Functions in the biological realm: The function of language as a case study

Eran Asoulin

DRAFT ONLY, PLEASE DO NOT CITE.
COMMENTS WELCOME.

Abstract: I argue that talk of functions is a theory-internal matter, and so nothing has its function essentially. Functions are ascribed to suit the explanatory purposes of particular theories. A revealing example of the way in which this is done is the function of language. In contrast to the claim that the function of language is communication, I argue that what one sees as the function of language is dependent on the explanatory purposes of one's theory, so that it is perfectly reasonable for evolutionary biology and generative linguistics to have different conceptions of what the function of language is.

1 Introduction

I argue that talk of functions in the biological realm is a theory-internal matter, and so nothing has its function essentially. In other words, functions are ascribed to suit the explanatory purposes of particular theories, for they are explanatory tools. Debates about functional ascriptions often fail to recognise that the systems to which the functions have been ascribed do not have those functions essentially: a function is neither a natural kind nor an objective aspect of the world that we can pick out independently of our theoretical considerations. The function of the heart as a pump is not an activity that we can characterise separate from the aim of evolutionary biology to explain the emergence and persistence of the heart in the species. A revealing example of the way in which functions are ascribed to suit a theory's explanatory purposes is the debate in regard to the function of language. In contrast to the common and implicit working assumption that *the* function of language is communication, I argue that what one sees as the function of language is dependent on the explanatory purposes of one's theory, so that it is perfectly reasonable for evolutionary biology and generative linguistics to have different conceptions of what the function of language is. I argue that if this is the case in that different theories ascribe different functions to *the same trait*, then it follows that functions themselves are not had essentially but are theory-internal in the sense to be spelled out below.

2 What is a function?

Broadly speaking, there are two main ways in which to address the problem of functions, each of which will be addressed in the sections that follow. The first is the selected effects account, which is now the majority position in the philosophical literature on the topic. The selected effects account claims that function talk is justified because it is true of biological systems (Wright 1973; Neander 1991; 2017). The second way to address the problem of functions is the causal role account (Cummins 1975; Amundson & Lauder 1991; Davies 2001), which claims that function talk is justified, but not because it is true of biological

systems per se. Rather, we attribute functions in a way that best fits the explanatory purposes of our theory, but this does not mean that a particular function is the only or true function of the biological system under investigation. This is because, as I argue below with the case study of the function of language, different theories can offer different explanations of the *same* biological system and as a result end up attributing different functions to it. There is no sense in which one of these theories is attributing the wrong function to the biological system whereas a different theory attributes the right function – this is because functional attributions are theory-internal and thus bounded by particular explanatory contexts. In other words, there is no theory-neutral way by which to judge the fruitfulness or veracity of particular functional attributions.

The selected effects account of functions differs markedly from the causal role account: the former is teleological, the latter is mechanistic. Each account has a different understanding of the nature of functions, how they should be ascribed to a system, and what the aim of functional ascription is. I think that when one adjudicates the debate between these two camps, it is important to remember that before we can answer the question of whether function talk benefits a particular theory's explanatory capacity we need to know what the theory is trying to explain. If the theory attempts to explain the continued presence of a biological trait in the species, then a teleological functional ascription is useful and explanatory. But if the theory is attempting to uncover the underlying mechanisms in virtue of which the observed behaviour of an organ is made possible, then a different functional ascription is required. In other words, I want to argue that the answer to the question *What is the function of X?* depends on the prior question of *What is the theory (that has X as one of its explananda) trying to explain?*

2.1 Functions are true of biological systems

There are a number of ways in which one can flesh out the claim that functions are true of biological systems, but what underlies them all is the claim that functions are true of biological systems in virtue of certain biological properties that we can, at least in principle, discover. A classic and oft-quoted exposition of the notion of function is Larry Wright's (1973; 1976), which now serves as the basis for the majority position in the philosophical literature on the topic. Godfrey-Smith (1993) remarked 25 years ago that much of the literature since Wright (1973) has been concerned with a refinement of and a convergence towards a view of functions that has Wright's idea at its core; the same is true today (see McLaughlin 2001 and Lewens 2004; 2007 for discussion). It is thus worth looking at the debate that Wright's paper initiated, for it sheds light on the details and motivations for the current selected effects account of functions. Wright argues that "conscious and natural functions are functions in the same sense, despite their obvious differences" (Wright 1973: 143). The task he sets for himself is to unify conscious and natural functions under the same analysis, thus showing that they are equivalent. He argues that an account of function is seriously mistaken if it cannot allow for such a unification. Wright says, for example, that the phrase *in order to* in functional ascriptions such as *The heart beats in order to circulate blood* is equivalent to the role that the same phrase plays in goal ascriptions such as *The rabbit is running in order to escape from the dog*. He argues that no analysis should begin by supposing that the two sorts of function are different. Moreover, he claims that on close inspection it is difficult to detect a difference in what is being requested in *What is the function of the human windpipe?* as compared with *What is the function of a car's exhaust pipe?* Thus, he says, we

should expect "functional ascriptions to be explanatory in something like the same way as goal ascriptions" (Wright 1973: 154).

Functional ascriptions are, on this view, explanatory in the following sense: to specify the function of, say, the kidneys, is to explain *why* humans have kidneys, it is to specify the *reason* humans have kidneys. This is an etiological analysis of functions, for it concerns an analysis of how a system we are studying got to where it is. As Wright (1973: 157) puts it, "it turns out that '*X* is there because it does *Z*' […] provides us with not only a necessary condition for the standard cases of functions, but also the kernel of an adequate analysis." This kernel has been further developed in the decades since, most notably by Millikan (1984; 1989), Neander (1991; 1991a) and Godfrey-Smith (1994; 1996), to form the current selected effects account of functions (see also Kaplan & Pigliucci 2001; Garson 2011; 2016).

