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

Eran Asoulin
Macquarie University, NSW, Australia

NOVEMBER 2017 DRAFT.
COMMENTS WELCOME.

Abstract: I argue that talk of functions in the biological realm is a theory-internal matter, and so nothing has its function essentially. Functions are ascribed to suit the explanatory purposes of particular theories. This is in contrast to functions of human-made objects, where functions are defined by reference to the intentions of the designer. This is not to say that there are no criteria by which to judge what a function is: rather, I want to suggest that the criteria of functional ascription are almost entirely theory-internal. A revealing example of the theory-internal nature of functional ascriptions is the function of language. 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, say, evolutionary biology and generative linguistics to have different conceptions of what the function of language is.

Keywords: functions, biological functions, function of language, scientific explanations, teleological explanations.

1 Introduction

The aim of this paper is to 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 and expository tools. This is in contrast to functions of human-made objects, where functions are defined by reference to the intentions of the designer. This is not to say that there are no criteria by which to judge what a function is in the biological realm: rather, I want to suggest that the criteria of functional ascription are almost entirely theory-internal. 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. The function of the heart as a pump is not an activity that we can characterise separate from our aims to explain the emergence and persistence of the heart in the species. A revealing example of the theory-internal nature of functional ascriptions is the debate in regard to the function of language. 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, say, evolutionary biology and generative linguistics to have different conceptions of what the function of language is.

2 What is a function?

There is an intuitive distinction between, say, the operation of a machine having certain effects and one of these effects being the machine's function. This is because the way in

which we ascribe functions to artefacts is inherently teleological – we know what the machine is *for* because of the intentions of its designer. Everyone can agree that the operation of a refrigerator has many effects (its motor makes a noise, say) but that only one (or very few) of these effects is its function. However, when one moves away from human-made objects, which were constructed with a certain purpose in mind, and into the biological realm this pre-theoretical intuition becomes problematic: it is not clear whether and to what extent this notion of function applies to biological entities.

Broadly speaking, there are three ways in which to address this problem, each of which will be addressed in the sections that follow. The first, which as Lewens (2007: 526) notes and as Neander (1991a: 168) predicted is now the majority position in the philosophical literature on the topic, is to claim that function talk is justified because it is true of biological systems. The second is to claim that function talk is not true of biological systems, but that it is a useful heuristic or metaphor that we use as aids in theory construction. The third way, which is the view I argue for here, is to claim 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 approaches to functional explanation that fall under the first two ways are teleological. It is a basic tenet of these approaches that the presence of, say, a biological organ can be explained by appeal to its function. The third way, however, is neutral in regard to whether functional explanations are explanatorily fruitful. This is because, as I argue below, 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, then a particular functional ascription will be 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 will be required. I argue below that this is true in the case of the function of language. 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, some perhaps conflicting with others, 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. They are all versions of what might be called a naturalised teleology that attempts to link functional ascriptions to naturalistic inquiry (see Walsh 2008; 2012). In work spanning several decades, Millikan has argued that function qua

purpose can be given a naturalistic reading. Human purposes, she says, when considered from a naturalistic perspective, "turn out to correspond to proper functions of a certain kind, so that a connection between human purposes and the natural purposes of body organs and instinctive behaviors will be established that does not rest on mere metaphor" (Millikan 1984: 4). I detail below her notion of proper function as it pertains to the function of language and to what she calls language devices, but for present purposes a proper function is for her defined in terms of the history of the system under study, and this is contrasted with the system's current structure or dispositions. In other words, her claim is that it is proper functions that define certain biological categories, so that the way in which we decide whether something is part of the biological category of, say, *kidney*, is by looking at its history of adaptation and selection.

A classic and oft-quoted exposition of the notion of function is Wright (1973), who argues that "conscious and natural functions are functions in the same sense, despite their obvious differences" (Wright 1973: 143). It is worth looking at the debate that Wright's paper initiated, for it sheds light on the details and motivations for arguing that functions are true of biological systems. The task Wright 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, 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."

Before continuing, note what sort of explanation this is: it is the sort of explanation that asks why a particular thing, say an organ, is present in an organism or system. Questions of the sort that ask why humans have hearts is claimed to be equivalent to questions of the sort that ask why car internal combustion engines have carburettors. Just as the answer to the latter is given by giving the function of the carburettor (it is there in order to blend air and fuel vapour), the answer to the former is given by giving the function of the heart (it is there in order to pump blood). This is a very different notion of scientific explanation to what Thagard (2012) calls the *mechanista* view of scientific explanation, according to which to explain a phenomenon is to describe a mechanism that produces it – in other words, to explain how it works. I think these two types of explanation map onto Mayr's (1961) distinction between the why-questions that form the basis of evolutionary biology and the how-questions that form the basis for what he calls functional biology. More on this below.

