wings to predict or explain selection. One must look instead to how well the various wings are functioning, and this means looking at the functions, not of the wings, but of something else: feather design, bone structure, musculature, and so on. Moreover, this argument iterates. It is the better of two muscle attachment schemes that gets selected; both the better scheme and the inferior scheme have the same function. Functions just do not track the factors driving selection. No doubt there *are* cases in which one subpopulation acquires some structure or behavior that the rest of the population just does not have, a biological analogue of adding a governor to steam engines, or an escapement to clocks. But such cases must be quite rare.[13] If they exhaust the proper domain of neo-teleology, then neo-teleology is insignificant at best. It comes out true as a kind of accident, a coincidence in the rare sort of case in which selective advantage happens to coincide with the introduction of something with a novel (in that context) function.

Selection can, to some degree at least, be explained by appeal to adaptiveness, although the connection between adaptiveness and selection is more indirect than is sometimes appreciated. What is uncontroversial is that a trait spreads because it is heritable and appears in a host that is more fit than the competition. Exactly the same thing can be said with equal truth about every trait of that host. Every trait of the winning host spreads, regardless of how adaptive it is – regardless, indeed, of whether it is adaptive (or has a function) at all. But this does not render adaptiveness irrelevant to selection, since the host in question was more fit than the competition because of some traits and in spite of others.[14] If H was a better design (in part) because of T, and all of H's traits spread because H was a better design, then T's positive contribution – its *adaptiveness*, in short – helps explain why it (and its neighbors) spread.

This suggests the possibility of saving neo-teleology by defining functions in terms of adaptiveness. This would turn neo-teleology into the idea that the proliferation and maintenence of some traits can be explained by appeal to the fact that they were adaptive. I certainly do not wish to take issue with that claim, though I think there are reasons for caution.[15] I do think, however, that there are good reasons to keep having a function and being adaptive distinct, and it is worth taking a brief detour to canvass these, for it will lead us back to the main point via another route.

Adaptiveness is a matter of degree; having a function is not. The more adaptive wing and its less adaptive competition both have the same function, but only the former is selected for. Functioning *better* is a matter of degree, and it is at least sometimes true that the more adaptive wing functions better. But this just makes it clear that functional analysis is prior to, and independent of, assessments of adaptiveness. When we have a system analyzed functionally, we are in a position to ask what sort of improvements could be made by substitution of functional equivalents.

The substitution of a functional equivalent that is (for example) more efficient, increases adaptiveness, but, by hypothesis, does not change anything's function.

The point here is not, as selectionist etiological accounts would have it, that only the selected wing has a function. After all, the worse wing was once the better one and was itself selected for. The point is rather that having a function is not what drives selection, but rather functioning better than the competition. What the function of a wing is should be distinguished from how well it performs it. The question of what function something has is evidently prior to the question of how well it is performed in a given organismic and environmental context, and hence prior to the question of how adaptive performing that function is for a given organism in a given environment. To repeat, the better and worse wings have the same function, but only the former spreads.

It might seem that there is a link of sorts between functions and adaptiveness, and hence between functions and selection. Knowing that the function of hearts is to circulate blood might be thought to constrain what sorts of variation in heart design would be adaptive, and hence what sorts of variations might be targets of selection. Indeed, I have been saying that it is the heart design that enables better circulation that gets selected. This suggests that when we identify the function of a trait, we have identified the dimension of performance that is relevant to assessing the adaptiveness of that trait. Circulation, not pulse production, is the function of the heart, and so it is variations in heart design that improve circulation, not variations that improve pulse production, that matter to adaptiveness.

Attractive as this line is, I do not think it will stand scrutiny. Better wing designs need not improve flight, but simply make it more efficient, or make development less error prone, or make the structure less fragile. Hence, selected changes in wing design that accumulate to yield the current design we seek to explain need not be related to the wing's function. Indeed, they may even compromise flight in the interest of other factors. Hence, if we are trying to understand why a given trait or structure is the way we find it, we cannot simply focus on variations that affect how well that trait or structure performs its function. We need, instead, to look at the complex economy of the whole unit of selection. This is precisely what a functional analysis of the whole unit facilitates, and is neglected when we focus on the function or functions of the trait in question.

6. Paley Questions

Even if we could make sense of the idea that things like wings and eyes – salient targets of functional attribution – spread through previously wingless and eyeless populations, the serious why-is-it-there question about such things as wings and eyes would remain untouched. How did there come to be such things in the first place? To harken back to Paley's famous example (1802), when we discover a watch in the wilderness, we are likely to infer a designer, not because we wonder why watches became so popular, but because we cannot otherwise understand how such a thing could come to exist at all.[16] And this is precisely the difficulty with eyes and wings. I propose to call this sort of why-is-it-there question a Paley question.

