10-1-1998

Etiological Theories of Function: A Geographical Survey

David J. Buller

Original Citation

Buller, D.J. Etiological Theories of Function: A Geographical Survey. Biology and Philosophy 13:4, 505-527.

Etiological Theories of Function: A Geographical Survey

DAVID J. BULLER
Department of Philosophy
Northern Illinois University
DeKalb, IL 60115
USA

Abstract. Formulations of the essential commitment of the etiological theory of functions have varied significantly, with some individual authors' formulations even varying from one place to another. The logical geography of these various formulations is different from what is standardly assumed; for they are not stylistic variants of the same essential commitment, but stylistic variants of *two* non-equivalent versions of the etiological theory. I distinguish these "strong" and "weak" versions of the etiological theory (which differ with respect to the role of selection in their definitions of *function*), draw out their respective implications, and argue that the weak version is to be preferred to the strong.

Key words: Cummins function, etiological theory, proper function, selection

1. Introduction

The etiological theory of biological functions – which defines the function of a trait in terms of its evolutionary history – has attracted a growing number of adherents (see, e.g., Millikan 1984, 1989b; Neander 1991a, 1991b; Griffiths 1993; Godfrey-Smith 1993, 1994; Mitchell 1995; Allen & Bekoff 1995). Although there have been some intramural disagreements among them regarding the details of a fully developed etiological theory, there have also been explicit statements of agreement regarding its essential commitment. But the formulations of this essential commitment have varied significantly, with some individual authors' formulations even varying from one place to another. Attention to these various formulations reveals a logical geography different from what is standardly assumed; for the various formulations are not stylistic variants of the same essential commitment, but stylistic variants of two non-equivalent versions of the etiological theory. There are in effect *two* etiological theories on the market, which have not been distinguished. In what follows, I will distinguish these two versions of the etiological theory, draw out their respective implications, and argue that one version is to be preferred over the other.

2. Two etiological theories

One standard formulation of the etiological theory is the following "selected effects" formulation:

> A current token of a trait *T* has the function of producing an effect of type *E* just in case, at some point in evolutionary history[1], there was selection for *T* (over alternative items) because of its having produced effects of type *E*.

The parenthetical clause in this formulation is explicit in Millikan: "Only if an item or trait has been *selected* for reproduction, *as over against other traits*, *because* it sometimes has a certain effect does that effect count as a function" (1993: pp. 35–36). But it is not explicit in some of Neander's formulations: "the central element of the etiological approach should be seen as the simple idea that a function of a trait is the effect for which that trait was selected" (1991b: p. 459; cf. 1991a: p. 173). Regardless of whether the parenthetical clause is explicit, however, I will show below that any formulation that appeals to selection in defining the function of a trait is committed not only to the parenthetical clause, but to much stronger conditions as well.

But first note an undesirable consequence of the "selected effects" formulation, which stems from phenomena such as the segregation distorter gene in *Drosophila melanogaster*. Although Mendelian segregation is normal in females carrying the segregation distorter gene, in heterozygote males the gene disrupts meiosis by causing its allele on the homologous chromosome to sabotage spermatogenesis in the homologue, thereby ensuring that its allele tends to be present only in dysfunctional sperm (Crow 1979; Lyttle 1993). As a result, the segregation distorter is present in 95–99% of a heterozygote's functional sperm and is consequently transmitted to nearly all its progeny. When a homozygous male is produced, it is either sterile or dies before maturing to reproductive viability. As Godfrey-Smith has pointed out, the segregation distorter gene thus satisfies the "selected effects" formulation of the etiological theory, since "disrupting meiosis is something that segregation distorter genes do, that explains their survival" (1994: p. 348). Yet biologists do not ascribe functions to such subversive traits of organisms, instead reserving function ascriptions for generally "helpful and constructive" traits (p. 347). Godfrey-Smith proposes to avoid this consequence by building into the etiological theory the requirement that the selected trait be contained within, and make a positive contribution to the fitness of, a larger system. This generates the following formulation, which I will call the *strong theory* (cf. Godfrey-Smith 1994: p. 359):

> A current token of a trait *T* in an organism *O* has the function of producing an effect of type *E* just in case past tokens of *T* contributed to the fitness

of *O's* ancestors by producing *E and* were selected for (over alternative items) *because of* this contribution to the fitness of *O's* ancestors.

While Godfrey-Smith is clearly right that the bare selected effects formulation has undesirable consequences, which can be avoided by the strong theory, it is likely that the selected effects formulation was only intended as an elliptical formulation of the strong theory. More explicit statements by Neander and Millikan seem to suggest that the selected effects formulation is indeed elliptical. For example, Neander says:

> It is the/a proper function of an item (*X*) of an organism (*O*) to do that which items of *X's* type did to contribute to the inclusive fitness of *O's* ancestors, and which caused the genotype, of which *X* is the phenotypic expression, to be selected by natural selection. (1991a: p. 174)

And Millikan says:

> items have functions when their being there depends on reproduction from ancestors having similar traits, these traits having been causally efficacious in helping to produce these items, and these traits having been selected at some point in this history for their capacity to make this kind of contribution. (1993: p. 41)

These formulations satisfy Godfrey-Smith's requirement that a functional trait make a contribution to its bearer (as does the formulation of Griffiths 1993: p. 415). So, in what follows, I will take the strong theory to be the official version of the selected effects formulation, and focus only on the strong theory.

In contrast to the strong theory, the etiological theory is sometimes articulated without mention of selection for the functional trait, but with an emphasis only on the requirement that the functional trait be a reproduction of items that had the same effect(s). For example, Millikan says that a function of a trait is "a function that its ancestors have performed that has helped account for proliferation of the genes responsible for it, hence helped account for its own existence" (1989b: p. 289). Of course, the way a functional trait "helped account for proliferation of the genes responsible for it" was by contributing to the fitness, hence reproduction, of the bearers of the genes for that trait (cf. Griffiths 1993: p. 412). This idea is encapsulated in the following formulation, which I will call the *weak theory*:

> A current token of a trait *T* in an organism *O* has the function of producing an effect of type *E* just in case past tokens of *T* contributed to the fitness of *O's* ancestors by producing *E*, and thereby causally contributed to the reproduction of *Ts* in *O's* lineage.

Like the strong theory, the weak theory appeals to a trait's history to characterize its function, defining the function of a current token of a trait in terms of the role played by ancestral tokens. In spite of this commonality, however, the two formulations are not equivalent; for, while the strong theory entails the weak, the converse does not hold.

