Brit. J. Phil. Sci. **69** (2018), 301–327

Evolution, Dysfunction, and Disease: A Reappraisal

Paul E. Griffiths and John Matthewson

ABSTRACT

Some 'naturalist' accounts of disease employ a biostatistical account of dysfunction, whilst others use a 'selected effect' account. Several recent authors have argued that the biostatistical account offers the best hope for a naturalist account of disease. We show that the selected effect account survives the criticisms levelled by these authors relatively unscathed, and has significant advantages over the BST. Moreover, unlike the BST, it has a strong theoretical rationale and can provide substantive reasons to decide difficult cases. This is illustrated by showing how life-history theory clarifies the status of so-called diseases of old age. The selected effect account of function deserves a more prominent place in the philosophy of medicine than it currently occupies.

1 Introduction

A core issue in the philosophy of medicine is whether non-evaluative facts play the key role in determining whether a state is pathological (Boorse [1977]; Wakefield [1992]; Kingma [2007]; Murphy [2009]). Call this the question of

doi:10.1093/bjps/axw021

Advance Access published on October 8, 2016

'naturalism' regarding disease (Kingma [2014]). Many authors on both sides of this question agree that some version of the so-called 'biostatistical theory' (BST) is the best naturalist account of disease (Schwartz [2007]; Kingma [2007]; Hausman [2012]; Boorse [1977], [2014]). We contend that this is a mistake. The BST is subject to some extremely damning criticisms, and so linking the fortunes of the naturalist project to the BST unnecessarily weakens that project.

Instead, we contend that a 'selected effect' view of biological function and dysfunction will be a component of the best naturalist view of disease (Wakefield [1992], [2007]). After describing the two accounts in Section 2, we argue for this conclusion along three lines. In Section 3, we defend the selected effect view against criticisms old and new. The selected effect view is very far from being 'dead in the water' as many authors appear to believe. In Section 4, we show how the selected effect view avoids fundamental problems that have dogged the BST since it was first proposed. Finally, in Sections 5 and 6, we show that the selected effect approach can do more than recover intuitive judgements about dysfunction. It can advance our understanding, providing a principled basis for demarcation decisions and generating new insights into the very idea of pathology.

Underlying our arguments is the idea that explication (Carnap [1950]), as opposed to mere analysis, is the proper goal of philosophy in this context.[1] It has proved extremely difficult to give an analysis that captures all cases of disease and only cases of disease, and still manages to say something substantive and informative about disease. In our view, the first two of these desiderata should not stand in the way of meeting the third. The aim of philosophical analysis should not be restricted to describing how disease is understood by ordinary folk, or by doctors, or by philosophers of medicine. It can and should have the more substantial aim of advancing our understanding of what it is for an organism to be in a normal or a pathological state. Ruth Millikan ([1989]) has advocated this view of philosophical analysis in her work on the concept of biological function more generally.

Towards the end of the article we show how evolutionary biology provides substantive reasons to consider some phenotypes as pathological and others as normal (see also Nesse [2001], [2007]). An evolutionary view of function can adjudicate difficult cases in a principled and independently motivated way. It therefore has the potential to advance our understanding of the subject, rather than merely recovering pre-existing intuitions. So this article includes both methodological and substantive elements: we advocate a particular approach

[1] Lemoine ([2013]) and Schwartz ([2014]) hold similarly pessimistic views of the effectiveness of conceptual anlysis in settling debates regarding the nature of disease.

to philosophical work, and we defend a particular application of this approach to dysfunction in the context of medicine.

Before we begin, it will be helpful to clarify the scope of the article. First, we do not offer an overall account of the concept of disease. Our focus is on the non-evaluative facts that will form at least part of an overall account of disease. The dominant view in the literature is that these are facts about biological function and dysfunction. So our goal is to defend a particular account of function and dysfunction—the selected effect account—in the context of medicine.

Second, nothing we say in this article should be taken to suggest there is a single correct account of biological function. Both authors are pluralists: we think there are a number of legitimate notions of function at play in the biological sciences, each with advantages in different contexts (Griffiths [1993], [2009]; Godfrey-Smith [1993]). Perhaps more than one account will be needed, even in medicine alone. The question is not, 'What is the correct account of biological function?', but rather, 'Which account(s) of biological function do a good job of demarcating pathological states?'.

Third, we aim to convince naturalists to seriously consider the selected effect account of function. We do not argue that those who reject naturalism must adopt a selected effects account! However, our arguments are relevant to anti-naturalists because, as we show below, many of them take a refutation of the biostatistical view of function to be *ipso facto* a refutation of naturalism.

Finally, we note that psychiatric disorders pose special challenges for any analysis of the normal and the pathological, so while we do not distinguish the psychiatric from the somatic domain in this article, we remain open to the possibility that they may require separate treatments.

With the throat clearing out of the way, we begin by outlining the two relevant accounts of biological function.

2 Biostatistical and Selected Effect Accounts of Function

The standard view of biological function and dysfunction in the setting of health and disease is that employed in Christopher Boorse's ([1977]) BST. In spite of being almost four decades old, this view is still widely discussed and hotly contested in the philosophical literature.[2]

Boorse ([1977], pp. 555, 567) gives the following account of disease:

(1) The reference class is a natural class of organisms of uniform functional design; specifically, an age group of a sex of a species.

[2] A selection of important recent philosophical discussions include (Cooper [2002]; Kingma [2007], [2010]; Hausman [2011]; Garson and Piccinini [2014]).

(2) A normal function of a part or process within members of the reference class is a statistically typical contribution by it to their individual survival and reproduction.

(3) A disease is a type of internal state that is either an impairment of normal functional ability—that is, reduction of one or more functional abilities below typical efficiency or a limitation on functional ability caused by environmental agents.

(4) Health is the absence of disease.

Our focus is on Boorse's definition of function, given in clause 2. Function is the contribution made to future survival and reproduction, and 'normal' functioning is simply statistical normality within a particular reference class. If a trait contributes to fitness in a manner typical of members of the reference class (or better), it is considered to be functioning correctly.

An alternative approach to biological function is the selected effect account (Papineau [1987]; Millikan [1984], [1989]; Neander [1991]; Griffiths [1993]; Godfrey-Smith [1994]). Details vary, but the core commitment of the selected effect account is simple: 'biological proper functions are effects for which traits were selected by natural selection' (Neander [1991], p. 168).