It is pretty generally conceded, I think, that Paley questions cannot be given neo-teleological answers (Godfrey-Smith 1994). Selection presupposes something to be selected. You cannot select for creatures with eyes unless eyes already exist. So it looks like selection cannot even address Paley questions. But, of course, this is much too quick.

Selection *can* address Paley questions, but only indirectly. The selection of eyes, or sighted organisms, is the wrong place to look. Selection builds a complex structure like a human eye or a sparrow wing by successive approximation (or what looks like it retrospectively) in relatively small steps beginning with an organism without an eye or wing and ending with what we observe today. Many of the fine details of such stories are unknown. Yet the in-principle possibility of the process is enough to provide the answer to Paley's original challenge: to explain how such things as the human eye came to exist in the first place without reference to the intentions, plans and actions of an intelligent creator and designer of eyes. Natural selection is clearly a central

player in the sort of story that has successfully met this challenge. But it enters in by accounting for the spread of small modifications to precursor structures. To think of the modern human eye or sparrow wing as itself selected is, to repeat, to conjure up a scenario in which there is a population of sightless primates or wingless songbirds into which is born a sighted or winged variation whose progeny take over the land or air. No one, of course, really believes anything like this. Yet something very like this is implied by neo-teleology – by the idea that eyes are there because they enable sight and wings because they enable flight. The modern human eye or sparrow wing never spread through any population. Some small changes to earlier structures very like the modern human eye or sparrow wing may have spread. And small changes to those structures may have spread. And so on. In short, as we have already seen, targets of functional characterization and targets of selection just do not match.

To summarize: if we ask why some complex structure is 'there', in the sense in which this means how it came to exist, appeal to its function or functions, as teleology (neo and classical) requires, is only going to be misleading. Such stories either run into the fact, fatal to classical teleology, that the crucial details of evolutionary (or ontogenic) development predate anything with the function that is supposed to do the explaining, or they founder on the fact that competing traits in selection scenarios typically have the same function. Things do not evolve because of their functions any more than they develop because of their functions.

It *is* generally conceded that teleology does not address Paley questions. But we are now in a position to see that Paley questions are all the questions there are about the evolution of traits. The idea that, although eyes and wings did not come to exist because of their functions, they nevertheless *spread* because of their functions, leaves us with a distorted picture of the role of selection. It makes us think that selection can spread only what is already there. While this is true in a sense, it is seriously misleading when we focus on the kinds of traits that have salient functions. It makes us think that eyes – eyes like ours – came to be somehow (some massive mutation?), and then were selected for because they were so adaptive. When we explain how eyes like ours came to be in the first place, we have said all there is to say about spread. When we have answered Paley's question, we have answered the evolutionary question. There is nothing left over for spread to do that it has not already done.

7. Conclusion

Let us consolidate our results. Neo-teleology, the idea that traits are there because of the effects that are their functions, is a non-starter when it comes to serious why-is-it-there questions: the questions I have called Paley questions. Appeals to function fail to address Paley questions, because nothing in the relevant lineage has the function in question until the trait in question is created. When it comes to Paley questions, neo-teleology has nothing to add to classical teleology. This is quite generally acknowledged. But neo-teleology fares no better as a story about why traits spread. Substantive neo-teleology misidentifies the targets of selection with the sort of complex generically defined traits – having eyes or wings – that have salient functional specifications.

Neo-teleology, I find, dies hard. Its rejection sounds to many like rejection of evolution by natural selection. But it is not. Darwin's brilliant achievement has no more need of neo-teleology than it has for its classical predecessor. What it needs is a conception of function that makes possession of a function logically independent of selection and adaptiveness. For it is only by articulating a reasonably illuminating functional analysis of a system that we can hope to understand *what* it is that evolution has created. If we want to understand *how* it was created as well, there is no avoiding the messy historical details by the cheap trick of assuming that all we have to do to understand trait proliferation and maintenance is to attribute a function. Neo-teleology thus amounts to a license to bypass the messy and difficult details, to jump over them in a way that makes it seem that the whole process was like the progress of a heat-sensing missile, arriving more or less inevitably at its goal regardless of the vicissitudes of wind and the meanderings of the target. The idea that evolution and development are goal oriented is precisely what makes classical teleology unacceptable. Neo-teleology creates the

same impression while masquerading as good Darwinian science.

There is another nexus of reasons why neo-teleology hangs on, at least in Philosophy. Twentieth-century empiricist philosophers such as Hempel (1959) were worried about function talk in science because it smacked of (classical) teleology. They set out to determine whether functions have a legitimate role in science. For reasons I am not clear about, they took this to be an issue about functional explanation, and interpreted *that* as a question about whether things could be explained by appeal to their functions. Thus, one important strand in the debate over functions simply assumed that the legitimacy of functions and the legitimacy of neo-teleology were one and the same. In Cummins (1975) I argued that this was a mistake; that functional attribution and functional analysis could be, and often are, decoupled from explaining why things are there by appeal to their functions. Still, the idea that functional attributions are equivalent to neo-teleological explanations remains widespread.