So the analysis of function that Wright is proposing can be summarised as follows: to say that the function of X is Z (i.e., that the function of X is to do Z) is to say that X is there because it does Z, and that this sort of explanation is an explanation of why X is there. This is of course a teleological explanation. 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.

As Lewens (2004) correctly notes, most contributions to the functions debate in the two or three decades following Wright (1973; 1976) were refinements and elaborations of Wright's etiological analysis.[1] One of the better known defences of the etiological theory is Neander (1991), 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 the notion of function can be both teleological and scientifically respectable, for to investigate the function of a trait is to investigate the trait's history of selection. If this is the case then "the apparent explanatory power of teleological explanations which appeal to biological functions is quite robust" (Neander 1991: 457). 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 by natural selection 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 2015: 16, emphasis in original).

In addition, like Wright, Neander insists that the functions of artefacts and the functions of biological entities are 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

[1] Notable earlier analyses of function, not necessarily similar to Wright's, include Hempel (1959), Nagel (1961), Sorabji (1964), Lehman (1965), Ayala (1970), Ruse (1971). More recent analyses that take a historical or etiological approach include Neander (1991; 1991a; 2015), Griffiths (1993), Godfrey-Smith (1993; 1994), Kaplan & Pigliucci (2001), Pigliucci & Kaplan (2006).

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., 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 (both in the biological and artefact cases), as I think they do, then the best way to do so is arguably to find a way to give biological 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 clearly 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 2015: 15). 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 one's 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 function) and it is thus possible for a certain theory to attribute the wrong function to a biological object. That is, it follows from this view 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 Functions are heuristics

In contrast to the above, where functions are seen as mind-independent and true of biological systems, there is the claim that function talk is merely a useful heuristic or metaphor that aids in theory construction and theoretical explanation. Dennett (1987) refers to this understanding of functions as artefact hermeneutics, and he endorses it as part of his intentional stance. Treating evolution from the intentional stance implies that functional attributions to humans "cannot be sustained […] without appeal to assumptions about 'what Mother Nature had in mind'" (Dennett 1987: 314). But we must realise, Dennett

argues, that these are nothing but assumptions: disagreements about correct functional attributions cannot be settled by appealing to any underlying facts, "*for there are no such deeper facts"* because "when we go to read Mother Nature's mind, there is no text to be interpreted" (Dennett 1987: 300, emphasis in original). That is, when a particular attribution of "proper function is controversial – when more than one interpretation is well supported – there is no fact of the matter" (Dennett 1987: 300). In other words, we must "abandon the idea that there has to be one, determinate, *right* answer to the question: What is it for?" (Dennett 1987: 319, emphasis in original). Dennett argues strongly for what he calls functional indeterminacy: "Mother Nature doesn't commit herself explicitly and objectively to *any* functional attributions; all such attributions depend on the mind-set of the intentional stance" (Dennett 1987: 320, emphasis in original).

In order to flesh out his reasoning for making this claim, Dennett gives a thought experiment, which is a variation of Dawkins's (1976) view of organisms being survival machines with a design that is intended to prolong the future existence of the "selfish genes" that make up the organisms. Dennett imagines a robot that is designed to prolong the life of a comatose person that is placed in a capsule inside the robot. The comatose person of course cannot guide the robot and instruct it when the need arises to make decisions about its survival, and so the robot must be designed at the outset to, amongst others, plan its future actions, learn from its mistakes, and communicate with its environment. Since the robot is a designed artefact, Dennett argues that there is no sense in which any of its states have anything other than derived intentionality – that is, the content of all of its states have no meaning outside of what the designer meant them to have. Dennett argues that if this is the case, and if Dawkins's (1976) view of organisms being survival machines is correct, then the conclusion is forced upon us that organisms such as humans are also artefacts in the same sense, the only difference being that the designer in the case of humans is evolution: "We are artifacts, in effect, designed over the eons as survival machines for genes that cannot act swiftly and informedly in their own interests" (Dennett 1987: 298). In other words, Dennett's argument is that since we ascribe functions to artefacts, and since humans are artefacts in the sense just described, we ascribe functions to humans in the same way as we ascribe functions to artefacts.