To see why, consider first the strong theory. By defining the function of a current token of a trait T as the production of an effect (in O's ancestors) because of which there was selection for past tokens of T, the strong theory is committed to the following, which are necessary conditions for the occurrence of natural selection: T must be hereditary; at some point in evolutionary history there must have been variation in T; and the bearers of T must have had greater fitness than bearers of T's variants (at least partly) because of possessing T. However, even these conditions, as stated, are not sufficient for there to have been selection for T. For, as Brandon (1990) has shown, selection requires variation *within a common selective environment*. If the variation in T was distributed through the population by group (such that all and only members of group G_i possessed the variant T_i) and each group faced different environmental demands, then whatever differential survival and reproduction occurred was due not to differences in fitness of the members of those groups, but to differences in the environments they inhabited. Since natural selection requires differential fitness, and differential fitness involves differences in the *ability* to survive and reproduce, only if differently endowed organisms face common environmental demands can their differential survival and reproduction be the result of differences in their abilities to survive and reproduce. To illustrate, if I lift a chair, but you fail to lift an identical chair that is bolted to the floor, the difference in our performances is due not to a difference in our lifting abilities (the strength of our muscles), but to a difference in our environments (chair bolted to the floor versus one that is not). Thus, by defining the function of a trait by appeal to selection for it, the strong theory attributes a function to a trait T only when (1) T is hereditary, (2) there has been variation in T within a common selective environment, and (3) the bearers of T had greater fitness within that common environment than bearers of T's variants (at least partly) because of possessing T.

The weak theory, in contrast, does not define the function of a trait in terms of selection for it, and thus does not make it a necessary condition of T's having a function that there was variation in T. And, since variation is not required, it is also not necessary that the bearers of T had greater fitness than any other organisms. The weak theory requires only that T contributed to the fitness of the ancestors of T's current bearers *by* producing an effect E and that this, in turn, contributed to the reproduction of Ts. The key to these requirements is the notion of a trait's *contributing to fitness* by producing its

effect(s). But it may appear that the role of this notion is problematic in the weak theory in a way that it is not in the strong theory. For, although the strong theory also invokes the notion of a trait's contribution to fitness, its invocation of selection may make it appear to require appeal only to the *relative* fitness of the organisms possessing variants of the functional trait: the functional trait was reproduced because its bearers had *greater fitness* than organisms possessing its variants. Since the weak theory does not require that functional traits ever exhibited variation, however, the weak theory needs some non-relative way of characterizing a trait's contribution to fitness; that is, the weak theory needs to be able to specify how traits that did not make their bearers more fit than other organisms nonetheless could have contributed to their fitness by producing certain effects. I will show in a moment that any difference between the strong and weak theories on this point is merely apparent – that the notion of contribution to fitness at work in the weak theory is operative in the strong as well. But first let me articulate this notion and its role in the weak theory.

Fitness is standardly taken to consist of the following components: viability, fertility, fecundity, and mating ability (Endler 1986: pp. 81–86). These components of fitness in turn consist of complex causal processes within and among an organism's various vital systems. A trait T can thus be said to make a contribution to fitness by producing an effect E if T's producing E constitutes a causal contribution to a complex process involved in a component of fitness. This requires, in turn, a way of precisely characterizing a trait's causal contribution to such a complex process. And there is really only one well developed way of characterizing an item's causal contribution to a complex process – namely, Cummins' theory of functions (Cummins 1975). According to Cummins' theory, an item X has a function only relative to a capacity C of a system S containing X, where the function of X emerges from a *functional analysis* of the capacity C of S. A functional analysis of C proceeds by analyzing C into the capacities of simpler components of S in such a way that C emerges as the "programmed manifestation" of the exercise of the capacities of those simpler components (where the latter capacities may themselves admit of functional analysis, until the analysis terminates in the dispositions of basic structural components of S). It is the *Cummins function* of X to produce effect E, then, just in case X is a component of S and X's producing E features in a functional analysis of some capacity C of S. In this way, Cummins functions are always relative to a system containing the functional item. To characterize an item's contribution to fitness, then, merely involves taking fitness (or one of its components) as the containing system on which the functional analysis is to be performed. Thus, if a trait T's producing some effect E appears in a Cummins-style functional analysis of a

component of fitness (that is, if T has the Cummins function of producing E within a component of fitness), then T causally contributes to that component of fitness by producing E and thereby contributes to fitness.

But the weak theory also requires that ancestral tokens of T must have "causally contributed to the reproduction of Ts". This requirement is satisfied provided that T is also hereditary. For any hereditary trait that makes a causal contribution to viability, fertility, fecundity, or mating ability will make a causal contribution to its own reproduction – by causally contributing to the reproduction of the genotype for that trait, which in turn causally contributes to the development of another token of that trait. The weak theory thus attributes a function to any trait that is (1) hereditary and (2) a descendant of tokens that performed a Cummins function within the fitness of ancestral bearers of the trait. For any trait that satisfies these conditions will have causally contributed to its own reproduction via genetic mechanisms of inheritance and development. In short, then, according to the weak theory the function of a hereditary trait T in an organism O is the Cummins function that T had within the fitness of O's ancestors, since performance of that function causally contributed to the reproduction of Ts in O's lineage.[2]

At this point it may appear that something has gone awry. For, with the notable exception of Griffiths (1993), it has become commonplace to sharply contrast the etiological theory with Cummins' theory of functions. So my importation of Cummins' theory into the characterization of the weak theory may appear to violate the spirit of the etiological theory. In the next section I will argue at length that the weak theory in fact satisfies not only the spirit, but also the letter of the etiological theory. For the moment, however, I would like to show that even the strong theory requires Cummins' theory to make sense of the notion of O's ancestors' being favored by selection *because of* possessing tokens of T that produced effects of type E (or, equivalently, of there having been selection for T because of its contribution to the fitness of O's ancestors).

To see why, consider first that this notion *cannot* be understood as meaning that bearers of T had greater fitness than bearers of its variants because T had a higher fitness value than its variants. For the fitness value of a trait is simply the average fitness of the bearers of that trait; so this interpretation would degenerate into the tautology that bearers of T had greater fitness than bearers of T's variants because they had higher average fitness than bearers of T's variants. In addition, since the fitness value of a trait is the average fitness of its bearers, correlated traits (products of pleiotropy or genetic linkage) always have the same fitness values. So, since warm coats are correlated with heavy coats in polar bears, an appeal to the relative fitness values of traits could never show that current polar bears' ancestors were favored by selection because

their coats produced greater warmth, rather than because their coats produced greater weight; for a heavy coat has the same fitness value as a warm coat. And this means that such an appeal could not show the coat's function to be the production of greater warmth, rather than greater weight.