However, even if you accept that functional explanation and functional description can be decoupled from teleological explanation (some do, some do not), it might seem that the original empiricist worry remains about functions. One might continue to think that they need, in current parlance, to be *naturalized*. But most, perhaps all, of the pressure to naturalize functions is really pressure to naturalize teleology. Once functions are separated from teleology, they do not look any more likely to offend empiricist scruples than any other dispositional properties. But this point is not widely appreciated, and therefore there still is, I think, a widespread feeling that functions need naturalizing, and that this amounts to naturalizing (neo-)teleology.

There is a different sort of philosophical problem that remains, however. It is pretty generally agreed that a thing's function (or functions) is some special class of its effects. The problem of analyzing functional attributions, then, seems to require some criterion for saying which effects count as functions. Why is blood circulation a function of the heart and not production of a pulse? Selectionist etiological accounts seem to many to provide an elegant solution to this problem: the functions of an X are those effects of an X that, historically, account for Xs having been selected.[17] I have, in effect, been arguing against the selectionist etiological account of functions in biology on the grounds that the targets of functional attribution are seldom the targets of selection. If I am right, then almost nothing has a function in the sense staked out by selectionist etiological accounts of what functions are. This, I think, is what Hempel (1959) did conclude. We are better off abandoning the selectionist etiological account of functions.

Notes

1 Paul Davies (2001) holds that something's function should not be identified with what it is for, since this builds an unacceptable sense of 'design' – one involving intentional considerations – into the concept of function. I have some sympathy with this, but am prepared, for the purposes of this chapter, to let "What is it for?" be a way of asking the same question as "What is its function?"

2 Most causal analysis is like this. When we describe causal interactions between functionally characterized components, the relevant causal generalizations pretty much come for free, since a functionally characterized component is a component identified by its relevant causal powers.

3 It might seem that there is mostly a difference in attitude between saying that masses follow geodesics unless disturbed and saying that their function is to follow their natural paths. The standard contemporary reply to this sort of worry is to say that functions are normative, and, since there is no question of non-geodesic motion being a malfunction, there is no place for functions in mechanics. But this misses the point. The point is rather that the grounding process winds up accounting for the motion by appeal to factors such as forces or mechanical constraints that could not be sensitive to function in any case.

4 It might also mean that teleology used to be uncontroversial, so that the two expressions seemed to mean the same thing.

5 See Buller (1999) for a collection of papers defending and elaborating some version or other of what I am calling neo-teleology. Notice that a defense of a selectionist account of functions is, in effect, a defense of neo-teleology, since selectionist accounts equate functional attribution with neo-teleological explanation.

6 Teleological appeal to designers, creators, and manufacturers to explain artifacts is still widespread.

The function of an escapement in a clock is said to explain its presence in a way that is grounded in the intentions of designers, and manufacturers. However, selectionist treatments are also popular: the escapement is said to be there because it solved a problem plaguing pre-escapement clocks, leading consumers to prefer escapement clocks. The resulting pressures of the marketplace then led to the (near) extinction of pre-escapement clocks. (These, in turn, are in the process of being replaced by electronic clocks that require no escapement, of course.)

7 I am going to ignore cultural evolution in this chapter. I think the points I make here against selectionist defenses of neo-teleology would apply to neo-teleological stories about cultural selection as well, but I have not investigated this issue.

8 It does not, in my view, explain why any particular plant is tall. See Sober (1984, 1995), and Pust (2001). Neander (1995*a*, *b*) argues for the opposing view. And it does not explain why all the pea plants in the garden are tall, since short ones will continue to occur, though they seldom reach maturity.

9 Even these scenarios are misleading in suggesting that flight or centralized circulators appeared suddenly on the scene. Circulation was probably centralized gradually, and early flight was no doubt a matter of short and ill-controled forays into the air.

10 Another way of putting this point is that complex structures such as human hearts and sparrow wings are not heritable traits. What is heritable, at most, are variations in these traits. This follows from the fact that heritability is a measure of how much of the variance is accounted for by genes. Hearts in humans and wings in sparrows are not heritable because there is no variance to account for.

11 One could deny the validity of neo-teleological explanation and still hold that functional attributions were disguised neo-teleological explanations. Presumably, someone holding this position would advocate abandoning functional attributions. Perhaps Hempel (1959) is an example. But contemporary defenders of selectionist etiological accounts of functions think of themselves as vindicating functional attribution by identifying if with what they take to be a form of viable evolutionary explanation.

12 Even this is too strong, and will be modified shortly.

13 Mutation, for example, is much more likely to change the size, density, shape, or attachment angle of a bone than to add a new bone. The altered bone will typically have the same function as its competitors.