Any capacity of a system that is of interest to a scientist can legitimately be subjected to functional analysis (Cummins [1975]). For example, Jared Diamond ([1997]) has analysed plants and animals in terms of their capacity to contribute to colonial expansion. The selected effect theory of biological function rests on the insight that one biological capacity, and the functions it generates, stands out from all the others. It is no mere quirk of psychology or sociology that leads the selected effect view—and biology generally—to focus on the capacity to survive and reproduce (see also Garson [2013]). According to the theory of natural selection, many parts and processes that we observe in living organisms exist because they contributed to survival and reproduction in ancestral organisms. Selected effect functions explain why organisms have the parts and processes that have those functions. They are often referred to as the 'proper functions' of those parts and processes (Millikan [1984]; Neander [1991]).

However, the selected effect account does not suggest that we look to the distant evolutionary past to discover the function of a trait. The current orthodoxy in philosophy of biology is that selected effect function derives from the 'modern history' of a trait: the evolutionary forces that have maintained the trait in a lineage under recent selection (Griffiths [1993]; Godfrey-Smith [1994]).

Karen Neander was a notable and early defender of the selected effect view in the context of medicine (Neander [1998]), but the most prominent advocate is Jerome Wakefield ([1992], [2000], [2007]). His 'harmful dysfunction' account

of medical disorder incorporates both evaluative and non-evaluative elements. It requires both that biological dysfunction occurs and that this dysfunction is considered harmful. In this article, we are not concerned with Wakefield's overall account of medical disorders, only with its non-evaluative portion. For Wakefield, biological dysfunction occurs when a structure fails to produce the effect for which it has been selected, so he subscribes to the selected effect theory. In the philosophy of medicine, Wakefield's account of dysfunction is frequently sidelined in favour of the BST account. In the next two sections, we argue that this is a mistake.

3 Objections to the Selected Effect Account

The view that the most viable naturalist account of dysfunction is the biostatistical account rests on arguments against selected effect accounts in general, or against Wakefield's version in particular. In this section, we show that these arguments are much less effective than their proponents suppose. Several counterexamples are thought to show that the selected effect account generates counterintuitive function ascriptions. Some of these counterexamples rely on an inadequate understanding of how natural selection works and evaporate once one considers a serious, contemporary attempt to give an evolutionary explanation of the relevant phenotype. Some target older versions of the selected effect view that its current adherents would not accept. And some raise good *prima facie* objections, but fail to recognize the replies available to their opponent.

3.1 Boorse

Boorse himself takes a superficial view of evolutionary theory when he states in his original paper that adaptedness cannot be an analysis of health, since 'parents hardly become healthier with each successive child, nor would anyone maintain that the healthiest traits are the ones that promote large families' (Boorse [1977], p. 548). However, ever since David Lack ([1947]) developed the concept of optimal clutch size, it has been understood that larger families do not necessarily mean higher fitness. They may lead to less robust offspring and lower rates of survival to reproductive maturity. Reproduction must also be traded against survival, as we discuss at length in Section 6. The selected effect view is not committed to the idea that a big family is equivalent to health. Perhaps this argument is intended to be light-hearted, but comments such as this undoubtedly contributed to the premature rejection of evolutionary views of dysfunction.

In his early writings, Boorse was attacking an early and inadequate version of the selected effects theory (Wright [1973]; Boorse [1976]), but selected effect

accounts have moved on considerably since then, as have Boorse's criticisms (Boorse [2002]). Some of these newer criticisms target the viability of selected effects as a general account of all function attributions, and so are not of concern here, as explained above. Some, however, target selected effect accounts in the medical context. For example, Boorse ([2002], pp. 91–2) claims that the selected effect account is unable to distinguish functional normality of different life-stages from one another. The mental abilities of an eight-year-old are less developed than those of a fourteen-year-old, but since the selected effect approach treats these individuals as members of the same evolutionary lineage, Boorse claims it cannot distinguish the two. But life history theory, a central aspect of modern evolutionary theory, is dedicated to making exactly such distinctions, and doing so with mathematical precision. There is no problem here. We discuss life history theory at length in Section 6.

Boorse also suggests that there is a problem with transitions in selective regimes. He uses the example of woodland in which the trees have been blackened by industrial pollution. There are some light-coloured moths, adapted for camouflage against the natural colour of the trees, and some darker moths, a recent adaptation for camouflage against the blackened bark. Boorse argues that the selected effect theory has no principled way to specify how much time must elapse before the function of light pigmentation ceases to be camouflage ([2002], p. 100).[3] The same phenomenon can occur spatially—in invasive cane toads (*Rhinella marina*) in Australia, large toads with lowered immune function and rapid, directional hopping are selectively favoured at the invasion frontier but disfavoured in settled populations behind the frontier (Phillips *et al.* [2010]). How far back do we need to go behind the frontier before suppressed immune function becomes dysfunctional?[4] The underlying issue is that selected effects functions can change gradually over evolutionary time and ecological space. But this is how biology works. It would only be problematic if it were a conceptual truth that diseases must have sharp boundaries, which does not seem to us to be at all evident. On the contrary, many sharp boundaries in medicine, such as body mass index cut-offs for obesity, or the fasting blood-glucose cut-offs for pre-diabetes and diabetes, are imposed on what are in reality continuous differences for practical, often administrative, reasons. Moreover, the fact that boundaries are not sharp does not mean they are unprincipled. Evolutionary biology is used to handling the vague boundaries that are inevitable in a Darwinian world in disciplined and rigorous ways. We return to this issue in Section 6.

[3] Boorse also discusses the problem of fixing a normal range of function for a trait, something dealt with at greater length in (Schwartz [2007]), discussed below.

[4] For an approach to this issue using expected mutation rates, see (Griffiths [1993]).

3.2 Kingma

Elselijn Kingma is a prominent critic of the biostatistical view of dysfunction (for example, Kingma [2007], [2010], [2016]). In her view, the BST is 'the best and only presently existing naturalistic account' of dysfunction and of disease ([2010], p. 262). Presumably, she excludes Wakefield's harmful dysfunction account in this statement because it is not purely naturalistic. However, in a recent book chapter on psychiatric disease, Kingma ([2013]) explicitly compares Boorse and Wakefield's accounts and argues that the BST is superior to the selected effect account of dysfunction.