Matthen (1997) is similarly skeptical of functional attribution (for other sorts of skepticism see Searle 1995 and Lewontin 1983; 2000), and he perhaps goes further than Dennett by arguing that explanation and prediction cannot be sustained by functional attribution at all. Matthen argues that teleological explanations exploit a similarity between designed objects and biological entities, but that this similarity rests merely "on the claim that the structure of relevance relations in the target phenomenon [in the biological world] remind us of that found in the products of intelligence and design" (Matthen 1997: 22). So to recognise such a similarity, which Matthen calls the product analogy, is not the same as to say "that natural selection works in anything like the same way as design", for the product analogy "is compatible with radical dissimilarities both between the products of selection and the products of design and among the processes by which they come to be" (Matthen 1997: 22). The consequence of this reasoning, Matthen argues, is that there is no such thing as teleological explanation, for to "indicate the explanatory relevance of certain factors is not the same as producing an explanation in which these factors figure" (Matthen 1997: 23). The claim here is that biologists make use of teleological explanations in cases where what a thing does is found to be relevant to explaining its existence. Such explanatory relevance

relations are exemplified by many artefacts, where of course each is made with a design and a view as to what it will do, and thus the analogy is made to biological entities. But this relevance relation is as far as we can take the product analogy, Matthen argues. It does not extend, he says, "to the way in which the functional feature came to be present, nor to the structure of the thing" (Matthen 1997: 35).

Matthen argues that teleological explanations merely point us in the direction of certain explanatory relevance relations that biologists make use of in their theory construction and explanations. In other words, functional attributions are "the attempt to show that the use to which a morphological structure is put, or the outcome of a biological process, has explanatory relevance to explaining the origins of the structure or process, not any particular hypothesis about how this could be so" (Matthen 1997: 36). Matthen argues that there can be any number of explanatory theories that are able to underwrite the explanatory relevance relations, some of which include Paley's design hypothesis, Lamarck's theory of evolution, and Darwin's theory of natural selection. It is these causal hypotheses, says Matthen, and not the functional attributions, that are doing the explanatory work. Thus, on this view, teleological explanations are not explanations per se, for they act only as a heuristic to guide the biologist in the quest to find explanations. Matthen goes further still by arguing that teleological attributions are dispensable, for "they describe a natural phenomenon by means of an analogy, and there is no reason why a more direct form of description should not be used in their place" (Matthen 1997: 37).

2.3 Cummins-functions

The third way of explaining the way in which functions can be ascribed to biological entities is to claim that function talk is justified not because it is true of biological systems but because it best fits the explanatory purposes of our theory. Despite some similarities, this way of thinking of functions is different to the functions qua heuristics view. This third way rejects functional indeterminism and argues that there *is* a fact of the matter about what function a particular biological entity has. So unlike Dennett, who claims that there is no fact of the matter to appeal to when we are faced with more than one functional attribution to the same biological entity, this third way claims that we can underwrite functional attributions by appealing to theory-internal explanatory contexts. One way in which to do this is the systematic account of function attribution (Cummins 1975; Cummins & Roth 2009). This account utilises functional analysis as its explanatory strategy, where the operation of systems is explained by the operation of their constituent parts. So complex systems are explained in terms of their (usually simpler) constituent parts as, for example, "amplification gets analyzed into the capacities of resistors, conductors, capacitors, power supplies, etc." (Cummins & Roth 2009: 74). The same is true in biology where 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 is to single out a capacity of the system by finding the role that that capacity plays in the analysis of the system. 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). So, unlike the second view, there *is* a fact of the matter as to what function a particular organ has, but this fact depends on the analytical context that relativises function-ascribing statements to the background of analysis. As I argue below, this relativisation is clear in the case of the function of language, where the background of analysis dictates what function is ascribed.

Notice that this third way takes the nature of scientific explanations, and the role of functions therein, to be something quite different from the first (etiological) way discussed earlier. The first way 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 third way, 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. One of the main roles of functions, on this view, is as causal relevance filters. That is, functional analyses specify the design of the system under investigation, thus providing an explanation of the system's capacities, but since any system under investigation will exhibit a multitude of effects we must have some way of filtering out the relevant effects from the irrelevant ones. As Roth & Cummins (2014) argue, functional concepts provide us with a way to isolate the causal consequences that are relevant to the capacities of the system we are investigating. In contrast to the way that Wright and Neander understand functions, this third way sees functions as nothing more than causal relevance filters, for "the explanatory power they provide has nothing to do with the subsumption of cause-effect pairs under functional laws" (Roth & Cummins 2014: 784; see also Cummins 2002; Cummins & Roth 2009; Roth & Cummins 2017).