Of course, the strong theory actually does solve this problem, since it defines *function* in terms of what there was "selection for"; and there was *selection for* warm coats, but merely *selection of* heavy coats (on the distinction, see Sober 1984: pp. 97–102). But the selection for/selection of distinction has become so well entrenched in the conceptual landscape that it is easy to forget what the distinction encapsulates and, consequently, to overlook precisely what is involved in defining a trait's function in terms of what it was selected for. To distinguish traits that there was selection for from those there was merely selection of requires a focus not on differences between organisms with variants of a trait, but on the differences between traits within a group of organisms that do not vary in those traits. Within such a group, traits that there was selection for are distinguished from traits there was merely selection of by the fact that the former *causally contributed to fitness* whereas the latter did not. As Sober puts it: "To say that there is selection for a given property means that having that property *causes* success in survival and reproduction" (1984: p. 100). But this means that specifying what there was selection for requires an antecedent and independent way of ascertaining which traits made a causal contribution to fitness. And here we are driven back to a Cummins-style functional analysis of fitness in order to identify which traits made causal contributions to fitness and how they did so. Thus, in virtue of defining the function of a trait as the effect that the trait was selected for producing in ancestral organisms, even the strong theory must invoke Cummins' theory in order to identify the selected-for effect (see Walsh & Ariew 1996: pp. 508–509, for a similar argument). In effect, there was selection for a trait *T's* producing some effect *E* only if *T* had the Cummins function of producing *E* within a component of fitness of *T's* bearers at that point in evolutionary history at which there was selection for *T*.

In light of all this, the difference between the strong and weak theories can be summarized in sharp focus. Both theories entail that a trait *T* has the function of producing an effect of type *E* in an organism *O* only if *T* satisfies the following two conditions:

(ET1)   *T* must have contributed to the fitness of *O's* ancestors by producing effects of type *E*, and

(ET2)   *T* must be hereditary.

Since both theories are committed to these two conditions, I will take them to represent the essential commitment of the etiological theory. But, whereas

these two conditions exhaust the requirements of the weak theory, the strong theory adds the following condition:

(ST3)   at some point in evolutionary history, there must have been selection for *T* because of its producing *E*.

(ST3) entails, in turn, the strong requirements that

(ST3.1) there must have been variation with respect to *T* at some point in evolutionary history, and

(ST3.2) *T's* bearers must have had greater fitness than bearers of its variants because of *T's* producing *E*.

The weak theory thus attributes functions to traits more liberally than the strong theory, since it counts as functional any hereditary trait of an organism that causally contributed to the fitness of that organism's ancestors, regardless of whether there was selection for that trait. Indeed, in accordance with the above requirements, the strong theory's function attributions constitute a proper subset of the weak theory's function attributions.

To illustrate the difference between the strong and weak theories, suppose that *T* is a hereditary physiological trait that plays a causal role in the process of gamete production, but that either (a) through genetic happenstance the necessary mutation(s) never occurred in the population to produce an alternative to *T* (cf. Kitcher 1993: p. 388), or (b) mutation produced alternatives to *T* in the population, but no two of the alternatives occurred within a common selective environment. In either case, there cannot have been selection for *T*. But, since *ex hypothesi T* does play a causal role in gamete production, *T* does causally contribute to the fitness of its bearers. Under these conditions, the strong theory would withhold a function attribution from *T*, since there has never been selection for *T*. But, since *T* does make a hereditary contribution to the fitness of its bearers, *T* does have some effect in virtue of which it gets reproduced across generations; so the weak theory would attribute a function to *T*. Thus, while the strong theory looks for a history of selection for a trait, the weak theory looks only for a history of a hereditary contribution to fitness.

It is in fact not merely a logical possibility that there are hereditary traits that contribute to fitness but for which there has not been selection. For, first, genetic drift sometimes produces hereditary traits that contribute to the fitness of their bearers (Brandon 1990: p. 175). Second, studies of selection "in the wild" have demonstrated the existence of selection for a large number of morphological traits, but for strikingly few physiological traits and virtually no biochemical traits. The most probable reason for this disparity is a far greater incidence of variation in morphological traits than in physiological and biochemical traits (see Endler 1986: pp. 154–160). And, if there is a very low incidence of variation in physiological and biochemical traits, there is

perhaps an even lower incidence of variation among such traits within common selective environments. But, if there are physiological or biochemical traits that have not varied in a common selective environment in natural populations, there cannot have been selection for those traits, regardless of their roles in contributing to fitness. In each of these cases, the strong theory entails that the relevant traits are functionless, while the weak theory can attribute functions to them.

There is a third type of case in which the strong and weak theories would diverge in their function ascriptions, and which is related to the second case mentioned above. Many functional traits are complex, involving an internal articulation into hereditary component parts that produce effects required for the functional trait's own proper effect. That is, if T is a complex trait that has the function of producing an effect E, then T contains component parts $t_1 - t_n$, each of which produces some effect that is necessary for T's producing E, and where the concerted production of the effects of $t_1 - t_n$ constitutes T's producing E. In such cases, even if there was variation in T at some point in evolutionary history and T was selected over its variants, it doesn't follow that there was variation in all the components of the variants of T. For example, the component t_1 may have been invariant in all the variants of T, those variants differing only with respect to other component parts. In that case, although there was selection for T, there would not have been selection for t_1 (cf. Amundson & Lauder 1994: p. 459). The strong theory, consequently, would not ascribe a function to t_1. But the effect it produces is causally necessary for T's producing its selected effect. So, if T has the function of producing E, then t_1 should be ascribed the function of producing its effect. Function ascriptions should "trickle in" from a complex trait to its hereditary components that produce effects that are necessary for the complex trait's own function. And the weak theory would generate such function ascriptions, although the strong theory would not.

Thus, the weak theory enjoys the following advantages over the strong: The weak theory can attribute functions to (1) traits produced by genetic drift, (2) traits that have not faced actual competition from alternatives during evolutionary history, and (3) all hereditary components of complex functional traits, while the strong theory withholds function attributions in all such cases.[3]

3. "An ambiguity in the notion 'function' "?

It is now time to address the objections that may have been accumulating in response to the previous section. These objections, I suspect, will have grown out of a single underlying source: the idea that the weak theory is the

product of an improper conflation of two distinct concepts of *function* and that any "advantages" of the weak theory are *ipso facto* spurious. For it has become common to claim that there are two distinct concepts of *function* at work in biology: "proper functions", which are explicated by the etiological theory, and "Cummins functions." It is claimed in addition that these two concepts are in the service of different explanatory projects in biology – *proper function* at work in evolutionary theory and behavioral ecology to explain the presence of a functional trait in its bearers, and *Cummins function* at work in physiology and developmental biology to explain the causal contribution of a functional item to a complex process (Millikan 1989a: pp. 175–176; Godfrey-Smith 1993: pp. 200–201, and 1994: pp. 344–345). Acceptance of this alleged bifurcation in the concept of *function* in biology has spawned attitudes ranging from liberal pluralism – "that there is room in biology for at least two different notions" (Allen & Bekoff 1995: p. 612) – to adamant separatism – that the two concepts must be kept distinct "in the interests of maintaining an accurate understanding of different explanatory strategies" (Godfrey-Smith 1993: p. 204).