Kingma targets the more recent versions of the selected effects account (Millikan [1989]; Neander [1991]). Her first concern is that the evolutionary history of a trait is difficult to confirm, and hence these accounts render function and dysfunction unknowable: 'we will never be in a position to access all those facts; we are lucky to have access to any' (Kingma [2013], p. 374). Her specific target in this chapter is psychiatry, but at least as far as somatic disease is concerned, this seems overly pessimistic, especially since the orthodox version of the selected effects account is now the 'modern history' version of Godfrey-Smith ([1994]). This account clearly distinguishes between evolutionary explanations of the origin of traits and evolutionary explanations of the maintenance of traits. For a great number of medically relevant traits, it is transparent how fitness would be reduced if they failed to perform their proper functions, and equally transparent that those fitness benefits have played a role in maintaining the trait in the recent past. It is clear that Northern Europeans lost much of their skin pigmentation to meet the challenge of synthesising vitamin D in that dark and benighted region (Jablonski and Chaplin [2000]). This selection pressure is still operative (Glass *et al.* [2009]).

Another of Kingma's objections to the selected effects account is that there are 'selected disorders'. These are traits that have been selected but which constitute disorders in modern circumstances. Kingma's examples are rape, being excessively violent, and attention-seeking behaviour. However, in each of these examples, either the traits are not considered diseases or they are paradigmatically controversial cases—cases where we would like our account of pathology to adjudicate, as there are no clear shared intuitions to guide us. For example, rape is not considered an illness: in the absence of some independently diagnosed mental disorder, rapists go to jail, not to hospital. Conversely, it is controversial whether undesirable personality traits such as attention-seeking are pathologies as opposed to normal but inconvenient human variation. So it is not a devastating criticism of Wakefield that he leaves it to future empirical research to determine whether these so-called personality disorders are really disorders.

Tim Lewens ([2015], p. 188) has presented a similar objection in counterfactual form. If rape were an evolved facultative response to low social status in males, then it would follow that males who have low social status but do not rape would be dysfunctional. This is intended to be a counterintuitive result. But Lewens is conflating two questions. First, does it seem counterintuitive now, as thing stand, to say that low status men who do not rape are dysfunctional? Yes, this is clearly counterintuitive. Second, if we had firm scientific grounds for believing that it is a monomorphic feature of the male brain (as opposed to a polymorphism, found in some men but not others) to have mechanisms designed to produce rape behaviour in response to low status, would it be counterintuitive to say that brains lacking these mechanisms (or with damaged mechanisms) were dysfunctional? No, it would not be counterintuitive to say dysfunction is present in that counterfactual scenario. It is a stipulation of the thought experiment that males are designed to exhibit rape behaviour in certain situations, and that some males cannot act as they are designed to act. Nevertheless, a propensity to rape—even if stipulated to be part of the design of human males—would be an immensely unfortunate aspect of human nature, and it would still be deplorable to act on it. The rhetorical power of this argument is that it makes the defender of SE function look 'soft on rape'. But this type of argument can be used to make any definition of dysfunction appear 'soft on rape'. For example, normativists are committed to the view that if everyone approved of rape, then not raping would be dysfunctional. Lewens defends this counterfactual mode of argument by saying that we should not make 'pathology hostage to evolutionary enquiry' ([2015], p. 188). In Section 6, we will argue the exact opposite.

Kingma's next criticism concerns traits that haven't been selected for, but are merely side effects of natural selection.[5] Reading is probably too recent in our lineage to have a selective history. If it is merely a by-product of some other cognitive abilities, then it has no proper function. This means that, according to the selected effect account, impaired ability to learn to read cannot be a dysfunction. Yet, Kingma claims, we see impaired ability to read as a disorder. Wakefield ([2000]) has previously argued that we consider such impairments to be disorders because (or at least when) we presume the impairment is explained by some underlying dysfunction in a cognitive trait that does have a proper function.[6] Kingma rejects this, arguing that something

[5] This distinction between 'selection of' and 'selection for' phenotypes is introduced in (Sober [1984]), and elaborated and defended in (Goode and Griffiths [1995]); Wilkins and Griffiths ([2013]).

[6] The example was actually introduced by Wakefield ([2000]), intended to be a case in his favour. Murphy and Woolfolk criticized his proposal along similar lines to Kingma (Murphy and Woolfolk [2000b]).

can go wrong with a side effect without anything going wrong with the traits of which it is a side effect.

However, it is very difficult to find plausible examples of this phenomenon, and reading is not one of them. Most people who cannot read have not had the opportunity to learn, so these cases are irrelevant. The relevant cases are people who have every opportunity to learn to read but are unable to do so or who require special education to do so. Kingma's proposal is that some of those people have no underlying deficit in any trait that has been selected, such as visuospatial recognition or facility with symbolic representation more generally. But in that case, it is a mystery why they can't learn to read in the normal way. Reading is not unique in this respect. The ability to acquire evolutionarily novel, culturally transmitted traits has been so critical in human evolution that it is very unlikely that substantial impairments to this ability are part of the normal variation in human populations. Minor impairments may well be part of normal human variation, but by the same token, they do not elicit the intuition about dysfunction.

Kingma finishes with a more general criticism: the selected effects account unduly restricts the set of traits that might exhibit pathology. Kingma claims that 'most if not all of our physical and the vast majority of our mental traits fall within the domain of health and disorder' (Kingma [2013], p. 379). She claims that this is a core commitment of the concept of a disorder, and that a restriction to selected traits fails to capture this core commitment.

But how large is the domain that the selected effects account excludes? Millikan ([1984]) has shown with admirable rigour that if our basic representational capacities have been selected, then all mental representations have proper functions derived from the functions of the system that produces them. The same is true for other physiological systems. No one in the past had ever run as fast as Usain Bolt, but the activities of his muscles as he runs the 100 m still have proper functions. The kinds of case Kingma has in mind certainly exist—the selected effects account cannot judge whether a mutation in a stretch of genuinely 'junk' DNA is dysfunctional—but for the most part, these cases do not generate strong intuitions about function and dysfunction.

Kingma's arguments are certainly not without merit, but they do not leave the selected effects account of dysfunction dead in the water and unworthy of further attention. At best, they establish that the account is unintuitive in rare and/or peripheral cases, and we will argue below that this level of disagreement with intuition is acceptable.

3.3 Hausman

Daniel Hausman's work has also focussed on the BST, although his goal is to update and defend it. Hausman ([2012], p. 520) explicitly rejects selected

effects accounts of functions, claiming that the BST is 'the best-developed naturalistic view of health'. Hausman's ([2012], p. 522) primary argument for this appears in a footnote:

> Suppose, e.g. (as might in fact be the case), that some of the many essential functions that the liver carries out were not selected for, but were side effects of other things that were selected for. There would then be no functional explanation for why the liver carries out these activities: the fact that the liver makes these contributions would not be a causal consequence of livers having done these things in the past [. . .] Yet these contributions of the liver would still be among its functions, and etiological [= selected effect] theories would mistakenly deny this.