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 explanatory theories explain phenomena by unearthing the mechanisms that produce them (see Thagard 2012). On this view one of the major aims of science is the discovery of mechanisms rather than laws (see Craver 2001 for discussion). 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.

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 indeed 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. Moreover, it is striking that despite their apparent prevalence, 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 etiological understanding of function 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 argue below in the case of the function of language, if we want to know how an organ works then the mechanistic understanding of functions is best. However, if we want to know why a particular organ has remained in the species with the structure it has, then the etiological understanding of function is best.

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 & 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 a function is in a particular explanatory context. Whilst the details of how to decide this are beyond the scope of this paper, I think the general outline of an account is clear in what follows, namely that what function is attributed depends on the 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 Cummins-functions and pave the way for my own analysis of functions, which is similar to Cummins's account in a number of respects. Then in the following section I discuss 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 by linguistics and developmental biology of language – 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 shifts as we move to a different explanatory theory.

3.1 Cummins-functions are relative to researcher interests

Neander (2015) is critical of Cummins-functions and argues that she cannot endorse such an understanding of functions because it leaves functions "drifting in the breeze of shifting explanatory aims" (Neander 2015: 15). 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 2015: 5). 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 2015: 5). 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 2015: 7). Moreover, "pragmatic features of the explanatory context determine […] the boundaries of the system under analysis" (Neander 2015: 7-8). In other words, the criticism is that "Cummins functions supervene on the mental states of people seeking explanations" (Neander 2015: 9). 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 2015: 5). In contrast, Neander argues that on Cummins's account "*if there are no relevant explanatory aims, then there are no functions*" (Neander 2015: 9, emphasis in original). Presumably, then, 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 2015: 14).

There is a problem with minimal functions, however, the cause of which is symptomatic of the way in which etiological accounts such as Neander's understand functional attributions. The problem is that it is assumed that the "mere activity" of a biological entity (that minimal functions are supposed to describe) exists independently 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 2015: 7). 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. If no theory or researcher interest is allowed to decide the matter, then the only other option is to argue that biological systems have their functions essentially, perhaps they are natural kinds or objective in some yet to be determined sense; but this is not a tenable position. 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 when we move to a different theory or science the same system being investigated can be given (and *is* given) a different functional attribution.

So 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 2015: 7).[2] 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 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".

I think that the main source of Neander's discomfort with Cummins-functions is that she, like other etiological theorists, wants there to be an objectively correct function of X that is not relative to explanatory contexts. However, functional attributions are inherently theory-internal, for different theories can ascribe different functions to the *same* biological system. In other words, the etiological theory and its notion of function is tailor-made for evolutionary biology, but once we move to functional biology (or chemistry, biolinguistics,

[2] Just to be clear, this is what I am referring to when I discuss the claim that biological objects have their functions essentially. In other words, the claim is that functions are among the set of properties that biological objects have that defines what they are, and that the theory of natural selection helps us uncover the normal-proper function of particular biological objects. Note that the claim that biological objects have their functions essentially implies that this is the case not only within particular theories but *across theories* too.

or any other science), it is not a given that an historical notion of function is appropriate.[3] 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 developmental biology. 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. Let us see how this is the case with the function of language.

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 attribution of function. That is, considering the way in which the structure of the language faculty is influenced by its internal linkages with our cognitive system suggests

[3] Though I don't have the space to discuss it here, I think the matter is even more complex than this. The selected effects theory (or any other theory of function) cannot work in the abstract, for it is an empirical question whether a particular trait has been selected for something or whether it has arisen by non-selectional mechanisms (perhaps the trait arose due to internal-to-the-organism developmental pathways that are largely unaffected by the environment). There is no question that selected effects play a role in evolution, but that does not tell us much until we specify (*and provide supporting empirical evidence*) that *this particular trait* exists because it was selected for something and not because, say, it arose due to developmental noise, mutation, or purely internal developmental factors. For discussion see Asoulin, Eran. 2017. *Functions, selected effects and empirical reality.* Unpublished manuscript, Macquarie University.