Even though Wright argues that "natural and conscious functions are functions by virtue of their being the reason the thing with the function 'is there'" (Wright 1973: 164), he is of course not suggesting here that there is a parallel in natural functions to conscious intent or to a designer. His understanding of function is an historical one, according to which the difference (which he sees as being quite minimal) between natural and conscious functions is the sort of reason that is appropriate to the particular functional ascription. That is, the minimal difference Wright thinks exists between the two sorts of function is whether a conscious agent is involved or not. So when we are explaining the presence of a particular organ or artefact by appeal to function, "the overriding consideration" is that the function "must be or create conditions conducive to the survival or maintenance" of the organ or artefact (Wright 1973: 164). But the specific nature of these conditions, Wright argues, is "mere etiological detail" that does not affect the essential form of the functional explanation (Wright 1973: 164). In other words, the history of an organ or artefact can come about in a number of ways, some of which involve natural selection and others that involve conscious agents, but the structure of the functional explanation in which this history participates remains unchanged.

One of the best known refinements and elaborations of Wright's etiological analysis is Karen Neander's (1991; 2017), who argues that "biological functions are intrinsically and universally teleological", which is to say that "function attributions universally and intrinsically justify teleological explanations" (Neander 1991: 458). She claims that to investigate the function of a trait is to investigate the trait's history of selection. On this view, to use Neander's example, that the pouch of the koala has the function of protecting its young explains why koalas have pouches. This is because having a pouch was selected because its function increased the fitness of koalas by allowing them to better protect their young. Such "selected functions are not *as-if* properties, nor are they ascribed to *fictional components*. Many real components of organic systems really have selected functions" (Neander 2017: 1162, emphasis in original).

In addition, like Wright, Neander insists that the functions of artefacts and the functions of biological entities should be cashed out in the same way, the only difference being that in the former there is intentional selection and in the latter there is natural selection. Regardless of what kind of selection is involved, the etiological approach sees the function of something as the effect for which it was selected. Neander argues that teleological explanations are "a species of ordinary causal explanation", for in the case of natural selection, effects of past instances "causally contribute to increased replication of the

trait" (Neander 1991: 463). It is in virtue of this history of prior causes that something has a function and thus what makes teleological explanations explanatory. It follows, then, that an effect of a system that does not have a history of the required kind (or no history at all) cannot have a function ascribed to it.

One might ask why Neander and Wright insist on the identity of function qua conscious selection with function qua natural selection. Why should we assume that the two sorts of selection (and function) are the same? What is gained by arguing for their equivalence? I think the reason is that Neander and Wright want to preserve the intuitive distinction between function and dysfunction and the intuitive distinction between non-accidental effects (i.e., selected effects functions) and accidental effects (i.e., non-functions). That is, if one wishes to preserve and justify the intuition that there is a truth of the matter as to what the function of something is, then the best way to do so is arguably to find a way to give functions a normative reading so that it becomes legitimate to single out certain effects as something a system is *supposed to* do. For Wright, one of the fundamental distinctions of function talk is that of accidental versus non-accidental effects. He says that "making a throbbing noise is not a function of the heart, it is just something it does – accidentally" (Wright 1973: 144). That is, "the heart throb, [is] our paradigm of non-function" (Wright 1973: 148).

One of course cannot make such statements without knowing the stable or normal state of the system in question (if such states exist for the system). One needs to know, in other words, the normal-proper function of the system in order to ascertain whether a particular effect is dysfunctional. If we want to understand a particular dysfunctional effect, says Neander, we need a description of the normal state of the system, from which this effect has deviated. Accordingly, the attempt to explain, say, normal human immunity or normal human vision "would be thrown into disarray if the function-dysfunction and accident-function distinctions were relativized to a researcher's idiosyncratic explanatory aims", for we "need to stabilize the notion [of function] and not leave it drifting in the breeze of shifting explanatory aims" (Neander 2017: 1161). In other words, the claim here is that there must be a truth of the matter as to what function a particular system has, and that this truth is not dependent on the theoretical or explanatory context. On this view, then, biological objects have their functions essentially: it is up to our theory to discover, not ascribe, their functions. If this is the case, then it is possible for a certain theory to attribute the wrong function to a biological object. That is, it follows from the selected effects account of functions that if a theory attributes to the heart the function of making thumping noises then it has made a mistake, for, as Wright argues, the heart throb is a paradigm case of non-function. However, if functions are relative to explanatory contexts, as I argue below, then it follows that the function-dysfunction distinction only makes sense within a particular explanatory context, and so what may be an accidental effect in relation to one theory might be a function in relation to a different theory.

2.2 Cummins-functions

The second way to address the problem of functions is the causal role account (Cummins 1975; Amundson & Lauder 1991; Craver 2001; Cummins & Roth 2009), which claims that we underwrite functional attributions by appealing to theory-internal explanatory contexts. This approach utilises an explanatory strategy where complex systems are explained in terms of their (usually simpler) constituent parts. In biology, explanations of organisms are

given in terms of a number of constituent systems, for example, the immune system, which in turn is analysed into constituent organs and structures. This strategy can be pursued until pure physiology takes over (Cummins 1975). Under this explanatory strategy (also called the analytical strategy) a function is understood in terms of an exercise of an analysed capacity *within a particular background of analysis.*