There are two ways we might understand this claim. One is biologically implausible and the other is not relevant to the current version of the selected effects account.

Hausman suggests that the liver might have some effects that are 'essential'—perhaps the organism will die or do very poorly without these effects—and yet there has been no selection for these effects, as they were side effects of other useful functions of the liver. But this misunderstands what it means to be an evolutionary side effect. If a trait enhances fitness in several different ways, the strength of selection acting on that trait will be greater than if the trait only enhanced fitness in a subset of those ways. This will show up in the evolutionary dynamics, and so all these effects will be selected effects and have proper functions.

Hausman's example is more plausible if we interpret it as pointing out that some of the essential effects that the liver now performs were not the functions that led to the original evolution of livers. However, interpreted this way, the argument targets a straw man. Current versions of selected effect accounts of function appeal to the selective forces that maintain traits in a population, not to the selective forces involved in the origin of traits.

Hausman also comments that selected effects theories cannot characterize 'partial dysfunction' (parts that work but do not work very well). This is a complex matter, and selected effects theories in philosophy certainly have not given it enough attention, but there is no reason to suppose that it cannot be done. Evolutionary theory does not imply that every organism in a population except the fittest organism in that population is dysfunctional with respect to some trait. We cannot run as fast as Usain Bolt, but that does not mean our leg muscles are dysfunctional. Amongst many other reasons, this is because the selective problem facing an organism is not how to be 'best', but how to make the best allocation of available resources to many different traits, given the available information about its circumstances. We return to the problem of 'partial dysfunction' in Section 4.

3.4 Murphy and Woolfolk

Dominic Murphy and Robert Woolfolk directly criticize Wakefield's views, rather than compare them with the BST. Like Wakefield, their primary concern is psychiatric disorders, but the points they make apply more widely. Murphy and Woolfolk recognize that in the current version of the selected effects account, functions arise from the selective maintenance of traits, rather than their selective origin (Murphy and Woolfolk [2000a]). Nevertheless, they believe that problems arise when we consider either spandrels or vestigial traits—either structures that are side effects of selected traits, or structures that are no longer selected but are still extant. An example of the former is the human chin, which apparently develops only as a side effect of the growth of other parts of the jaw. An example of the latter is the human appendix, a reduced form of the caecum that houses cellulose-digesting bacteria in herbivores. Since neither kind of structure has effects for which it was recently selected, neither has selected effect functions. Nevertheless, we can have disorders of these structures. Therefore, Murphy and Woolfolk argue, disorder does not require this type of dysfunction.

Wakefield ([2000]) gives a similar reply to both cases. In disorders that arise from these structures, there is dysfunction present, just not where Murphy and Woolfolk look for it. In the case of the spandrels, given that these structures arise as side effects of other selected traits, there will have been some dysfunction in those other traits. For example, the hereditary prognathic chin of the Habsburg imperial family made it difficult for them to chew. One Emperor supposedly starved as a result. The ability to masticate food has, obviously, been selected and the other parts of the jaw of which the chin is a side effect are designed to facilitate this. In cases of disordered vestigial traits, Wakefield argues that the organ itself isn't dysfunctional, since it doesn't have a job to do, but some of the tissues that make up the organ are dysfunctional. Even if the appendix has no selected function, for example, the tissues of the appendix are severely inflamed (and so dysfunctional) in cases of appendicitis.

We are sympathetic to these replies. However, there does seem something rather *ad hoc* about them. If Wakefield can cite tissue inflammation as a type of dysfunction, then the requirement that dysfunction must be present is dangerously close to trivial, since this can easily apply to all structures, selected or not. Additionally, such responses may be in danger of mislocating the pathology. Appendicitis seems to be a disorder of the appendix, regardless of the status of its components.

At this point, we might wonder how such clashes of intuition can be resolved. Murphy and Woolfolk claim appendicitis is a disorder of the appendix, while Wakefield claims it is a disorder of the appendix's tissues. Murphy and Woolfolk claim that a strangely formed chin is a disorder, while

Wakefield claims it is just a sign that something closely related is disordered. It seems unlikely that intuitions about cases will solve this disagreement and, indeed, perhaps we have reached a limit for the role intuitions can play in the debate. This brings us back to a point we made in our introduction. An account of pathology should capture paradigmatic cases, but it should do more than that. For example, it 'ought to provide a way of integrating research [. . .], to underlie a heuristically fruitful taxonomy of mental disorders, and generally do a good job at furthering our understanding of phenomena labelled pathological' (Murphy and Woolfolk [2000a], p. 242). If using an account that delivers on this desideratum costs us disputed intuitions, or even clear intuitions regarding peripheral or rare cases, perhaps that is a price worth paying.

We do not pretend that this short discussion has conclusively vindicated the selected effects account of dysfunction.[7] We do take ourselves to have shown that the selected effect account of dysfunction has not been given a fair go in the philosophy of medicine. Perhaps the selected effect account of dysfunction doesn't apply to each and every intuitive case of pathology, but it has not been decisively refuted, nor are its failings so evident that it can be disposed of in a footnote. In the next section, we outline a reason for taking the selected effects account very seriously: it avoids deep and intractable problems for the biostatistical account.

4 Problems for the Biostatistical Account

The BST has endured sustained criticism since 1977, and a great many purported counterexamples to the view have been produced. Two of the most persistent issues highlighted by these examples are the 'epidemic problem' and the 'reference class problem'. First, it seems plausible that there are instances where a disease is so widespread that it is statistically normal for members of the reference class to be affected (Neander [1991]; Schwartz [2007]; Kraemer [2013]). For example, body lice were almost ubiquitous in humans in the not so distant past and are still ubiquitous in many other species. Looking to the future, if obesity continues to increase worldwide, then metabolic problems such as type II diabetes may become statistically typical in at least certain populations. It seems odd to say that the effects of parasitic infection or diabetes are not disease states merely because they predominate statistically.