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 valid 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 validity 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. 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). Ellis argues that "any definition of language must include a communicative function", for "it is impossible to define communication without reference to a linguistic component" (Ellis 1999: 2). In other words, language serves communication, for language "is only useful or practical to the extent that it ministers to communicative goals" (Ellis 1999: 2). 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.[4] Millikan sees linguistic phenomena such as words, surface syntactic forms, tonal inflections, stress patterns, punctuations, the imperative mood, and the indicative mood, amongst others, as all being what she terms language devices. Note that, as diverse as these linguistic features are, they are all features of language use as opposed to the underlying mechanisms in virtue of which this use is made possible; more on this below. Millikan explicitly defines language devices as including the ones listed above "and any other significant surface elements that a natural spoken or written language may contain" (Millikan 1984: 3). In what follows I consider Millikan's analysis of the functions of the grammatical moods in order to bring to light the explanatory context within which operates the claim that the function of language is communication. In the next section I will show that a change in explanatory context brings with it a change in the function ascribed to language.

Millikan argues that, in English, the syntactic forms identified with the indicative mood "are proliferated in the service of a number of different coordinating conventions

[4] For other versions of the argument see Pinker & Bloom (1990), Jackendoff (1999), Pinker & Jackendoff (2005).

having different functions" (Millikan 2003: 230). These functions, according to Millikan, include conveying information, giving orders, conveying norms, and making declarations. The original function of the indicative, she claims, was to use it for conveying information, and this use followed a certain pattern. This pattern has become conventional because it was reproduced and thus it proliferated. Before looking at the nature of this pattern, note that a reproduction for Millikan "must be such that had the original been different in specifiable respects the reproduction would have differed accordingly" (Millikan 2003: 219). Moreover, all that is required for one thing to be reproduced from another is there to "be a mechanism that produces the second on the model of the first, such that, had the first been different in specifiable respects, that would have caused the second to differ accordingly" (Millikan 2003: 219). Under this way of understanding reproduction, the process can of course occur both consciously and unconsciously.

Millikan writes that the pattern of the use of the indicative mood for conveying information begins with "a speaker S believing some proposition p and accordingly speaking an indicative mood sentence that expresses p, given the truth-conditional part of the semantics of the language", and the pattern concludes when "a hearer H, following the truth-conditional semantics for the language, translates the sentence into the thought that p, and accordingly believes that p" (Millikan 2003: 230). This pattern, then, produces a coordination between speaker and hearer, the function of which is for H to become informed about p. For this pattern to have this function, Millikan argues, it must be a reproduced pattern, it must continue to be reproduced only because it makes a contribution servicing coordinations, and this contribution would make it the case that H would become informed about p. This is the case for a specific pattern, but for "the indicative mood form itself to have the transfer of information from speaker to hearer as a function, it must be that the contribution it makes to the whole is to utilize the functions of other reproduced parts of sentences exemplifying it to that end" (Millikan 2003: 230).

It is clear in the above that the explanatory context of Millikan's discussion is the adaptive value of language use. This is explicit immediately following the above quote, where she argues that the reason H responds to indicative sentences by translating them into beliefs is that "in H's experience, responding selectively to indicative sentences in this way has often enough resulted in the appropriation of useful information" (Millikan 2003: 231). In other words, "had H not lived where speakers often enough expressed true beliefs using the indicative pattern with these truth-conditional semantic rules, H would not translate from indicative sentences into belief in this way" (Millikan 2003: 231). Thus, speakers and hearers collectively learn from each other the way in which to speak and respond so that each of their purposes is served, so that the speaker's function relies on the hearer's function and vice versa. Considering this, Millikan 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).[5] 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,

[5] Such comparisons are common in Millikan's work. See also: "Because linguistic conventions can be more or less compatible in various ways, and because they are often built on one another, exactly like genes in gene pools, then tend to get together in stable clusters" (Millikan 2003: 225).

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).

The explanatory context in which Millikan understands the attribution of a function to language is one in which function is understood as "the hypothesized function of a given language device that accounts for the continued proliferation both of speaker utterances and of stable cooperative (overt or covert) hearer responses" (Millikan 1984: 3). In other words, in this context the question that the functional attribution is supposed to help answer is why the language faculty remained in the human species. This is analogous to the case of the function of the heart. When we are trying to explain why the heart remained in the species then the functional attribution of it being a pump is explanatory. It helps us understand why the heart does what it does when seen from this perspective. But when we switch explanatory contexts, say when we move to medical diagnostics, the function of the heart making thumping noises makes much more sense.