In other words, function-ascribing statements "make no provision for speaking of the function of an organism except against a background analysis of a containing system" (Cummins 1975: 763). We appeal to functions, says Cummins, in order to explain the underlying capacities of the system under investigation, and so talk of functions is appropriate only when we use the analytical strategy in this way. Another way to put the matter is thus: to ascribe a function to an item "is to ascribe a capacity to it which is singled out by its role in an analysis of some capacity of a containing system" (Cummins 1975: 765). This capacity can then be further analysed into a number of other capacities that together explain or manifest the former (higher level) capacity. The claim is that when "a capacity of a containing system is appropriately explained by analyzing it into a number of other capacities whose programmed exercise yields a manifestation of the analyzed capacity, the analyzing capacities emerge as functions" (Cummins 1975: 765). On this view, then, the reason why *The function of the heart is to pump blood* appears to give the correct functional attribution whereas *The function of the heart is to make thumping noises* appears to give the wrong functional attribution is that the usual explanatory context we work with includes "the background of an analysis of the circulatory system's capacity to transport food, oxygen, wastes, and so on, which appeals to the fact that the heart is capable of pumping" (Cummins 1975: 762). The function of a particular organ depends on the analytical context that relativises function-ascribing statements to the background of analysis (see also Hardcastle 1999; 2002). As I argue below, this is clear in the case of the function of language, where the background of analysis dictates what function is ascribed.

Notice that this way takes the nature of scientific explanations, and the role of functions therein, to be something quite different from the selected effects account. The latter claims that functional ascriptions are explanatory in the sense that to specify the function of, say, the kidneys, is to explain *why* humans have kidneys, it is to specify the *reason* humans have kidneys. The causal role approach, on the other hand, understands the nature of scientific explanation to be that of unearthing the mechanisms in virtue of which the phenomenon under investigation is made possible.

This way of understanding functions does not see functional analysis as explaining events by subsuming them under laws. Rather, functional analysis explains the dispositional properties of complex systems. Specifically, it "explains how a complex system works, and, consequently, why it has the property that is the target explanandum", but it "does not do this by identifying the causes of the system's acquisition of that property, but by specifying the abstract design of the system, a design the having of which amounts to having the target property" (Roth & Cummins 2014: 784). This is of course a kind of mechanistic explanation. I think that one can make the strong claim that a significant number of scientific theories explain phenomena by unearthing the mechanisms that produce them. On this view, one of the major aims of science is the discovery of mechanisms rather than laws (Thagard 2012). Machamer et al. (2000) argue that much of the practice of science can be understood in this way, and they stress that this is a different project to that of the discovery of laws. They give an example from biology according to which if a single base were changed in DNA and the

mechanism of protein synthesis operated as usual, then a counterfactual would be supported. "No philosophical work is done," they say, "by positing some further thing, a law, that underwrites the productivity of activities" (Machamer et al. 2000: 8). Activities are constitutive of mechanisms, and it is they that make phenomena intelligible. In other words, it is not regularities or laws that explain. Rather, what does the explaining are the mechanisms in virtue of which such regularities are made possible. Indeed, as Craver and Darden (2013) show, biology is in large part a search for mechanisms.

It should be noted that mechanistic explanations are not reductive explanations – one cannot use them to deductively predict from a lower level what will occur at a higher level. The decomposition into mechanisms (and into mechanisms of mechanisms) preserves the higher levels, and a mechanistic explanation would be incomplete without a hierarchy of levels. In other words, multiple levels are required in order to properly explain a particular phenomenon, and it is the integration of different levels that makes phenomena intelligible. A description of a mechanism is also more than merely a list of entities and their activities, for it involves not only such a list but also how the entities and activities are organised together actively, spatially, and temporally such that they yield the capacity under investigation (Craver 2001). Moreover, it is striking that despite their prevalence in science, mechanistic explanations do not have the status they deserve in the philosophy of science. Bechtel (2009) shows how biologists and psychologists rarely make use of laws in giving explanations, and in the relatively few cases in which they do the laws tend to be those of physics or chemistry (see also Bechtel 2008). In the case of biology, there is an "ubiquity of references to mechanism" and a "sparseness of references to laws" (Bechtel & Abrahamsen 2005: 423). Cummins (2000) speaks of the scandal in regard to the widespread belief that scientific explanation is subsumption under law: laws tell us what something does, not how it does it.

This discussion should not be taken to mean that the selected effects account of functions can or should be ruled out. Rather, what I want to highlight here is the crucial link between functional attributions and explanatory aims. Thus, as I show below in the case of the function of language, if we want to know how an organ works then the mechanistic understanding of functions (the causal role account) is best. However, if we want to know why a particular organ has remained in the species, then the selected effects understanding of functions is best. Construed in this way, etiological accounts of function are not compatible with causal role accounts of functions, for it is not the case that everything that is etiologically functional is also functional in a mechanistic or casual role sense. As was noted early on by Cummins (1975) and others, what we have here are two different modes of explanation applying different methodologies. As Millikan remarks, "it is important to see that Cummins' sense of function does not fit *with the theory of natural selection* to help explain the current presence of a trait in a species" (1989a: 175, emphasis in original), nor does it aim to. We should not expect the two senses of function to be (nor should we strive to make them) compatible. Godfrey-Smith (1993) argues that we should keep both senses of function strictly distinct and that any attempt to make one sense of function apply to both evolutionary biology and physiology is misguided: each field is asking different questions and offering different types of explanation. This is well understood by physiologists: Roux (2014), for example, argues that since selected effects accounts "restrict the functional attribution of a trait to its past selective value and not its current properties, these theories are inconsistent with the concept of function in physiology" (Roux 2014: 2245). A more

adequate account for physiology, Roux argues, is the causal role account of functions. Thus, there is no need (and often no way) to unify the different uses of the notion of function, for each notion of function was developed to suit the explanatory needs of the theory in which it is used.