It is certainly true that the concept of a Cummins function *per se* is distinct from the concept of a proper function. For Cummins' theory wasn't even intended specifically as a theory of *biological* functions; indeed, nothing in Cummins' theory restricts functional analyses to *any* specific types of containing system. Cummins' theory carves out a very broad range of phenomena: wherever there is a complex system – economic, administrative, cognitive, respiratory, or internal combustion – there are Cummins functions contributing to its overall capacities and dispositions. So Cummins' theory is properly thought of as providing a fully general way of characterizing the functional contributions of components to the systems containing them. But what if we stipulate organisms as the containing systems on which to perform functional analysis? Are proper functions then simply Cummins functions within organisms? The common answer is unequivocally "no" – that, although there is considerable overlap in the extensions of the two concepts in biological systems, the cases of extensional agreement are nonetheless generated in response to very different theoretical questions (Millikan 1989a: pp. 175–176; Godfrey-Smith 1993: p. 201, and 1994: p. 345). Ascription of a proper function purportedly answers a question about why the functionally characterized item is present in its bearer, while ascription of a Cummins function purportedly only answers a question about how the functionally characterized item contributes to a feature of its bearer.

But cases of agreement in the assignment of proper functions and Cummins-functions-within-organisms are, I think, more principled than the pluralist view makes them out to be. The significance of the etiological the-

ory should be seen not in its alleged total departure from a Cummins-style understanding of functions, but in the types of principled constraint it places on the application of Cummins-style functional analysis. The argument of the previous section – that defining an item's function in terms of what it was *selected for* doing must resort to a Cummins-style functional analysis of fitness in order to characterize the item's selected-for effect – went part way toward demonstrating this. In addition, however, the pluralist maneuver of relegating physiological functions to Cummins' theory, rather than modifying the etiological theory to encompass them, is rather unsatisfactory. For, since Cummins' theory recognizes no distinctions among kinds – natural or otherwise – of containing system, this maneuver has the effect of treating physiological functions as having more in common *qua functions* with changes in the interest rate than with changes in pigmentation to avoid predation. This chops the domain of functional phenomena through the bone, rather than carving it at the joints; *prima facie*, cases of biological function should have more in common with one another than with functions in non-biological domains. So, an account of the concept of *function* in biology that provides a principled unity and shows what distinguishes biological from non-biological functions should be preferred to a bifurcated account that fails to properly separate some biological functions from non-biological functions.

The weak theory provides such a unified account of functions in biology; for its function ascriptions encompass not only the functions assigned in evolutionary theory and behavioral ecology (where selection for the functional item features prominently), but those assigned in physiology and developmental biology as well (where selection for the functional item is not an important consideration *if* it has occurred at all). Thus, in addition to the advantages detailed in the last section, the weak theory also has parsimony in its favor. To show this, however, and to show that the weak theory is not simply the product of an improper conflation of distinct concepts of *function*, it is necessary to reexamine the arguments that have motivated and sustained the sharp contrast between proper functions and Cummins functions. These arguments show that Cummins' theory fails to satisfy certain important desiderata for a theory of biological functions. But they have typically also been taken to show that these desiderata can only be satisfied by a theory of functions that appeals to selection for the functionally characterized item – in other words, by the strong theory. By reexamining these arguments, I will show that in fact the weak theory satisfies all the desiderata for a theory of functions set by the standard etiological critiques of Cummins' theory, and that appeal to selection for the functional item (as in the strong theory) is unnecessary. Consider, then, the three main objections to Cummins' theory and how the weak theory meets them.

516

3.1. *The liberality objection*

Millikan argues that Cummins' theory is too liberal in assigning functions, because "there [is] no limit to the numbers and kinds of things that can be looked upon as being 'systems' or 'parts of systems'" (1989a: p. 175). For example, if we take the Earth's water-cycle as the system with which we begin functional analysis, then, since clouds produce rain that makes vegetation grow, Cummins' theory entails that a function of clouds is to make vegetation grow (1989b: p. 294). But it is clearly not the *purpose* of clouds to make vegetation grow (1989b: p. 294); clouds are not the product of *design* for making vegetation grow (1989a: p. 175). Thus, having a purpose or having been designed for some end is not a condition an item must satisfy in order to have a Cummins function. However, in the paradigm case in biology, having a function is a matter of having a purpose (1989b: p. 293). And function ascriptions can capture the purposes of functionally characterized items only if they are made with a concept of *function* that builds in reference to the history of selection that shaped the functional item (1993: p. 38).

The weak theory, however, attributes a function to a current token of a trait only when past tokens of the trait performed a Cummins function within a component of ancestral fitness. Consequently, the weak theory applies Cummins-style functional analyses only to those systems that possess the property of fitness – that is, only to *systems that undergo natural selection* (cf. the strong theory of Griffiths 1993: pp. 412–413). Like the strong theory, then, the weak theory's function attributions are grounded in selection, although in a different way. One way to look at the difference between the theories is as follows. The weak theory sees functionally integrated systems undergoing natural selection, and assigns functions to parts of those systems in virtue of the contributions the parts made to overall ancestral system fitness. A current token of a trait has the "purpose" of doing *X*, then, if past tokens made a contribution to the reproductive success of their bearers by doing *X*. In contrast, with its emphasis on direct selection for traits, the strong theory finds the purpose of a trait in *the trait's differential success* in the process of natural selection. In short, the weak theory's function assignments are grounded in *selection of adapted systems*, whereas those of the strong theory are grounded in *selection for traits* of those systems.

Organisms, of course, are the paradigmatic examples of adapted systems; but the weak theory can attribute functions within any type of system that biological theory finds it necessary to represent as possessing fitness. Thus, the weak theory is neutral with respect to whether DNA segments, groups, or populations are adapted systems (as is the strong theory of Godfrey-Smith 1994; see p. 349). But, however biological theory settles the issue of which are the adapted systems, the weak theory's restriction of Cummins-style

functional analyses to adapted systems places a significant and principled constraint on the types of thing that "can be looked upon as being systems." Unlike Cummins' theory, then, the weak theory is not overly liberal with respect the types of system to be functionally analyzed.

But liberality may still seem to be a threat "from within"; for it may seem that the weak theory will still be overly liberal in its assignment of functions to *parts* of adapted systems. The objection lurking here concerns cases like the flying fish discussed by Williams (1966: pp. 11–12). Although there is an adaptive advantage to its flight (presumably predation avoidance), it is physiologically necessary that a flying fish return to water, since it can't survive in air. The fish's weight contributes to its returning to water, and hence contributes to its survival. A functional analysis of the fitness of flying fish, then, might seem to lead the weak theory to attribute to the fish's weight the function of returning it to water. But, as Williams pointed out, the fish's return to water is a simple consequence of laws of physics and it would be "absurd" to postulate special *biological* functions "to achieve the mechanically inevitable" (p. 12). Attributions of biological function "should not be invoked when less onerous principles, such as those of physics and chemistry or that of unspecific cause and effect, are sufficient for a complete explanation" (p. 11). Since the return to water is completely explainable by physical principles, Williams concludes, from a functional standpoint the "real problem is not how it manages to come down, but why it takes it so long to do so" (p. 12).