4.2 The function of language as an instrument of thought

As I noted above, in much of the discussion of the function of language there is an implicit assumption according to which the explanatory context is the appearance and subsequent persistence in the species of the language faculty. But this explanatory context is not the only possible context. Thus, consider generative linguistics and evolutionary theory. The latter argues that the language faculty remained in the species due to its selectional advantage in fostering better communication and co-operation. This is part of the explanatory context, part of the main question that evolutionary theory attempts to answer. As Mayr put it in his hugely influential paper almost sixty years ago, the evolutionary biologist has in mind the historical question "How come?", for the field's main preoccupation is "to find the causes for the existing characteristics and particularly adaptations of organisms" (Mayr 1961: 1502). The main aim of the field is to look for the historical background that has led to particular organisms and particular features of these organisms to be the way they are and to persist in the species. That is, the evolutionary biologist "studies the forces that bring about changes in faunas and floras" and "the steps by which have evolved the miraculous adaptations so characteristic of every aspect of the organic world" (Mayr 1961: 1502).

The function of language from the evolutionary biologist explanatory perspective, then, focusses on the way in which language is used, thus fostering the argument that since language was used for communication, and since communication was fitness enhancing, language remained in the species. 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. 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 functional accounts of language use, for communicative systems are consistent with more than one sort of language faculty (see Reinhart 2006 for discussion). In other words, fulfilling the role of a communication system is compatible with more than one way in which the language faculty could be structured. Since this is the case, if we want to explain why the language faculty has the particular structure it has and not some other structure that is just as useful in human communication, we need to abandon the evolutionary biologist perspective (with its proprietary functional attributions) and replace it with the perspective of what Mayr called functional biology.

Functional biology, Mayr wrote, is "vitally concerned with the operation and interaction of structural elements, from molecules up to organs and whole organisms" and its "ever-repeated question is 'How?' How does something operate, how does it function?" (Mayr 1961: 1502). That is, the functional biologist "attempts to isolate the particular component he studies, and in any given study he fully deals with a single individual, a single organ, a single cell, or a single part of a cell" (Mayr 1961: 1502). Taking the functional biologist perspective on the function of language yields 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 (for an overview of biolinguistics see, amongst others, 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. My aim here is not to justify that the language as an instrument of thought hypothesis is a fecund explanation (see, for example, Chomsky 2013; 2013a or Asoulin 2016 for such an argument), but rather to show that different sciences can and do attribute different functions to the same biological system.

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 (see Chomsky 1995 or 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 (an I-language) that is encoded in individual brains.[6]

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 (for more discussion on human thought in this regard see Hinzen 2006; 2013). There is a discontinuity, a partial overlap, between animal thought and human thought, and what accounts for this discontinuity are the underlying mechanisms of language that structure this subset of thought in a particular way. In other words, while there is no doubt that animals think, there is little evidence that their thoughts display productivity and systematicity[7] that are the distinguishing marks of human thought (see Hauser et al. 2014 for a review of the literature showing the dearth of evidence in this regard; see also Berwick et al. 2011). 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 functional biologists. That is, the former 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 the latter 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, 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.

The explanatory context within which biolinguistics operates focusses on the nature of thought. One of the major aims of biolinguistics is to understand how the language faculty fits into the human cognitive system. It should be clear that the underlying mechanisms of language are of course not linguistic in nature, for if they were they would be explanatorily vacuous in regard to explaining the place of language in the human

---

[6] As an aside, note that this approach is certainly not unique to biolinguistics. For example, the study of human vision as a computational system, most famously articulated by Marr (1982), has been a very successful research programme.

[7] Linguistic productivity is part of the creative aspect of language use (see Asoulin 2013 for references and discussion). It is the ability to produce and understand an unlimited number of sentences. Linguistic systematicity refers to the fact that our ability to produce and comprehend expressions of a certain kind guarantees that we can produce or comprehend other systematically related expressions. Fodor (1975) is a classic discussion of the productivity and systematicity of thought and language. For systematicity, see Cummins (1996), Cummins et al. (2001), McLaughlin (2009).

cognitive system. Rather, these mechanisms encapsulate the process that takes place *before* expressions are given a phonological or semantic interpretation in a particular natural language. Notice again the different explanatory aims as compared with the above discussion of those who attribute to language the function of communication. In order to further elucidate this difference, let us consider the thought experiment from Reinhart (2006: 2ff.), which fleshes out one of Chomsky's thought experiments. She imagines a primate that has acquired the full set of human cognitive abilities but that does not have the language faculty. This fictitious primate would have, in addition to its cognitive abilities that allow it to think like its fellow primates, a set of concepts that is the same as that of humans and a set of sensorimotor systems that enable it to perceive and code information in sounds. Moreover, Reinhart imagines that this primate would also have the human system of logic, the abstract formal system that contains an inventory of abstract symbols, connectives, functions, and definitions necessary for inference. Given the nature of this primate, then, what would it be able to do with these systems? Can it make use of a human concept in its thought processes? That is, given all these additions but lacking the language faculty, can this fictitious primate add to its thinking abilities the sorts of thoughts that display productivity and systematicity and that at present appear to be unique to humans? Reinhart argues that it could not.