14 This is Sober's distinction (1984*a*) between being selected and being selected for.

15 One needs to be careful with the idea that traits spread because they are adaptive. The underlying rationale is that adaptive traits are likely to give their hosts the kind of advantages that lead to greater reproductive success. Hence, over the long haul, the subpopulation that has the trait in question is likely to grow relative to the rest of the population. The resulting *spread* of the trait in question through the population is the essence of selection.

Two points need mentioning. First, for this story to be substantive, we require a conception of adaptiveness that makes it independent of fitness. Secondly, whether adaptive traits spread depends on the extent to which conditions approach what we might call 'full-shuffle' conditions – i.e. conditions under which there is a fixed pool of heritable traits that do not interact, and every combination of them gets tried out in the fullness of time in a fixed environment. (See Kaufman 1989, 1993, on the importance of trait interaction.) That these conditions are seldom if ever satisfied in complex organisms is evident. The wonder is that natural selection works at all, given the poor working conditions with which it is faced.

16 Of course, if watches are popular, you are more likely to find one. But Paley's beachcomber did not want an explanation of why a watch was found, but of why there were any watches to find.

17 This is sometimes confused with the idea that the functions of X are those effects of X that were adaptive – i.e. that contributed to the fitness of their hosts. This could be true, even though Xs were not selected because of those effects, or even though Xs were not selected at all. Selection presupposes variability; positive contributions to fitness do not.

References

Allen, Colin, and Beckoff, Marc (1995), 'Biological Function, Adaptation and Natural Design', *Philosophy of Science*, 62: 609–22.

Buller, David (1999) (ed.), *Function, Selection and Design: Philosophical Essays* (Albany, NY: SUNY Press).

Cummins, Robert (1975), 'Functional Analysis', *Journal of Philosophy*, 72/20: 741–65.

Cummins, Robert (1977), 'Programs in the explanation of behavior', *Philosophy of Science*, 44: 269–87.

Cummins, Robert (1983), *The Nature of Psychological Explanation* (Cambridge, Mass.: MIT Press).

Davies, Paul S. (2001), *Norms of Nature: Naturalism and the Nature of Functions* (Cambridge, Mass.: MIT Press).

Dennett, Daniel (1995), *Darwin's Dangerous Idea: Evolution and the Meanings of Life* (New York: Simon & Schuster).

Godfrey-Smith, Peter (1994), 'A Modern History Theory of Functions', *Noûs*, 28: 344–62.

Griffiths, Paul (1993), 'Functional Analysis and Proper Functions', *British Journal for the Philosophy of Science*, 44: 409–22.

Hempel, C. G. (1959), 'The Logic of Functional Analysis', repr. in C. G. Hempel, *Aspects of Scientific Explanation, and Other Essays in the Philosophy of Science* (New York: Macmillan, 1965), 297–330.

Kauffman, S. A. (1989), 'Origin of Order in Evolution: Self-Organization and Selection', in B. C. Goodwin and P. Saunders (eds.), *Theoretical Biology: Epigenetic and Evolutionary Order from Complex Systems* (Edinburgh: Edinburgh University Press).

Kauffman, S. A. (1993), *The Origins of Order: Self-Organization and Selection in Evolution* (New York: Oxford University Press).

Kitcher, Phillip (1993), 'Function and Design', *Midwest Studies in Philosophy*, 18: 379–97.

Millikan, Ruth Garrett (1984), *Language, Thought, and Other Biological Categories: New Foundations for Realism* (Cambridge, Mass.: MIT Press).

Neander, Karen (1991), 'The Teleological Notion of "Function"', *Australasian Journal of Philosophy*, 69: 454–68.

Neander, Karen (1995*a*), 'Pruning the Tree of Life', *British Journal for the Philosophy of Science*, 46: 59–80.

Neander, Karen (1995*b*), 'Explaining Complex Adaptations: A Reply to Sober's Reply to Neander', *British Journal for the Philosophy of Science*, 46: 583–87.

Paley, William (1802), *Natural Theology: or, Evidences of the Esistence and Attributes of the Diety, Collected from the Appearances of Nature* (London: Faulder).

Pust, Joel (2001), 'Natural Selection Explanation and Origin Essentialism', *Canadian Journal of Philosophy*, 31 (June), 201–20.

Sober, Elliot (1984*a*), 'Force and Disposition in Evolutionary Theory', in C. Hookway (ed.), *Minds, Machines and Evolution* (Cambridge: Cambridge University Press), 43–62.

Sober, Elliot (1984*b*), *The Nature of Selection: Evolutionary Theory in Philosophical Focus* (Cambridge, Mass.: MIT Press).

Sober, Elliot (1995), 'Natural Selection and Distributive Explanation: A Reply to Neander', *British Journal for the Philosophy of Science*, 46: 384–97.