3 Functions are attributed to suit theoretical explanatory purposes

Arguing that, say, the heart's function can be relativised to the context of analysis is of course not new. Frankfurt and Poole (1966) argued that the way in which we decide on the usefulness of a biological item partly depends upon the context in which this item is used. So in one context we say that the usefulness of the heart consists in its ability to pump blood, but that is not the only context in which the heart proves useful. That is, "the present environments of many vertebrates include physicians, and the practice of physicians involves making diagnoses which often rely on the character of their patients' heart sounds", and thus the "patients presumably have better chances of survival and of reproducing if this diagnostic technique can be used on them than if, *ceteris paribus*, their health could not be evaluated by listening to the sounds which their hearts make" (Frankfurt & Poole 1966: 72). In this context, then, the function of the heart in vertebrates *is* to produce heart sounds.

It is of course tendentious to argue that functional attributions are context dependent by listing examples where this can occur, for what is needed is a principled way in which to decide what an appropriate function is in a particular explanatory context. This does not mean that one can assign anything any function so long as an explanatory context exists. Rather, the fruitfulness of particular functional ascriptions is judged on how well the ascriptions contribute to the aims of the explanatory context, to the aims of what the theory is trying to explain. I want now to discuss and answer some of the criticisms levelled at Cummins-functions, for this will clarify the nature of the causal role account of functions and pave the way for the discussion of the case study of the function of language that fleshes out my analysis of functions. I show that when one considers two sciences investigating the same system (as is the case with the investigation of language by biolinguistics and evolutionary biology) it is evident that our functional attributions are only stable within a particular explanatory theory, for what we take the function of something to be (and what justifies functional attributions) shifts as we move to a different explanatory theory.

3.1 Cummins-functions are relative to researcher interests

Neander (2016; 2017) argues that she cannot endorse the causal role understanding of functions because it leaves functions "drifting in the breeze of shifting explanatory aims" (Neander 2017: 1161). She insists that the notion of function must be sensitive to the function-dysfunction or function-accident distinction, and so it must underwrite talk of "normal function, of systems functioning properly, of malfunction, dysfunction, abnormal functioning, impaired functioning and functional deficits" (Neander 2017: 1151). Cummins-functions cannot do this, whereas Neander's normal-proper functions can, for, she argues, "the notion of normal-proper function is […] the notion that most centrally underwrites this kind of talk [of function-dysfunction] in biology" (Neander 2017: 1151). The main problem Neander sees with Cummins-functions is that they are mind-dependent in that "[r]esearcher interests determine the complex […] capacity of a system to which a Cummins function contributes" (Neander 2017: 1153). Moreover, "pragmatic features of the explanatory context

determine […] the boundaries of the system under analysis" (Neander 2017: 1153-1154). In other words, the criticism is that "Cummins functions supervene on the mental states of people seeking explanations" (Neander 2017: 1155). This is a curious way of understanding *mind-dependent*, especially so when seen in comparison to Neander's claim that what she calls minimal functions are not mind-dependent. The minimal function of a biological entity is "just something that it does", minimal functions are thus "mere doings, mere activities, which can contribute to outcomes that we might or might not be interested in explaining" (Neander 2017: 1151). In contrast, Neander argues that on Cummins's account "*if there are no relevant explanatory aims, then there are no functions*" (Neander 2017: 1155, emphasis in original). Presumably, then, Neander's minimal functions are mind-*independent* or objective in some sense, they are essential parts of biological objects, and thus provide a foundation unaffected by the breeze of shifting explanatory aims. Minimal functions, which purportedly do not supervene on the mental states of researchers seeking explanations, can then be used as a foundation for "a stable notion of normal-proper function that does not shift with researcher interests" (Neander 2017: 1160).

There is a problem with minimal functions, however, the cause of which is symptomatic of the way in which selected effects accounts such as Neander's understand functional attributions. The problem is that it is assumed that the "mere activity" of a biological entity, which minimal functions are supposed to describe, can be picked out from all the other things that the entity does in a way that is independent of researcher interests. Suppose one agrees that "the ostensive definition for 'normal-proper function' does not presuppose the identification of normal-proper function with selected function" (Neander 2017: 1153). If this is the case, then there is no way to move from "mere activity" to normal-proper function, for we are blocked from using the theory of natural selection as a way in which to cash out the justification for the functional attribution. We need some explanatory theory in the context of which we can pair a particular "mere activity" with a normal-proper function, for every complex biological entity and most of its component parts do more than one thing.

If no theory or researcher interest is allowed to decide between the various mere activities, then it might seem appealing to argue that biological systems have their functions essentially, perhaps they are natural kinds or objective in some yet to be determined sense. It might be appealing to argue that it is up to our theory to discover, not ascribe, functions. But this is not a tenable position, for the same system can be ascribed different functions by different explanatory theories. In other words, once we ascribe a function to a biological system we need to be able to justify the ascription. If the notion of function is tacitly understood to be proprietary to evolutionary biology then it might seem like theoretical considerations do not play a role in the uncovering and justification of the function. It might appear like the justification for the functional ascription is unrelated to the tacitly assumed theory of natural selection. But as we'll see in the case study of the function of language, when we move away from evolutionary biology to biolinguistics, the ascribed function of the same biological system changes, which shows that functional ascriptions are relativised to the theories in which they take part.

Note that the issue here is not whether mere activities are essential or contingent properties of biological entities but rather whether one of the many activities that a system does can be picked out without reference to a background theory and thus independent of researcher interests. What other criteria can be used to pick out one activity from all the

other activities that a biological system does? The way out of this dilemma is to allow the theory of natural selection to link "mere activity" with normal-proper function, for which Neander of course has made a strong case, but this shows that functional attributions *are* theory-internal, for (to repeat) when we move to a different theory the same system being investigated can be given (and often *is* given) a different functional attribution. I think we should be wary of making the following conflation. An object can have many properties, and one that the theory picks out depends on the theory's explanatory purposes. But this should not be conflated with the claim, that causal role theories do *not* make, that the properties of the object itself depend on the explanatory interests of (or the mind of) the theorist.