What prevents this fictitious primate from making use of the new systems and concepts it has acquired? What prevents it from thinking human thoughts? The primate can communicate its earlier preexisting concepts to its fellows, and it can make inferences typical of primates, but since it does not have the ability to create recursive and hierarchically structured expressions it cannot construct or comprehend propositions necessary for *higher-order* inference. In other words, this imagined primate has concepts and knowledge of first-order logic, which it can use and comprehend, but that is not enough to produce and comprehend propositions nor to make second-order and higher-order inferences. In order to be able to do the latter, the primate in the thought experiment must (but does not) possess recursion, a feature necessary for the abstract linguistic computational system (for more on recursion see Tomalin 2007; Zwart 2011). A fortiori, this primate cannot comprehend the entailment relations between propositions: it cannot think those sorts of thoughts. Now compare this fictitious primate to real world humans: we *can* think those sorts of thoughts. This is because the way in which the underlying mechanisms of language work in humans is by providing us with higher-order logic (see Crain 2012 on the relation between natural language and classical logic), by providing us with a computational system that creates recursive and hierarchically structured expressions that display productivity and systematicity and that we use to, amongst other uses, talk and think about the world.

It is clear from the above that, 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 we need to know the function of language in this particular explanatory context.

5 Concluding remarks

Biolinguistics and evolutionary theory 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. In other words, the same structure is given a different functional attribution depending on the explanatory theory within which it is embedded. There is no theory-neutral sense in which one could argue that biolinguistics has attributed the "wrong" function to language whereas evolutionary theory has attributed the "right" function. I think it is misleading to even phrase the debate in such a way. 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 particular functional attribution doing in a particular theory. That is, how does a particular functional attribution fit within an explanatory theory, and does this functional attribution help in deepening and making the explanation more fecund. The functional attribution that best fits the explanatory purposes of biolinguistics is that of language being an instrument of thought, whereas within evolutionary theory the function of language is communication. Both are valid within each theory and both theories add to our understanding of language and the role it has both in our cognition and the way in which we interact with the world. This argument generalises, for one can ask these two kinds of questions of other biological objects. For example, a theory aiming at an explanation of the role the eye plays in the visual system internal to our cognitive system will attribute a particular function to the eye. In contrast, a theory investigating the survival value that the eye offers a particular organism in a particular environment will probably attribute a different function to the eye. These are two valid questions asked about *the same* object, yielding different functional attributions.

References

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. 2013. The creative aspect of language use and the implications for linguistic science. *Biolinguistics* 7. 228–248.

Asoulin, Eran. 2016. Language as an instrument of thought. *Glossa: a journal of general linguistics* 1(1): 46. 1–23. DOI: http://dx.doi. org/10.5334/gjgl.34

Asoulin, Eran. 2017. *Functions, selected effects and empirical reality*. Unpublished manuscript, Macquarie University.

Ayala, Francisco J. 1970. Teleological explanation in evolutionary biology. *Philosophy of Science* 37(1). 1-15.

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.

Berwick, Robert C., Kazuo Okanoya, Gabriel J.L. Beckers & Johan J. Bolhuis. 2011. Songs to syntax: The linguistics of birdsong. *Trends in Cognitive Sciences* 15(3). 113–121

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, Mass.: The 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.

Crain, Stephen. 2012. *The emergence of meaning*. Cambridge: Cambridge University Press.

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

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

Cummins, Robert. 1996. Systematicity. *The Journal of Philosophy* 93(12). 591–614.

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. 2002. Neo-teleology. In Andre Ariew, Robert Cummins & Mark Perlman (eds.), *Functions: New essays in the philosophy of psychology and biology,* 157–172. Oxford: Oxford University Press.

Cummins, Robert, James Blackmon, David Byrd, Pierre Poirier, Martin Roth & Georg Schwarz. 2001. Systematicity and the cognition of structured domains. *The Journal of Philosophy* 98(4). 167–185.

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.

Dawkins, Richard. 1976. *The selfish gene.* Oxford: Oxford University Press.