However, *every* functional trait of an organism produces its functional effect as a matter of "physical inevitability", at least in the sense that the causal process by which it produces that effect falls under the laws of physics and chemistry. So what is the distinction that Williams is after? Earlier in his discussion, Williams argues that a phenomenon is proprietarily biological only if it undergoes the process of natural selection in addition to processes explained by the laws of physics and chemistry (p. 5). Of course, the feature of the process of natural selection that cannot be completely explained solely in terms of the laws of physics and chemistry is the *differential reproduction of variants* in a population. Thus, the reason that returning to the water is not a properly biological process is that the individual fish in a population of flying fish could not possibly *vary* with respect to whether they return to water. That is, there could never be an allele for weightlessness, so natural selection could never act to favor weighted fish over weightless fish. Thus, since the principle of natural selection could not *possibly* play any role in explaining the return to water, it is "completely explainable" by the laws of physics and the return to water is, in this sense, "physically inevitable." The point of Williams' argument, then, is that we should never consider an item

to be functional (in the proprietary biological sense) if it could not possibly exhibit the sort of variation that is required for natural selection – that is, if it could not possibly be acted on by natural selection.

Seen in this light, it becomes clear that the weak theory is not as liberal in its assignment of functions as the objection suspected. For the weak theory assigns functions *only* to *hereditary* traits that made contributions to ancestral fitness (recall (ET2)). A trait is hereditary, however, only if it is the developmental product of a gene (or genes) transmitted across generations. There being "a gene for" some trait in turn entails the possibility of allelic variation at the locus occupied by that gene. So the weak theory assigns functions only to traits that could vary as a result of genetic variation. The difference between the weak theory and the strong theory is that the latter requires of a functional trait that it *actually has* exhibited variation, while the weak theory requires only that it is *possible* for the trait to exhibit variation. The weak theory, then, would not assign a function to the flying fish's weight *per se*; for having weight is not hereditary, since there can be no allele for weightlessness.

The weak theory would, however, under certain circumstances assign a function to the *specific weight* of the fish – not, that is, to its weighing *something*, but to its weighing, say, 7.6 pounds. For the specific weight of flying fish is something that can vary as a result of genetic variation; so specific weight is hereditary. But there is no reason why Williams' arguments should be taken as providing a principled block to such function ascriptions. As Williams said, the functional problem regarding flying fish is not that they return to water, but *how long* they take to do so. So, if specific weight were a contributing factor to the degree of efficiency of the fish's gliding mechanism (i.e., the duration of flight), the weak theory would assign it the function of making its specific contribution. But this would not actually violate any of Williams' defended restrictions on function ascriptions. Flying fish aside, now, the weak theory's requirement that functional traits be hereditary prevents it from being overly liberal in its assignment of functions to parts of adapted systems.

3.2. *The malfunction objection*

It is sometimes argued that diseased or malformed organs have functions, although Cummins' theory fails to assign them any (see, e.g., Millikan 1989b: pp. 294–295; Neander 1991a: pp. 181–182). For example, it is the function of the kidneys to filter metabolic wastes from the blood; but many current tokens of kidneys are diseased and, hence, unable to filter metabolic wastes from the blood. Nonetheless, it is still the function of diseased kidneys to filter metabolic wastes; diseased kidneys are just malfunctioning or not performing their function. But, since a Cummins-style functional analysis of a body

containing a diseased kidney would not reveal that kidney to be producing the effect of filtering metabolic wastes from the blood, Cummins' theory would not assign it the function of filtering metabolic wastes. Thus, it is argued, only by appealing to what the kidneys are *supposed to do* – and not to what some current token actually does – can we correctly say that it is the function of a diseased kidney to filter metabolic wastes and that it is failing to perform its function. And this normative "supposed to" can only be captured by a concept of *function* that builds in what a trait was selected for doing.

The weak theory, however, assigns a current token of a trait the function of producing some effect *E* just in case that trait's producing *E* appears in a Cummins-style functional analysis *of ancestral systems* possessing that trait. That is, the weak theory does not apply Cummins-style functional analyses to current systems themselves, but to the ancestors of those systems. The function of a trait is then the Cummins function it performed within the fitness of those ancestral systems and in virtue of which it got reproduced. Thus, the weak theory would also entail that a diseased current token of a kidney has the function of filtering metabolic wastes from the blood, since that is the Cummins function that kidneys performed in – the contribution they made to the fitness of – ancestral systems and in virtue of which kidneys got reproduced. So, like the strong theory (and unlike Cummins' theory), the weak theory assigns functions to malfunctioning or non-functioning traits. The strong theory's appeal to selection for the functional trait over alternatives plays no essential role in such function assignments.

3.3. *The "unexplained presence" objection*

Another objection to Cummins' theory is that one of the principal explanatory goals of function attributions in biology is to help explain the *presence* of the functional item, and the function attributions licensed by Cummins' theory do not achieve this explanatory goal (Millikan 1989a: p. 175; Godfrey-Smith 1993: pp. 200–201 & 204). Considering the water-cycle again, making vegetation grow is not an effect of clouds that helps explain why there are clouds. However, when an item is attributed the function of producing some effect when, and only when, there has been selection for that item because of its producing that effect, the function attribution does help explain the presence of the functional item in the organisms possessing it. Thus, it is argued, function attributions that explain the presence of the functionally characterized item in its bearers must appeal to selection for that item.

When the weak theory attributes a function to a trait, however, past tokens of that trait will have made some hereditary contribution to the overall fitness of the ancestral systems that possessed the trait; and this hereditary contribution to the fitness of those ancestral systems helps explain why that trait

got reproduced, hence why it is present in descendant systems. Thus, the weak theory's function ascriptions explain the presence of the functionally characterized items just as do the strong theory's function ascriptions. In fact, even under the function ascriptions of the strong theory the presence of a trait in its bearer is explained only by the combination of (1) the fact that the trait is hereditary and (2) the fact that it contributed to the fitness of ancestral systems. What the strong theory adds to these two conditions – namely, past selection for that trait over its variants – does not in fact contribute to the explanation of the presence of a trait in any of its current bearers. Even Millikan, at one point, acknowledges this fact: "The requirement on a certain kind of cause of the functional item's being there is fulfilled by the cycling [of reproduction and development], not the selection" (1993: p. 41; see p. 37 for the bracketed insertion). This, however, should not be surprising. For natural selection causes only changes in gene frequencies and, hence, changes in the *frequencies* of traits in a population; selection is not a cause of the presence of *any* trait in any individual organism (Cummins 1975: pp. 750–751; Endler 1986: pp. 46 & 241; Maynard Smith 1993: p. 20; Mayr 1988: p. 98; Sober 1984: pp. 149–152).