In sum, unless we assume that biological entities have their minimal functions essentially, there is no way in which we can objectively pick out a particular activity of a biological entity and label it as a minimal function. Unless every activity that the entity is capable of doing is one of its minimal functions, how are we to know which minimal function can form the basis for a normal-proper function? Biological entities do many things and are engaged in many activities, and so the very act of choosing one mere activity over another is already mind-dependent in Neander's sense. Neander claims that since minimal functions are "mere doings, activities or causal contributions to outcomes", they exist "whether or not these outcomes are complexly achieved and whether or not anyone wants to explain them" (Neander 2017: 1153). That is correct, but the conclusion she draws from this does not follow. No one denies that biological entities do a great variety of things (whether or not we want to label these as minimal functions), but the moment we choose a particular activity or effect to explain, the moment we construct an explanatory theory in order to explain a particular activity or a certain set of activities, our theoretical constructs become mind-dependent. Though notice that it is mind-dependent only in the uninteresting and benign sense that all theoretical constructs in scientific theories are mind-dependent. As Hardcastle (1999: 39) argues, "letting the discipline set the boundaries on relevant effects is pragmatic, but it is neither mysterious nor arbitrary" because "scientists define exactly what they mean by fitness or health" and "these are operationalized to sets of recognizable, rigorous, and robust criteria." These criteria include accepted theories, the research community, the background context of the research, and the biases of individual scientists.

I think that the main source of Neander's discomfort with casual role accounts of functions is that she, like other selected effects theorists, wants there to be an objectively correct function of X that is not relative to explanatory contexts. However, if different theories can ascribe different functions to the *same* biological system, then it follows that functional attributions are inherently theory-internal. In other words, the selected effects account of functions is tailor-made for evolutionary biology, but once we move to a different science it is not a given that an historical notion of function is appropriate. It appears that some critics of Cummins-functions assume that the notion of function as it applies to evolutionary biology is the only sort of function there is. Under this assumption it of course follows that functions cannot be relativised to researcher interests. Matthen (1997), for example, complains that Cummins "makes *every* result of a thing's activities into a function of that thing: for example, it is a Cummins-function of the heart to make diagnostically useful sounds. This is surely a counter-intuitive result" (Matthen 1997: 23, fn. 5, emphasis in original). This is incorrect, for such a criticism misses the relativisation of each Cummins-function to a particular explanatory context. In other words, it is only counterintuitive that

the function of the heart is to make diagnostically useful sounds because our intuition in this case implicitly assumes the explanatory context of evolutionary biology (Hardcastle 2002). But when we switch explanatory contexts, and when we are explicit about what our explanatory context is, then there is no problem with claiming that the heart's function is to make diagnostically useful sounds.

4 The function of language

As I argued above, talk of functions is only appropriate relative to a certain explanatory context. Systems and their constituent parts have many effects, but the effect that counts as the system's function does so because of the explanatory role it plays in a theory. This also applies to the function of language. The language faculty has many effects, one of which is that of allowing us to communicate, but we cannot select this effect as the function of language without answering the prior question of what explanatory role, if any, this functional attribution plays within a particular theory. In other words, how does the claim that the function of language is communication fit in with a particular explanatory theory of language? When this question is answered we realise that the function of language is communication if our theory is seeking to explain the reason for why the language faculty remained in the species, but *at the same time* the function of language is also that of an instrument of thought if our theory is trying to unearth the mechanisms in virtue of which language use is made possible. Functional attributions are theory-internal in this sense. So we can ask two questions simultaneously: (1) Why did the language faculty remain in the species? Answer: Because it had a selectional advantage. Evolutionary biologists (and evolutionary psychologists) argue that it was the communicative function of language that gave it such a selectional advantage. At the same time, however, we can also ask a different question: (2) What is the structure of the language faculty in virtue of which language production and comprehension is made possible? This is a very different question that calls for a different functional attribution. That is, considering the way in which the structure of the language faculty is influenced by its internal linkages within our cognitive system suggests that the function of language is that of an instrument of thought. Let us look at both these questions in turn. I will show that both attributions of function are equally legitimate within their respective explanatory context. If this is the case, then it follows that biological objects do not have their functions essentially, for it is not mind or theory independent factors (in Neander's sense discussed above) that determine the legitimacy of functional attributions but rather each theory's explanatory context.

4.1 Communication as the function of language

Much of the theoretical and empirical work into language has taken it for granted that the function of language is communication. Indeed, just as in the case of the function of the heart, many find it counterintuitive that the function of language could be anything but communication. This is the case in linguistics, philosophy, psychology, and cognitive science (see Asoulin 2016 for references and more discussion). Often this is the starting point of the discussion and an implicit working assumption. Jackendoff writes in this vain that "the basic function of language is to convert thoughts into communicable form; the virtue of human language over other natural communication systems is that the range of messages it can convey is so broad" (Jackendoff 2007: 69). Millikan concurs when she argues that it "is primarily for the service of coordination between speakers and hearers that language

patterns are selected to be proliferated as conventions" and thus "were it not for the fact that employing its conventions sometimes serves purposes common to both speaker and hearer, language as we know it would shrivel and die" (Millikan 2003: 229). In other words, "were it not for their roles in the achievement of communicative coordinations, there is every reason to suppose that the individual language faculties of individual humans would atrophy" (Millikan 2003: 229).