Second, BST lacks a principled basis for determining the reference class within which to assess statistical normality. Boorse's reference classes are

[7] One important omission here is Kenneth Schaffner's ([1993]) discussion of the role of selected effect function in physiological research. Schaffner claims that it plays no essential role at all. We don't have space to discuss this properly, but our belief is that the role evolutionary thinking plays in discovering physiological mechanisms is separate from the role it plays in determining which mechanisms are normal and which are pathological (Griffiths [2009]). Here, we are only concerned with the latter issue.

partitioned by age and sex. This seems intuitively correct, but we are not told why it is correct (see Kingma [2007]). Boorse states that 'the reference class was restricted by sex and age because of differences in normal physiology between males and females, young and old' (Boorse [1997], p. 8). But this reliance on 'normal differences' is unhelpfully circular. We need an independently motivated reason to partition the population in one way as opposed to another.

These two problems coincide in so-called diseases of old age. If we partition the population according to age as Boorse suggests, then deleterious physical states that become widespread at certain ages will not be classed as dysfunctional; they will be 'epidemics' within a particular reference class. For example, if it is typical to have marked osteoarthritis at age eighty, the BST will declare osteoarthritis to be normal functioning—a physiological variant, rather than an instance of pathology. But biomedicine certainly classifies osteoarthritis as pathology, and it seems right to do so. Boorse is aware of this issue, and has attempted at various times to either resolve it or bite the bullet. For example, Boorse ([2002], p. 103) claimed that if a particular physical state is typical within some age group, that state is not dysfunctional within that age group: 'medicine is wrong [. . .] what is pathological is only age-excessive atherosclerosis, premature prostate cancer, and so on'.

We believe that a philosophical account of dysfunction can be revisionary: it ought to have the resources to resist some intuitive judgements regarding pathology where necessary. However, such resistance needs to be justified in terms of the broader goals for which we are explicating the concept. Osteoarthritis is not some peripheral case of pathology—it characteristically manifests as (potentially very severe) immobility and pain, caused by bone surfaces rubbing against each other instead of being cushioned by cartilage. Moreover, osteoarthritis is only one of many 'diseases of old age'. Boorse ([2014]) has recently considered an alternative view, according to which normal functioning is indexed to the young adult phenotype. It is unclear why one would adopt this addition, beyond the fact that it saves some intuitions. Moreover, it would encounter the reverse problem, making a great many aspects of normal childhood and normal ageing pathological. We show in Section 6 that a principled solution to systematic age-related change is available using the selected effects account.

The selected effect account of function successfully negotiates these classic objections to the statistical account. First, epidemics of disease present no difficulty for the selected effect view. Every instance of a trait in the current population may fail to produce the effect for which it was selected in the past, and in such an instance, the selected effect account diagnoses an epidemic of dysfunction. Second, the selected effect account avoids the reference class problem. The classes relevant to attributions of function and dysfunction

are objectively and independently determined: they are the lineages that feature in evolutionary explanations of the prevalence of traits in a population.[8]

In the previous section, we saw that the selected effects account can effectively rebut many of the arguments offered against it. Here, we have seen that it effectively addresses the two main problems of the biostatistical framework. These two themes come together in a surprising way in the work of Peter Schwartz, another critic of the selected effects account.

4.1 Schwartz

Schwartz ([2007]) raises the issue of 'partial dysfunction' that was mentioned earlier. He argues that the performance of many traits fall on a gradient with no clear point at which to declare the trait functional or dysfunctional. For example, heart ejection fraction—the efficiency with which the heart pumps blood—can vary continuously. How low must the ejection fraction be before someone is classified as having a dysfunctional heart? Schwartz argues that the selected effect approach is unable to negotiate this issue. It is indeterminate if any particular ejection fraction has been selected because whether a trait is favoured by natural selection depends on the environment and the competing alternatives. For example, a particular ejection fraction might be selected for in a population of conspecifics with poor heart activity, but selected against in a population with excellent hearts.

This is, indeed, a neglected issue in the philosophical literature on selected effect function. However, there are plenty of intellectual resources in evolutionary biology to meet this challenge, so, like many of the authors discussed above, Schwartz presents only a *prima facie* objection. Selection in spatially and temporally heterogeneous environments (where 'environment' can include conspecifics) is a major topic in evolutionary theory. It can lead to the evolution of polymorphisms, in which many types coexist in a population as result of selection, or phenotypic plasticity, in which the optimal phenotype manifests itself differently across different environments. Phenotypic plasticity may be either intra-generational or inter-generational, with parents producing different offspring on the basis of their 'experience' of the environment. Another form of phenotypic plasticity occurs in the physiological mechanisms of homeostasis and allostasis, in which the proper function of a trait is defined by the present or predicted value of one or more environmental variables, respectively. In all these ways, many different values of a trait may coexist

[8] There are some tough conceptual issues about the individuation of evolutionary processes (see, for example, Brandon [1990]) that critics of the selected effects account could use to argue that it is indeterminate or unknowable which lineage features in the explanation of the recent maintenance of a trait. This, however, would be to deny that biology can produce well-confirmed evolutionary explanations of the recent trajectory of traits in populations, something even many creationists are willing to allow, so we suggest that the burden of proof here is on the critic.

in a single population, each of them performing a different proper function. Ruth Millikan's ([1984]) apparatus of 'derived and adapted' proper functions is probably the only parallel in the philosophical literature to the sophisticated ways in which biology approaches these issues. Another relevant aspect of evolutionary theory is that organisms are designed to allocate different amounts of resources to their various traits according to their circumstances. The 'right' amount of resources for an organism to devote to one of the many activities in which it must engage to survive and reproduce, and hence the 'right' level of performance, differs according to age and many other circumstances (see Section 6). So although this is certainly a complex issue, it is not a reason to discard an evolutionary approach.

Having rejected the selected effects account, Schwartz turns to the biostatistical account. Schwartz has two criticisms: statistical normality does not guarantee correct functioning, and statistical rarity is no guarantee of poor functioning. For example, if heart failure became more prevalent, this would not make heart failure a case of correct functioning. Conversely, the ejection fraction in the lowest 1% of healthy young men might still provide a perfectly adequate blood supply, so it should not be classified as dysfunctional, rare though it is. So according to Schwartz ([2007], pp. 375–6), whether a trait is dysfunctional is independent of the statistical prevalence of the trait. However, rather than discard the biostatistical view, Schwartz claims it can be repaired. A weak heart is dysfunctional, he argues, when (and because) it causes other problems. The fact that a young man's heart is functioning in the bottom 1% for his age group is irrelevant as long as there are no negative consequences of this.