Neander (1995), however, gives an interesting argument that selection *is* a cause of the presence of a trait in its bearers.[4] If her argument is right, then the weak theory's function attributions *do* fail to explain the presence of its functionally characterized traits. Neander's reasons derive from the fact that most traits of any complexity are the developmental products of "gene sequences," which are in turn the product of cumulative selection. Her argument is clearest if we consider a simple example of a hypothetical species of haploid, uniparental organisms. Suppose that all organisms in the current generation of this species have the gene sequence <A,B,C>, which has evolved from <a,b,c> through successive mutations (*a* to *A*, etc.), where selection so strongly favored each new sequence that it was driven to fixation generations before another mutation. Suppose further that the sequence <A,B,C> produces the trait *T* (during normal development). To explain the presence of *T* in an organism *S* of the current generation, we would cite the developmental process by which <A,B,C> produces *T*. But Neander argues that a more complete causal explanation would cite the reproductive process by which *S* inherited <A,B,C> from its parent. An even more complete causal explanation, she argues, would trace the chain of inheritance of <A,B,C> back through *S's* lineage. Eventually, however, we will have to explain the *causal origin* of the <A,B,C> sequence, with some organism *Z*, by citing the fact that *c* mutated to *C*. But, Neander argues, mutation is only *part* of the causal explanation of the origin of <A,B,C>. The other part of the explanation involves *selection*, since <A,B,C> would not

have arisen by mutation of *c* to *C* with *Z* unless there had been prior selection for <A,B,c>. If only <A,b,c> had been available in the generation before *Z*, the production of <A,B,C> would have required not only the mutation of *c* to *C*, but that of *b* to *B* as well. Since the probability of two mutations occurring is much smaller than the probability of only one, the fact that selection drove <A,B,c> to fixation by the generation before *Z* – as opposed to favoring <A,b,c> or favoring neither over the other – greatly increased the probability that <A,B,C> would occur. The fact that selection thus increased the probability that <A,B,C> would occur, Neander concludes, shows that it is a causal factor in the origin of the <A,B,C> sequence and, hence, that it is a historical cause of *S's* possessing *T*.

But the fact that selection for <A,B,c> increased the probability that <A,B,C> would occur is *not* in fact relevant specifically to explaining the presence of *T in S*. For selection for <A,B,c> merely increased the probability that <A,B,C> would occur *in the population* of which *Z* was a member. Indeed, there were no two individual organisms, α and β, in *Z's* generation (or later) such that selection for <A,B,c> made the probability that α would have the <A,B,C> sequence greater than the probability that β would have it. So selection for <A,B,c> did not make it more probable *that Z would have* the <A,B,C> sequence *rather than* some other organism in *Z's* generation (or a later generation). And, if it is supposed to be selection's increasing the probability of an occurrence that implicates selection in the causation of that occurrence (as per Neander's argument), selection for <A,B,c> was not a cause *of Z's having* <A,B,C>, since selection did not in fact increase the probability *that Z would have* <A,B,C>. An explanation of how *Z* came to have <A,B,C> need thus cite only the fact that *Z's* parent had <A,B,c> and that there was mutation of *c* to *C* in the production of *Z*. (Neander would argue that selection was a cause of *Z's* parent's having <A,B,c>; but my counterargument, *mutatis mutandis*, would again apply.) If selection was not a cause of *Z's* having <A,B,C>, however, it was not among the historical causes of *S's* having *T*. Selection would be a cause of *S's* possessing *T* only if it not only increased the probability that <A,B,C> would occur, but also made it *more probable that an ancestor than a non-ancestor* of *S* would possess <A,B,C>; for, had <A,B,C> originated outside of *S's* ancestry, *S* would not have inherited that sequence and, hence, would not possess *T*. So selection for <A,B,c> merely produced a *population level* increase in the probability that <A,B,C> would occur; it did not differentially increase the probability that any particular *individual organism* in that population would have the <A,B,C> sequence. Thus, selection is not a cause of the presence of a trait in any of its bearers and, hence, does not help explain the presence of a trait

in its bearers. (See Sober (1995) for further related criticisms of Neander's position.)

Of course, the argument to this point has presupposed that the purpose of function ascriptions is to explain the presence of a *trait token in its bearer*, rather than to explain the presence of a *trait type in a species* or population. But the literature is ambiguous on this issue. Neander (1991b) discusses function ascriptions that explain the presence of "the switch on the wall" (p. 460) and "*your* opposable thumb" (pp. 460–461; emphasis added), which are clearly cases of explaining the presence of tokens in their bearers. Others explicitly take function ascriptions to explain the presence of a trait type in a species or population (Millikan 1989a: p. 175; Griffiths 1993: pp. 414–415; Allen & Bekoff 1995: p. 612). And Godfrey-Smith (1993) discusses explaining why "the thing in question is there" (p. 198) and why "that particular part is there" (p. 201), where both locutions could be interpreted in either way. One might, then, disambiguate and insist that function ascriptions explain not the presence of a token of a trait in its bearer, but the presence of a type of trait in a species. Then the explanatory job is explaining why the trait has reached the degree of frequency in a species it has, and this is an explanatory job in which selection *is* involved. Consequently, one could argue, the argument above fails to show that selection is not a necessary ingredient in explaining "the presence" of a functionally characterized trait, since it is the presence of the trait *in a species* that is explained by function ascriptions.

But there is something paradoxical about claiming that the etiological theory's function ascriptions explain the presence of a trait type in a species, rather than trait tokens in their bearers. For etiological definitions of *function* are formulated for *current tokens* of a trait. And there is good reason for this focus. The core insight of the etiological theory, which enables it to avoid any kind of supernatural final causation, is that an item's presence is explained not by its *own* effect, but by the fact that past items of the same type produced the same type of effect and that these effects of past items were among the causes of the existence of the current item. If *function* were defined for a trait *type*, without distinguishing past tokens of that trait type from current tokens, the definition would assign functions to the original token(s) of that trait type, which had no ancestors, and would thus explain their presence in terms of their own effects. So the core insight of the etiological theory requires defining *function* for a current token of a trait. To insist that etiological function ascriptions explain the presence of a trait in a species is thus to cleave the etiological definition of *function* (explaining what it means for *my heart* to have the function of pumping blood) from its putative explanatory role (at the level of a trait type in a species). If etiological function ascriptions, defined for token traits in their bearers, are to explain the

presence of functionally characterized items, that explanation must consist in accounting for the presence of the functionally characterized trait token in its bearer. And, as shown above, this is a form of explanation in which selection does not figure.