The main argument that attempts to ground the claim that language is for communication claims that the adaptive value of language use is its communicative function. That is, language fitness is said to correspond to communicative success. Millikan's version of this argument is, I think, the most interesting and in-depth investigation of this claim (see also Pinker & Bloom 1990; Jackendoff 1999; Pinker & Jackendoff 2005). Millikan sees linguistic phenomena such as words, surface syntactic forms, tonal inflections, stress patterns, "and any other significant surface elements that a natural spoken or written language may contain" as being what she terms language devices (Millikan 1984: 3). For example, she argues that the original function of the indicative mood was to produce a coordination between speaker and hearer, the function of which was for the hearer to become informed about a proposition and thus believe its content. This pattern became conventional because it was reproduced and thus the syntactic form of the indicative mood proliferated.

The explanatory context of Millikan's discussion is thus the adaptive value of language use. She writes that the "parallels with the evolution and fixation of symbiotic relations between animal species and with the evolution of animal signal systems should be apparent" (Millikan 2003: 231). Language devices in general, says Millikan, should be understood as undergoing evolution by natural selection in effectively the same way as biological systems do, the only difference being that "[u]nlike the lineages that make up animal species, linguistic lineages frequently acquire new functions without changing their physical forms" (Millikan 2005: 61). She says that, similar "to mutations in biological evolution are novel uses of conventional linguistic forms introduced by speakers through figures of speech or through Gricean implicature" (Millikan 2005: 61). So if the hearer interprets an implicature in a new way, this novel use of the language device will result in a new coordinating function. This new function will then be copied by other speakers, the result of which will be that a "new lineage of tokens with a different stabilizing function has branched off from the original lineage but without any change in physical form" (Millikan 2005: 61).

4.2 The function of language as an instrument of thought

The function of language from the evolutionary biologist perspective is justified by arguing that since language was used for communication, and since communication was fitness enhancing, language remained in the species. If this particular explanatory context remains implicit or unnoticed, then one can perhaps agree with Millikan (2005: 25) that "a primary function of the human language faculty is to support linguistic conventions, and that these have an essentially communicative function." But what about the mechanisms in virtue of which language use is made possible? An account of these mechanisms is indispensable for a full understanding of the language faculty and the way it is used. But when we change our focus to the mechanisms in virtue of which language use is made possible, when we consider the structure of the language faculty itself, we change the explanatory context and

as a result the function of language being communication no longer provides a basis for a fecund explanation. In fact, the claim that the language faculty remained in the species because its use in communication bestowed upon it an adaptive fitness tells us little about the *structure* of the language faculty. This is because it is near impossible to derive the properties of the underlying mechanisms of language from (selected effects) functional accounts of language use, for communication systems are consistent with more than one sort of language faculty (see Reinhart 2006 for discussion). That is, fulfilling the role of a communication system is compatible with more than one way in which the language faculty could be structured; it's also compatible with no language faculty at all, for non-human animals communicate but have no language faculty. In other words, communication systems can be instantiated in various different ways, only one of which includes the human faculty of language.

Notice that the claim here is not that the language faculty is multiply realisable, but rather that one effect of the use of the language faculty (communication) is multiply realisable, for the language faculty is not necessary for communication. Analogously, it is near impossible to derive the properties of the underlying mechanisms of flight from selected effects accounts of wings, for flight is consistent with more than one sort of physiology (insect wings, bird wings, etc.). Thus, if we want to study how the language faculty works (analogously, if we want to study how the insect wing works), then we need to move away from the evolutionary biologist perspective. Another way to put the matter is as follows. Explaining the presence of the language faculty by appeal to what it does (allowing communication between humans) does not explain why a functional equivalent of the language faculty with a different structure isn't there instead. There is no deductive inference from communication to the language faculty. What we have here instead is an inference to the best explanation (Cummins 1975). That is, given that communication is occurring in humans, we can legitimately infer that the language faculty is present in humans precisely because the language faculty enters into our best explanation of communication in humans. But if we are investigating the structure of the current human language faculty, as opposed to investigating why it remained in the species, then the selected effects account of functions is unhelpful.

A different perspective on the function of language is that of generative linguistics, and also the broader, generative-oriented, biolinguistics program. Biolinguistics treats language as an internal computational system, a recursive mechanism that produces a potentially infinite set of hierarchically structured expressions that are employed by the conceptual-intentional systems (systems of thought) and the sensorimotor systems to yield language production and comprehension. This view of language is strongly shaped by its interface with the systems of thought, rather than by the peripheral process of externalisation inherent in the link with the sensorimotor systems (Chomsky 2007; Boeckx 2011; Di Sciullo & Boeckx 2011). In order to understand why the language faculty has the structure that it does, it is necessary to understand how it is integrated into and how it is linked with other cognitive systems in the mind. In this explanatory context, the functional attribution that is best suited is that of language as an instrument of thought (Chomsky 2013; 2013a; Asoulin 2016).

Biolinguistics does not deny that *a* function of language is communication, it does not deny that we use language to communicate. Rather, the claim is that from the perspective of its internal structure the function of the language faculty is that of an

instrument of thought. This may sound counterintuitive or fallacious, but that is only because the usual implicit explanatory context within which we consider functional attributions is one that focusses on the appearance and subsequent persistence in the species of organs such as the language faculty. However, there is another way in which we can study language, namely, by regarding language growth as analogous to the growth of bodily organs (Chomsky 1995; Anderson & Lightfoot 2000; 2002). From this perspective the language faculty is seen as a biological organ in the same sense that the immune system, the skin, or the circulatory system are seen as organs of the body. That is, these organs are "not objects that can be removed leaving the rest intact, but subsystems of a more complex structure that we hope to understand by investigating parts that have distinctive characteristics, and their interactions" (Chomsky 2000: 90). Biolinguistics takes its object of study to be the underlying mechanisms of language, which are a subsystem of our cognitive system and are composed of a computational system (called an I-language) that is encoded in individual brains.