Given that statistical considerations drop out of his account, it is surprising that Schwartz calls this a 'fix' for the BST, rather than a replacement.[9] How Schwartz defines negative consequences is even more surprising: these are the consequences that 'diminish the ability of a part or process in the organism [. . .] to carry out an activity that is generally standard in the species and has been for a long period of time', and this will usually be 'an activity or capacity that has been subject to the process of natural selection in the species' (Schwartz [2007], p. 379). Schwartz's final view thus seems much closer to the selected effects account than to the biostatistical account. This confirms our impression that the selected effects account deserves to be taken more seriously in the philosophy of medicine than it is at present. It has been rejected so definitively that it cannot even be recognized when it is reinvented.

[9] See also (Hausman [2014], p. 635).

5 Analysis versus Explication

In this section, we argue that the selected effects view has a strong theoretical rationale that is lacking in at least recent versions of the biostatistical account. It can serve as an explication of dysfunction, not merely as an analysis. As such, it can guide our intuitions, rather than merely confirm them where they are clear and leave them unclear where people disagree. The history of the biostatistical account is one of ongoing modification to fit counterexamples (Boorse [1997]; Hausman [2011]). Updated versions have then encountered further counterexamples, necessitating further modification (Kingma [2010]). The current state-of-the-art is a version due to Justin Garson and Gualtiero Piccinini ([2014]).[10] Their account ([2014], pp. 6–13) can be summarized as follows, using the standard notation for conditional probability:

> Let *RC* be a reference class, *X* a trait that is distributed over that reference class, and *Y* a function of that trait. Let *S* be a situation-type (a set of instances of situations), such as asphyxiation due to an obstructed throat, or a disjunction of such situation-types, such as asphyxiation due to an obstructed throat or ingestion of a toxic substance.

(1) 'Function'
A function of *X* in *RC* is *Y* $=_{\text{def}}$
- (i) *P*(X contributes to survival or inclusive fitness in *RC*) is non-negligible.
- (ii) *P*(X contributes to survival or inclusive fitness in *RC* by doing *Y*|*X* contributes to survival or inclusive fitness in *RC*) is non-negligible.

(3) Situations where *Y* is 'appropriate'
Relative to *RC*, the set of situations, *S*, under which *X*'s exercise of *Y* is appropriate can be defined by the following two (independently necessary and jointly sufficient) conditions:
- (i) Inclusivity: *P*(*X* is in *S*|*X*'s doing *Y* contributes to survival or inclusive fitness) ∼1.
- (ii) Specificity: there is no *S*′ that is a proper subset of *S* such that (i) is true of *S*′.

(3) 'Appropriate rate of function'
Trait *X* performs function *Y* at a rate of functioning that is appropriate in a situation *s*$=_{\text{def}}$

[10] Garson and Piccinini ([2014]) are pluralists about the concept of function. They defend the biostatistical view here because it is treated in other recent work in philosophy of medicine as the best hope for a naturalist account of dysfunction.

(i) If an organism in *RC* possessing *X* is in *s* and *s* [is not a member of] *S*, then *X* performs function *Y* at a rate of zero (or close to zero) [and]

(ii) If an organism in *RC* possessing *X* is in *s* and *s* [is a member of] *S*, then *X*'s rate of functioning provides an adequate contribution to survival or inclusive fitness in *s*, relative to other rates that are physiologically possible for *X* in *s*.

(3) 'Dysfunction'
Let *X* be a token of a trait-type whose tokens are distributed over reference class *RC*, *Y* a function of *X*, and *S* the set of situations that are appropriate for *X* performing *Y*.
Then: *X* is dysfunctional with respect to $Y =_{\text{def.}}$ *X* cannot perform *Y* in at least one situation, *s*, where *s* [is a member of] *S*, at the rate that is appropriate in *s*.

Summarizing still further, Garson and Piccinini's analysis states that in any setting that is not extremely rare, where *X* performing *Y* might contribute to the organism's survival or inclusive fitness, if *X* fails to perform *Y* adequately with respect to what it is physiologically possible for *X* to do, then *X* is dysfunctional. Note the switch here from a consideration of the statistical population norm to what is 'possible' for that trait. However, 'possibility' is presumably still indexed to what members of a relevant reference class are capable of, rather than the individual in question. It might not be physiologically possible for an individual's severely damaged cardiac muscle to contract effectively, but this cannot be a reason to deny that their heart is dysfunctional. Rather, it must be that their damaged cardiac muscle is dysfunctional because it is possible within some wider human population for cardiac muscle to contract more effectively.

Garson and Piccinini's account deals with many counterexamples to the biostatistical account. Perhaps it will turn out to be not just the latest but the ultimate version, immune to the best efforts of philosophers to imagine scenarios in which it delivers a counterintuitive verdict. In fact, it would be surprising if there were obvious counterexamples because this analysis has been developed to deal with all the problem cases suggested over the past thirty-seven years. And herein, we believe, lies a problem. Perhaps this account can recover our intuitions about whether *X* is dysfunctional by asking whether *X* performs *Y* 'adequately' with respect to what is 'physiologically possible' in some set of scenarios. But the more important issue is why these cases count as dysfunctional, and why it matters that something is dysfunctional. If the only rationale for an analysis is that it fits all of our intuitions about cases, why not simply rely on intuition?

Moreover, the analysis still requires intuition to deliver a determinate verdict about any particular case. The issue of choosing the reference class in an independently motivated way is not addressed, and the notion of 'physiological possibility' in clause 3(ii) is not fully defined. In the absence of principled criteria, these are free variables that can be used to fit the definition to a pre-determined, intuitive result.

A philosophical account of dysfunction should give more guidance when intuitions are unclear, and deliver greater theoretical contributions to our understanding of dysfunction and pathology. Rather than refining an analysis through a series of counterexamples, we ought to find a theoretically grounded explication of the concept (Carnap [1950]) and see what this tells us about both the clear and the difficult cases. An explication needs to accord with our intuitions regarding paradigm examples, but it does not have to achieve this across the board. It certainly does not have to accord with our intuitions regarding esoteric thought experiments. Rather, it is more important that in those cases where the explication disagrees with intuition (and also in those where it agrees), the account ought to be able to explain why it disagrees (or agrees). In this way, we will increase our understanding of the straightforward cases, be able to assess whether we should or shouldn't change our minds in the ones that clash, and generate some further avenues for investigation.