Thus, although the weak theory employs Cummins-style functional analyses in its function assignments, the standard objections to Cummins' theory do not apply to the weak theory. Indeed, the weak theory satisfies all the desiderata for an etiological theory of functions set by the standard objections to Cummins' theory: The weak theory is not overly liberal in its function assignments, the weak theory assigns functions to malfunctioning items, and its function ascriptions explain the presence of the functionally characterized traits in their bearers. So the strong theory's appeal to selection for a trait is not a necessary ingredient of a thoroughly etiological theory. In addition, the weak theory provides a unity to the concept of *function* lacking under the strong theory. For the strong theory, with its requirement of selection for a trait, does not assign functions to many physiological and developmental traits for which there has not been direct selection, instead relegating them to subsumption under Cummins' theory. But physiological and developmental features of interest to biologists are hereditary features of organisms in lineages that have undergone evolution. And, since the weak theory assigns functions to all hereditary traits that produce effects appearing in a Cummins-style functional analysis of the fitness of the ancestral bearers of those traits, the weak theory generates, in a principled way, function assignments to all such physiological and developmental traits. The weak theory thus provides a unified account of biological functions, without separating uses of *function* in evolutionary theory and behavioral ecology (which have been the privileged domain of the strong theory) from those in physiology and developmental biology (which have been relegated to Cummins' theory).

4. Conclusion

Kitcher (1993) has proposed a non-etiological theory of functions that boasts a similar unified account of biological usage across the range from behavioral ecology to physiology. In concluding, I would like to compare Kitcher's theory with what I have called the "weak theory", by way of partially indicating the weak theory's position within the logical geography of the functions literature beyond the borders of the etiological theory.

Kitcher defends a design-based theory of functions, according to which "the function of S is what S is designed to do" (1993: p. 380). Kitcher takes design, in the primary case, to be a feature of organisms and takes natural selection to be the source of that design (pp. 380–381). In Kitcher's view,

natural selection has shaped organisms that are successfully responsive to the environmental demands to which they are subjected. The selection pressures acting on organisms *as wholes* generate "a hierarchy of ever more specific selection pressures" acting on the constituent features of organisms – their anatomy, physiology, morphology, and behavior (p. 390). Such constituent features themselves exhibit design, according to Kitcher, when they are responsive to their more specific selection pressures and thereby contribute to an organism's overall response to environmental demands. That is, features "of organisms have functions in virtue of their making a causal contribution to responses to those [selection] pressures", where the notion of a "causal contribution" is to be understood in exactly Cummins' sense (p. 390). In short, then, for Kitcher a trait T in an organism O has the function of producing some effect E just in case T has the Cummins function of producing E within $O's$ responsiveness to environmental demands.

Like the weak theory, Kitcher's theory emphasizes an item's Cummins contribution to its bearer, rather than direct selection for that item, and does not make selection for a trait a necessary condition of its being assigned a function. Indeed, since both the weak theory and Kitcher's theory rely on a Cummins-style functional analysis (of fitness and responsiveness to environmental demands respectively) in order to identify the function of a trait, they would agree in the majority of their function assignments. But the two theories diverge in their function ascriptions with respect to some core cases; and the divergences stem from two key differences between the theories.

First, the weak theory functionally analyzes the fitness of *ancestral* organisms – identifying the function of a current token of a trait in terms of the role that past tokens played within ancestral fitness –, whereas Kitcher's theory permits functional analyses of the responsiveness of *current* organisms to their environmental demands. This makes Kitcher's theory subject to the standard "malfunction objection." That is, since diseased kidneys, for example, do not contribute to their bearers' responses to environmental demands, Kitcher's theory would withhold assigning to a diseased kidney the function of filtering metabolic wastes from the blood. So Kitcher's theory fails to capture the normativity of function ascriptions. Second, the weak theory assigns functions only to *hereditary* traits, whereas Kitcher's theory does not have the requirement of hereditariness. This means that Kitcher's theory is threatened with "liberality from within." Kitcher's theory, for example, cannot rule out assigning to the weight of Williams' flying fish the function of returning the fish to water, since the weight makes a causal contribution to the fish's responsiveness to the environmental demand of returning to water.

Finally, the fact that Kitcher's theory neither restricts functional analyses to ancestral organisms nor requires hereditariness means that the function ascriptions generated by Kitcher's theory do not explain the presence of the functionally characterized items in their bearers (see Godfrey-Smith 1993: pp. 202–205, for an extended argument to this effect). Consider a polyandrous population in which male sperm are monomorphic and function only to fertilize eggs. Suppose that a mutation occurs in generation *X* of this population that produces a second type of sperm in its bearer *S*; these are "killer" sperm, which do not swim toward and attempt to fertilize eggs, but which instead kill the sperm of other males encountered in a female reproductive tract.[5] *S* would thus possess a great advantage in reproductive competition and would, consequently, have an adaptive advantage vis-à-vis the other males of generation *X*. Since the killer sperm obviously contribute to *S's* responsiveness to an environmental demand (competition with other males for paternity), Kitcher's theory would entail that the killer sperm have the *function* in *S* of killing other males' sperm. But the fact that the killer sperm do that does not in any way help to explain the fact that they are present *in S*. In fact, the only reason they are present in *S* is just plain *luck* of the mutation draw. They are *not*, however, present in *S's* male offspring due to luck; they are present in *S's* male offspring because they contributed to *S's* reproductive success and they were hereditary. So Kitcher's theory cannot distinguish lucky cases (like that of *S*) from cases (like those of *S's* male offspring) in which the functionally characterized trait is present in its bearer because it contributed to ancestral responsiveness to environmental demands and was, because of that, passed from parent(s) to offspring; for Kitcher's theory generates function ascriptions that cut across both types of case.

These limitations of Kitcher's theory, of course, stem from the three standard etiological objections to Cummins' theory, and they are limitations faced by any theory that defines *function* in terms of a contribution to fitness or to responsiveness to environmental demands. It has been the great virtue of etiological theories that they have overcome these limitations. But the strong theory's appeal to selection for a trait, by way of overcoming these limitations, forces an unnecessary dichotomy between senses of *function* in biology. The appeal to selection for a trait in defining *function* has thus hindered, rather than facilitated, the development of a fully satisfactory theory of biological functions. In contrast, the weak theory, which does not appeal to selection for a trait, is able to provide the sort of unity that is the virtue of Kitcher's theory while simultaneously possessing all the hallmark characteristics of the etiological theory. Thus, the weak theory provides everything that etiological theorists have ever wanted from a theory of functions – and a little more: the possibility of consensus *with* unity.

Acknowledgements

Notes

[1] Most proponents of the etiological theory take this point to be the recent, as opposed to the remote, evolutionary past of the population possessing T (see, e.g., Millikan 1989a, 1993; Griffiths 1993; Godfrey-Smith 1994).