The claim that language is an instrument of thought amounts to much more than the claim that the function of language is to take pre-formed thoughts and then externalise them. That is, the claim regards the way in which the underlying mechanisms of language structure thought in a particular way that is unique to humans. More specifically, the language faculty and its computational operations allow humans to have a distinct type of thought that we do not share with other animals *in addition* to the types of thought that we do share with them (Hinzen 2006; 2013). There are two clarificatory remarks to be made here. First, note the stress on the underlying mechanisms of language and not on any particular natural language – this distinction mirrors the different perspectives on language that are taken by evolutionary biologists and biolinguists. That is, evolutionary biologists focus on the use of the language faculty (in conjunction with other systems) in the context of particular natural languages such as English or Italian, whereas biolinguists focus on the underlying mechanisms in virtue of which this use is made possible. Each explanatory context requires its own functional attribution, so that in the context of biolinguistics the fundamental property of language is the internal construction of infinitely many expressions by a generative procedure that yields uniquely human thought processes and perspectives on the world. According to this view, then, the externalisation of language with intent to communicate is a secondary property of language. The underlying mechanisms of language provide us with a unique way of structuring the world around us, which we use for various purposes such as thinking and talking about the world. Second, this claim is not a Whorfian one of linguistic determinism, thought is certainly independent of language, and what can be expressed or thought by a speaker of one language can certainly be expressed or thought by a speaker of a very different language.

Thus, as opposed to the explanatory context that attributes to language the function of communication, the best functional attribution that fits in with the explanatory aims of biolinguistics is that of language as an instrument of thought. That is, if we want to explain how the computational system that is the language faculty fits in with the human cognitive system then language as an instrument of thought is the more fecund functional attribution.

5 Concluding remarks

Biolinguistics and evolutionary biology are different theories with different explanatory aims in regard to language. The latter focusses on the appearance and subsequent

persistence in the species of the language faculty, whereas the former sees language as an instrument of thought. The result is that they each attribute a different function to the *same* entity, namely to the language faculty. The same structure is given a different functional attribution depending on the explanatory theory within which it is embedded. This stands in contrast to the claims of Neander and others according to which there must be a truth of the matter as to what function a particular system has. I argued above that we cannot attribute functions independently of our theoretical considerations. Therefore, for example, there is no theory-neural sense in which we can claim, as Wright does, that the heart throb is a paradigm case of non-function. Similarly, there is no theory-neutral sense in which we can claim, as Millikan does, that the primary function of the language faculty is communication. The question is not whether a particular theory is correct in attributing a particular function to X or not. Rather, the question is what explanatory work is a functional attribution doing in a particular theory. The functional attribution that best fits the explanatory purposes of biolinguistics is that of language being an instrument of thought, whereas the functional attribution that best fits the explanatory purposes of evolutionary biology is that of communication being the function of language. Both are legitimate within each theory and both theories add to our understanding of language and the role it has in our cognition and the way in which we interact with the world. Thus, functions are not essential parts of biological objects but are theory-internal in the sense that the fruitfulness of particular functional ascriptions is judged on how well the ascriptions contribute to the aims of the explanatory theory, to what the theory is trying to explain.

References

Amundson, Ron & George V. Lauder. 1994. Function without purpose: The uses of causal role function in evolutionary biology. *Biology and Philosophy* 9. 443-469.

Anderson, Stephen R. & David Lightfoot. 2000. The human language faculty as an organ. *Annual Review of Physiology* 62. 697-722.

Anderson, Stephen R. & David Lightfoot. 2002. *The language organ: Linguistics as cognitive physiology*. Cambridge: Cambridge University Press.

Asoulin, Eran. 2016. Language as an instrument of thought. *Glossa: a journal of general linguistics* 1(1): 46. 1–23.

Bechtel, William. 2008. *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. New York: Routledge.

Bechtel, William. 2009. Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology*, 22(5), 543-564.

Bechtel, William & Adele Abrahamsen. 2005. Explanation: A mechanist alternative. *Studies in History and Philosophy of Biology & Biomedical Sciences*, 36(2). 421-441.

Boeckx, Cedric. 2011. Biolinguistics: A brief guide for the perplexed. *Linguistic Sciences* 10(5). 449–463.

Chomsky, Noam. 1995. Language and nature. *Mind* 104(413). 1–61.

Chomsky, Noam. 2000. Linguistics and brain science. In Alec P. Marantz, Yasushi Miyashita & Wayne O'Neil (eds.), *Image, language, brain*, 13-28. Cambridge, MA: MIT Press.

Chomsky, Noam. 2007. Biolinguistic explorations: Design, development, evolution. *International Journal of Philosophical Studies* 15(1). 1–21.

Chomsky, Noam. 2013. What kind of creatures are we? *The Journal of Philosophy* 90(12). 645–700.

Chomsky, Noam. 2013a. Problems of projection. *Lingua* 130. 33-49.

Craver, Carl F. 2001. Role functions, mechanisms, and hierarchy. *Philosophy of Science* 68(1). 53-74.

Craver, Carl F. & Lindley Darden. 2013. *In search of mechanisms: Discoveries across the life sciences*. Chicago: The University of Chicago Press.

Cummins, Robert. 1975. Functional analysis. *The Journal of Philosophy* 72(20). 741–765.

Cummins, Robert. 2000. "How does it work?" versus "What are the laws?": Two conceptions of psychological explanation. In Frank C. Keil & Robert A. Wilson (eds.), *Explanation and cognition*, 117–145. Cambridge, MA: MIT Press.