The selected effects account of dysfunction is a good candidate for an explication of dysfunction because the underlying evolutionary rationale for the definition is transparent. It simply applies the orthodox evolutionary concept of adaptation. This means we can address the balance between theoretical coherence and correspondence with intuition in a constructive and well-motivated manner. Also, we will not be open to criticisms of circularity or covert subjectivity regarding the constituent concepts in the account, as these concepts are derived from a mature science, namely, evolutionary biology, to which we can look for additional clarification when it is needed.

In the next section, we illustrate these advantages by showing how life history theory can illuminate the 'diseases of old age' that caused such trouble for the biostatistical account.

6 Explicating Dysfunction: Life History Theory and Senescence

Life history theory (Roff [1992], [2002]; Stearns [1992]) is a good place to start when examining the adaptive functions of traits that are important for health and illness. This branch of evolutionary theory is based on the insight that an organism is a process lasting from birth to death. Natural selection does not design a single, adult phenotype. It designs a changing series of phenotypes—a

life-history.[11] In life-history theory, the basic evolutionary problem facing the organism is to find the optimal way to parcel the resources available to it into offspring. This problem is modelled as the simultaneous optimization of two parameters, the probability of surviving to each age class and the number of offspring produced in each age class, integrated across all age classes (there are both continuous and discrete versions of these models). The primary constraint on this optimization problem is the quantity of resources available to the organism. But it is also constrained by multiple trade-offs between the two key parameters: there is an overall trade-off between survival and reproduction; reproduction in the current age class must be traded off against reproduction later; current reproduction must be also traded off against growth, and against condition (the maintenance of structures that have already developed); current growth or condition may trade off with survival to later age classes. The mechanisms that induce these trade-offs may be either genetic—the same alleles influence both traits—or physiological, for example, nutrition allocated to making somatic cells is not available to make germ cells.

Solving this complex optimization problem under different sets of constraints and in different ecological settings leads to the many different life-history strategies observed in nature. Organisms do not start reproducing the moment they are born because they do not have enough resources to reproduce successfully—they will have a higher lifetime reproductive output if they invest in growth before commencing reproduction. This raises the question of when to stop growing and start reproducing. Many organisms are semelparous—they complete all their growth before engaging in a single round of reproductive activity to which they commit all their resources. A famous example is the Australian marsupial genus *Antechinus*, in which males die at the end of the breeding season, and in one species of which all females die after weaning their offspring. Other organisms, including humans, are iteroparous—they have several rounds of reproduction, which implies staying alive and maintaining condition between each round. One basic evolutionary dynamic affecting the choice between semelparity and iteroparity is that high juvenile survival and low adult survival favours semelparity, whilst the opposite favours iteroparity. There are many other life history choices to be made. For example, humans exhibit determinate growth—they do not continue to increase in size after reaching reproductive maturity, as many fish and reptiles do.

Life history strategies are implemented by the organism's morphology, physiology, and behaviour. This implementation lends itself to functional analysis. For example, we can identify the proper function of shutting down the immune system in male antechinus during the breeding season: it is to free

[11] Recall Boorse's objection that the normal mental abilities of an eight-year old are not normal for a fourteen-year old. According to life-history theory, mental ability at eight and fourteen are separate traits of humans, as separate as legs and arms, and each has its own selection pressures and constraints.

up energy for sperm production, fighting, and copulating. Shutting down the immune system is an adaptation for that purpose: modern antechinus have this trait because of the selective advantage it conferred on their ancestors.

When paired with life history theory, the selected effect account of biological function illuminates a number of otherwise problematic issues regarding dysfunction. One of the most dramatic demonstrations of this is senescence. As mentioned earlier, certain phenotypes, such as osteoarthritis, are statistically normal in later age classes. The selected effects theory confirms the intuition that these states are dysfunctional. However, the advantage of this approach is not just that it coincides better with our intuitions. Rather, the evolutionary theory of senescence provides solid reasons why we should distinguish between phenotypes that manifest in later age classes, recognizing many of them of them as dysfunctional but, more speculatively, recognizing some as adaptations to the specific ecological demands of those age classes. Drawing these distinctions is not only necessary to explain why we observe these phenotypes, they also have an important heuristic role for biomedical research and potential implications for treatment.

As an example of the evolutionary theory of senescence, consider three age-classes in a contemporary East Asian population (Table 1). Those aged two years can almost all digest lactose and none of them have osteoarthritis. Most of those aged twenty no longer produce lactase and so cannot digest lactose, but few of them have developed osteoarthritis. Almost all of those aged eighty have developed osteoarthritis. So are lactose maldigestion and osteoarthritis dysfunctional or not? The person on the street in Western countries regards both lactose maldigestion and osteoarthritis as diseases, while Boorse's original version of BST regarded neither as diseases. How are we to decide?

The two ways in which phenotypes change over time in Table 1 represent two very different biological phenomena. Understanding these two phenomena provides a rationale for classifying the phenotypes as functional or dysfunctional.

Lactose digestion in mammals is an adaptation to the ecological demands of infancy. Mammals switch off production of lactase when they wean because to do otherwise would waste resources. Lactase production is an adaptation for a specific stage of development and ceases operation as part of normal mammalian development. However, when food sources containing lactose became important for adults in some ancestral human populations, then, in the most famous example of gene–culture co-evolution, continued expression of lactase in later age classes was favoured by natural selection. Those human populations, which include northern Europeans, are an exception to the ancestral pattern of mammalian development. Most human populations, however, have not been under selective pressure to digest lactose, and so retain the ancestral pattern and switch off the lactase gene at weaning. Problems occur when

Table 1. Prevalence of lactose maldigestion and osteoarthritis in three age classes in a hypothetical but realistic contemporary East Asian population. (Actual numbers are not definite due to varying diagnostic criteria and distinctions between affected joints.)

Age class	Lactose maldigestion	Osteoarthritis
2	<1%	<1%
20	90%	10%
80	90%	80%

people from these populations are in an environment that strongly favours lactose digestion, potentially resulting in the problematic symptoms of lactose intolerance.

This leads to some interesting results. First, the selected effects account appears to suggest that whether lactose maldigestion is dysfunctional depends on one's ancestry. If an individual with lactose maldigestion has a simple pattern of geographic ancestry, they can be objectively assigned to a specific evolutionary lineage in which this is either part of normal functioning, as in East Asians, or dysfunctional, as in northern Europeans.[12]

Second, if the individual has complex geographic ancestry, as is increasingly likely today, then whether their lactose maldigestion is a dysfunctional phenotype becomes objectively vague. By 'objectively vague' we mean that there are factual grounds for assigning an individual to a zone between two categories.[13] Doctors may need to declare a sharp cut-off for official diagnosis, but this is often a pragmatic clinical decision, rather than an application of a determinate concept of pathology. Here, perhaps, lies the best answer to Schwartz's ([2007]) concern with 'drawing a line': sometimes, it is indeed vague whether dysfunction is occurring or not, but that simply means we are representing the biological facts accurately.