[2] This idea of incorporating Cummins' theory into the etiological theory is not original. Griffiths proposes a sort of synthesis of Cummins' theory with the etiological theory: "The proper functions of a biological trait are the functions it is ascribed in a [Cummins-style] functional analysis of the capacity to survive and reproduce (fitness) which has been displayed by animals with that feature" (1993: p. 412). Griffiths' theory, however, is explicitly a version of the strong etiological theory: "The proper functions of a trait . . . are the effects in virtue of which the trait was selected" (p. 412). But, as this section will show, the two sentences quoted are not equivalent and the weak theory opts only for the first. (Kitcher also proposes a theory of functions that incorporates Cummins' theory, which I will discuss in section 4.)

[3] The weak theory also captures a far broader range of biological *usage* of the concept of *function*. For example, Gould and Lewontin (1979) argue that there are many traits for which there has been no selection, but to which the concept of *function* nonetheless applies (pp. 585 & 590–593). In addition, Williams (1966) contends that the "prime question" of "teleonomy", the study of the functions of traits, would be "What is its function?" and that any "plausible demonstration of design in relation to a goal would provide the answer to the teleonomist's prime question" (pp. 258–259). Only after answering the prime question would the teleonomist inquire as to what maintains the trait with that function in current populations and make an "initial attempt . . . to explain the mechanism as the inevitable consequence of the natural selection of alternative alleles" (p. 259; see also p. 260). Also, Endler (1986) argues that unless we know the functions of physiological and biochemical traits, we cannot determine whether there has been selection for those traits (p. 158). Indeed, according to Endler, knowing the functions of alternative traits will enable us to know which alternative is functionally superior and thus to predict whether selection will occur (p. 89). None of these uses of the concept of *function* take selection for a trait to be a necessary condition for its having a function; they are thus compatible only with the weak theory.

[4] This paragraph and the next are adapted from Buller (1997).

[5] This type of sperm dimorphism is an oversimplification for the purpose of example. Baker and Bellis (1995: chap. 11) catalog three functional types of sperm in *humans* (and that still leaves several morphological types of human sperm with undetermined functions). In addition to "egg-getters", whose function is to fertilize the ovum (and which make up only about 1 in 200 of the sperm in an ejaculate [see chap. 12]), there are "blockers" and "seek-and-destroy" sperm. Blockers enter the channels of the cervical mucus and stop, thereby barring passage of any other sperm. Seek-and-destroy sperm, on which the "killer" sperm of the example are modeled, kill the sperm of other males through two primary mechanisms: head-to-head combat, in which two sperm become attached at the head and one injects a lethal enzyme into the other, and what is in effect an enzyme bomb, where a sperm releases a lethal enzyme in the vicinity of the sperm of another male. Sperm polymorphism is also well documented in several other species (see chap. 11). In Lepidoptera, for example, apyrene sperm completely lack a genome, yet are functional, serving only to engage in sperm warfare (see Sivinski 1984).

References

Allen, C. and Bekoff, M.: 1995, 'Biological Function, Adaptation, and Natural Design', *Philosophy of Science* **62**, 609–622.

Amundson, R. and Lauder, G.V.: 1994, 'Function Without Purpose: The Uses of Causal Role Function in Evolutionary Biology', *Biology and Philosophy* **9**, 443–469.

Baker, R.R. and Bellis, M.A.: 1995, *Human Sperm Competition: Copulation, Masturbation, and Infidelity*, Chapman & Hall, New York.

Brandon, R.N.: 1990, *Adaptation and Environment*, Princeton University Press, Princeton.

Buller, D.J.: 1997, 'Individualism and Evolutionary Psychology (or: In Defense of "Narrow" Functions)', *Philosophy of Science* **64**, 74–95.

Crow, J.F.: 1979, 'Genes that Violate Mendel's Rules', *Scientific American* **240**(2), 134–146.

Cummins, R.: 1975, 'Functional Analysis', *Journal of Philosophy* **72**, 741–765.

Endler, J.A.: 1986, *Natural Selection in the Wild*, Princeton University Press, Princeton.

Godfrey-Smith, P.: 1993, 'Functions: Consensus Without Unity', *Pacific Philosophical Quarterly* **74**, 196–208.

Godfrey-Smith, P.: 1994, 'A Modern History Theory of Functions', *Noûs* **28**, 344–362.

Gould, S.J. and Lewontin, R.C.: 1979, 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme', *Proceedings of the Royal Society of London* **B 205**, 581–598.

Griffiths, P.E.: 1993, 'Functional Analysis and Proper Function', *British Journal for the Philosophy of Science* **44**, 409–422.

Kitcher, P.: 1993, 'Function and Design', *Midwest Studies in Philosophy* **18**, 379–397.

Lyttle, T.W.: 1993, 'Cheaters Sometimes Prosper: Distortion of Mendelian Segregation by Meiotic Drive', *Trends in Genetics* **9**, 205–210.

Maynard Smith, J.: 1993, *The Theory of Evolution* (Canto ed.), Cambridge University Press, Cambridge, UK.

Mayr, E.: 1988, 'An Analysis of the Concept of Natural Selection', in *Toward a New Philosophy of Biology*, Harvard University Press, Cambridge, MA, pp. 95–115.

Millikan, R.G.: 1984, *Language, Thought, and Other Biological Categories*, MIT Press, Cambridge, MA.

Millikan, R.G.: 1989a, 'An Ambiguity in the Notion "Function"', *Biology and Philosophy* **4**, 172–176.

Millikan, R.G.: 1989b, 'In Defense of Proper Functions', *Philosophy of Science* **56**, 288–302.

Millikan, R.G.: 1993, 'Propensities, Exaptations, and the Brain', in *White Queen Psychology and Other Essays for Alice*, MIT Press, Cambridge, MA, pp. 31–50.

Mitchell, S.D.: 1995, 'Function, Fitness and Disposition', *Biology and Philosophy* **10**, 39–54.

Neander, K.: 1991a, 'Functions as Selected Effects: The Conceptual Analyst's Defense', *Philosophy of Science* **58**, 168–184.

Neander, K.: 1991b, 'The Teleological Notion of "Function"', *Australasian Journal of Philosophy* **69**, 454–468.

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

Sivinski, J.: 1984, 'Sperm in Competition', in R.L. Smith (ed.), *Sperm Competition and the Evolution of Animal Mating Systems*, Academic Press, New York, pp. 85–115.

Sober, E.: 1984, *The Nature of Selection*, MIT Press, Cambridge, MA.

Sober, E.: 1995, 'Natural Selection and Distributive Explanation: A Reply to Neander', *British Journal for the Philosophy of Science* **46**, 384–397.

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

Williams, G.C.: 1966, *Adaptation and Natural Selection*, Princeton University Press, Princeton.