Cummins, Robert & Martin Roth. 2009. Traits have not evolved to function the way they do because of a past advantage. In Francisco J. Ayala & Robert Arp (eds.), *Contemporary debates in philosophy of biology*, 72–85. Oxford: Wiley-Blackwell.

Davies, Paul Sheldon. 2001. *Norms of nature: Naturalism and the nature of functions.* Cambridge, MA: MIT Press.

Di Sciullo, Anna Maria & Cedric Boeckx (eds.). 2011. *The biolinguistic enterprise: New perspectives on the evolution and nature of the human language faculty*. Oxford: Oxford University Press.

Frankfurt, Harry G. & Brian Poole. 1966. Functional analyses in biology. *The British Journal for the Philosophy of Science* 17(1). 69-72.

Garson, Justin. 2011. Selected effects and causal role functions in the brain: The case for an etiological approach to neuroscience. *Biology & Philosophy* 26(4). 547–565.

Garson, Justin. 2016. *A Critical Overview of Biological Functions.* Dordrecht: Springer.

Godfrey-Smith, Peter. 1993. Functions: Consensus without unity. *Pacific Philosophical Quarterly* 74. 196-208.
Godfrey-Smith, Peter. 1994. A modern history theory of functions. *Noûs* 28(3). 344-362.
Godfrey-Smith, Peter. 1996. *Complexity and the function of mind in nature.* Cambridge: Cambridge University Press.
Hardcastle, Valerie Gray. 1999. Understanding functions: A pragmatic approach. In Valerie Gray Hardcastle (ed.), *Where biology meets psychology: Philosophical essays*, 27-44. Cambridge, MA: MIT Press.
Hardcastle, Valerie Gray. 2002. On the normatively of functions. In André Ariew, Robert Cummins & Mark Perlman (eds.), *Functions: New essays in the philosophy of psychology and biology*, 144-156. Oxford: Oxford University Press.
Hinzen, Wolfram. 2006. *Mind design and minimal syntax*. Oxford: Oxford University Press.
Hinzen, Wolfram. 2013. Narrow syntax and the language of thought. *Philosophical Psychology* 26(1). 1–23.
Jackendoff, Ray. 1999. Possible stages in the evolution of the language capacity. *Trends in Cognitive Sciences* 3(7). 272–279.
Jackendoff, Ray. 2007. *Language, consciousness, culture: Essays on mental structure.* Cambridge, MA: MIT Press.
Kaplan, Jonathan M. & Massimo Pigliucci. 2001. Genes `for' phenotypes: A modern history view. *Biology & Philosophy* 16. 189-213.
Lewens, Tim. 2004. *Organisms and artifacts: Design in nature and elsewhere.* Cambridge, MA: MIT Press.
Lewens, Tim. 2007. Functions. In Mohan Matthen & Christopher Stephens (eds.), *Philosophy of biology*, 525-547. Amsterdam: Elsevier.
Machamer, Peter, Lindley Darden & Carl F. Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67(1). 1-25.
Matthen, Mohan. 1997. Teleology and the product analogy. *Australasian Journal of Philosophy* 75(1). 21-37.
McLaughlin, Peter. 2001. *What functions explain: Functional explanation and self-reproducing systems*. Cambridge: Cambridge University Press.
Millikan, Ruth Garrett. 1984. *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
Millikan, Ruth Garrett. 1989. In defense of proper functions. *Philosophy of Science* 56. 288–302.
Millikan, Ruth Garrett. 1989a. An ambiguity in the notion "function". *Biology and Philosophy* 4. 172-176.
Millikan, Ruth Garrett. 2003. In defense of public language. In Louise Antony & Norbert Hornstein (eds.), *Chomsky and his critics*, 215-237. Oxford: Blackwell.
Millikan, Ruth Garrett. 2005. *Language: A biological model*. Oxford: Oxford University Press.
Neander, Karen. 1991. The teleological notion of "function". *Australasian Journal of Philosophy* 69(4). 454-468.
Neander, Karen. 1991a. Functions as selected effects: The conceptual analyst's defense. *Philosophy of Science* 58(2). 168-184.
Neander, Karen. 2016. Kitcher's two design stances. In Mark Couch & Jessica Pfeifer (eds.), *The philosophy of Philip Kitcher*, 45-67. Oxford: Oxford University Press.
Neander, Karen. 2017. Functional analysis and the species design. *Synthese* 194. 1147-1168.

Pinker, Steven & Paul Bloom. 1990. Natural language and natural selection. *Behavioral and Brain Science* 13. 707–784.

Pinker, Steven & Ray Jackendoff. 2005. The faculty of language: What's special about it? *Cognition* 95. 201–236.

Reinhart, Tanya. 2006. *Interface strategies: Optimal and costly computations.* Cambridge, MA: MIT Press.

Roth, Martin & Robert Cummins. 2014. Two tales of functional explanation. *Philosophical Psychology* 27(6). 773-788.

Roth, Martin & Robert Cummins. 2017. Neuroscience, psychology, reduction and functional analysis. In David M. Kaplan (ed.), *Explanation and integration in mind and brain science*, 29-43. Oxford: Oxford University Press.

Roux, Etienne. 2014. The concept of function in modern physiology. *The Journal of Physiology* 592(11). 2245–2249.

Thagard, Paul. 2012. *The cognitive science of science: Explanation, discovery, and conceptual change*. Cambridge, MA: MIT Press.

Wright, Larry. 1973. Functions. *The Philosophical Review* 82(2). 139-168.

Wright, Larry. 1976. *Teleological explanations: An etiological analysis of goals and functions*. Berkeley: University of California Press.