Dennett, Daniel. 1987. *The intentional stance.* Cambridge, Mass.: The 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.

Ellis, Donald G. 1999. *From language to communication*. London: Lawrence Erlbaum Associates.

Fodor, Jerry A. 1975. *The language of thought.* New York: Crowell.

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

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.

Griffiths, Paul. 1993. Functional analysis and proper functions. *British Journal for the Philosophy of Science* 44. 409-422.

Hauser, Marc D., Charles Yang, Robert C. Berwick, Ian Tattersall, Michael J. Ryan, Jeffrey Watumull, Noam Chomsky & Richard C. Lewontin. 2014. The mystery of language evolution. *Frontiers in Psychology* 5. 1–12.

Hempel, Carl. 1959. The logic of functional analysis. In Llewellyn Gross (ed.), *Symposium on Sociological Theory*, 271-307. New York: Harper & Row.

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, Mass.: The MIT Press.

Kaplan, Jonathan M. & Massimo Pigliucci. 2001. Genes `for' phenotypes: A modern history view. *Biology & Philosophy* 16. 189-213. DOI:10.1023/A:1006773112047

Lehman, Hugh 1965. Functional explanations in biology. *Philosophy of Science* 32(1). 1-20.

Lewens, Tim. 2004. *Organisms and artifacts: Design in nature and elsewhere.* Cambridge, Mass.: The MIT Press.

Lewens, Tim. 2007. Functions. In Mohan Matthens & Christopher Stephens (eds.), *Philosophy of biology*, 525-547. Amsterdam: Elsevier.

Lewontin, Richard. 1983. Gene, organism and environment. In D. S. Bendall (ed.), *Evolution from molecules to men*, 273-285. Cambridge: Cambridge University Press.

Lewontin, Richard. 2000. *The triple helix: Gene, organism, and environment*. Cambridge, Mass.: Harvard University Press.

Machamer, Peter, Lindley Darden & Carl F. Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67(1). 1-25.

Marr, David. 1982. *Vision: A computational investigation into the human representation and processing of visual information.* New York: W. H. Freeman and Company.

Matthen, Mohan. 1997. Teleology and the product analogy. *Australasian Journal of Philosophy* 75(1). 21-37. DOI: 10.1080/00048409712347651

Mayr, Ernst. 1961. Cause and effect in biology: Kinds of causes, predictability, and teleology are viewed by a practicing biologist. *Science,* New Series, 134(3489). 1501–1506.
McLaughlin, Brian P. 2009. Systematicity redux. *Synthese* (170). 251–274.
Millikan, Ruth Garrett. 1984. *Language, thought, and other biological categories*. Cambridge, Mass.: The MIT Press.
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.
Nagel, Ernest. 1961. *The structure of science: Problems in the logic of scientific explanation.* New York, NY: Harcourt, Brace and World.
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. 2015. Functional analysis and the species design. *Synthese*. 1-22. DOI: 10.1007/s11229-015-0940-9
Pigliucci, Massimo & Jonathan M. Kaplan. 2006. *Making sense of evolution: The conceptual foundations of evolutionary biology.* Chicago: The University of Chicago Press.
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. DOI: 10.1080/09515089.2013.800454
Roth, Martin & Robert Cummins. 2017. Neuroscience, psychology, reduction and functional analysis. In David M. Kaplan (ed.), *Explanation and integration in mind and brain science*, XX-XX. Oxford: Oxford University Press.
Ruse, Michael E. 1971. Function statements in biology. *Philosophy of Science* 38. 87-95.
Searle, John. 1995. *The construction of social reality*. New York: The Free Press.
Sorabji, Richard. 1964. Function. *Philosophical Quarterly* 14(57). 289-302.
Thagard, Paul. 2012. *The cognitive science of science: Explanation, discovery, and conceptual change*. Cambridge, MA: MIT Press.
Tomalin, Marcus. 2007. Reconsidering recursion in syntactic theory. *Lingua* 117. 1784–1800.
Walsh, Denis M. 2008. Function. In Stathis Psillos & Martin Curd (eds.), *The Routledge companion to philosophy of science*, 349–357. Abingdon, UK: Routledge.
Walsh, Denis M. 2012. Mechanism and purpose: A case for natural teleology. *Studies in History and Philosophy of Biological and Biomedical Sciences* 43. 173-181.
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.
Zwart, Jan-Wouter. 2011. Recursion in language: A layered-derivation approach. *Biolinguistics* 5(1-2). 43–56.