In the case of lactose maldigestion, there is no clinical need to determine whether individuals with complex geographic ancestry have dysfunctional alleles or selected alleles. That distinction could at least sometimes be made by

[12] It was pointed out by an anonymous reviewer that this means two patients with lactose intolerance might be in differing states of health and disease purely due to different ancestry. Our first reaction to this point was that it is the correct result. However, we have since realized that the issue is more subtle than this, not least because maldigestion of lactose does not necessarily lead to lactose intolerance, which might be viewed as a state of pathology regardless of the root cause. For most humans, it is not abnormal to stop producing lactase, but it is abnormal to have resulting gastrointestinal dysfunction, even if this is the effect of an evolutionarily novel environment (see also Gluckman *et al.* [2009], p. 5). Our thanks to the reviewer for pushing us on this question.

[13] In a Darwinian world, many biological categories grade into one another. For example, during a speciation event or in a hybrid zone, there are individuals who cannot be clearly assigned to any one species. But other individuals are definitively in one species and not the other.

identifying the geographic origins of the haplotypes containing the alleles in each individual, but it would serve no clinical purpose. In the context of clinical care, it may still be medically appropriate to advise someone with lactose intolerance to avoid dairy products, even when their lactose intolerance does not arise due to a dysfunction. Similarly, advising light-skinned individuals to wear sunscreen when they visit Australia does not imply any dysfunction is occurring.

The development of osteoarthritis is a phenomenon of a very different kind, but is also illuminated by life history theory. Like many iteroparous species, humans undergo senescence, or deterioration in condition in later age classes. Senescence is explained in life history theory in a number of ways. The most famous is 'antagonistic pleiotropy' (Williams [1957]), in which some allele increases fitness in earlier age classes, but reduces fitness in later age classes. Because younger organisms have a greater reproductive potential than older organisms, such alleles may be favoured by selection, leading to the evolution of species-typical disease outcomes in older age-classes. Senescence may also be explained simply by a high rate of extrinsic mortality. An organism that is less likely to survive to a later age class because of, for example, a high rate of predation, should allocate more resources to reproduction and less to maintaining condition.

Senescence differs from development (ontogeny) because there is no selection for senescent phenotypes. In development, earlier phenotypes are succeeded by later ones because the later phenotype increases fitness in the age class in which the phenotype is manifested. Embryonic haemoglobins are just what is needed in the early months of gestation, but adult haemoglobins are, unsurprisingly, better when you are an adult. In stark contrast, people do not develop osteoarthritis because stiff and painful joints are better suited to the ecological demands of later life. Evolutionary explanations of senescence recognize that the senescent phenotype always imposes a cost—the organism would be fitter if it retained the phenotype seen in earlier age classes. The senescent phenotype is explained by genetic or physiological linkage, either of which can tie this phenotype to another, fitness-enhancing phenotype. Here, we have a well-motivated, definitive answer why osteoarthritis is a pathological state no matter how prevalent it becomes: joints are selected to move freely within a certain range, and never selected to stiffen and become painful.

In normal English usage, the term 'senescence' refers to the entire suite of phenotypic changes characteristic of later age classes. But not all such phenotypes need be senescent in the sense defined by evolutionary explanations of ageing. One advantage of a biologically informed, selected effect approach to dysfunction is that it reveals an important distinction between genuinely senescent phenotypes and age-specific adaptations to the ecological demands of

later age classes. For example, menopause has traditionally been regarded as simply one more example of reduced function in old age. Most mammals, however, remain fertile for a much larger proportion of their lifespan—in many cases, until shortly before death. So human menopause, which can occur less than two-thirds into the lifespan, cries out for explanation. If the 'grandmother hypothesis' (Hawkes [2003]) is correct, then menopause is an adaptation for multi-generational childcare—an age-specific adaptation like infant lactase production or embryonic haemoglobin (for other examples of adaptations to old age, see Le Couteur and Simpson [2011]).

It is worth noting that this kind of insight may have real consequences for medicine. If menopause is an adaptation, for example, then we might expect to find a single mechanism driving the complex changes in gene expression that underlie menopause. Moreover, when investigating the causes of individual differences in the onset of menopause, we might look for environmental cues that carry information that would be relevant in light of our hypothesis about its evolutionary function.

The lesson we draw from these examples is that even deeply seated intuitions about dysfunction are defeasible in light of new discoveries, and particularly new discoveries about the evolutionary reasons for the existence of phenotypes.

7 Conclusion

The selected effect approach to dysfunction in the philosophy of medicine is not 'dead in the water', as many authors appear to believe. On the contrary, once we adopt the orthodox version of the selected effects account, the so-called modern history version, and take evolutionary theory seriously, the main criticisms of the view can be deflected.

There are three desiderata for a philosophical account of dysfunction in the context of biomedicine. It should capture all the cases of dysfunction, only cases of dysfunction, and it should say something substantive and informative about dysfunction. Much of the existing literature has emphasized the first two desiderata. Because of this, the current version of the popular biostatistical view, honed on many counterexamples, takes almost a page to summarize in relatively technical language and still includes free parameters that can be adjusted to fit what seems intuitive.

Conversely, we have emphasized the third desideratum. Our approach is an explication of the concept of dysfunction, not merely an analysis of the concept as we find it today. Because it has a strong theoretical rationale, our evolutionary approach to dysfunction can challenge established ways of thinking and offers substantive reasons for reversing intuitive judgements about dysfunction in some cases. Our discussion of senescence illustrates

this, showing that an evolutionary explication of the concept of dysfunction can improve our understanding of dysfunction in the context of the diseases of old age.

Our approach also grounds the ideas of function and dysfunction in a mature scientific discipline. It emerges naturally from the core concept of adaptation in evolutionary theory. When determining the function of a specific trait we can appeal not to our intuitions, but to an established scientific practice, which can offer substantive evidence for one function ascription rather than another.

The selected effect approach to biological function is not without its problems. But we believe that it deserves a more prominent place in the philosophy of medicine than it currently occupies. To discover if a naturalist account of dysfunction and pathology is possible, we must be sure we have settled on the best naturalist account. Until the selected effect view is given proper attention, this will not occur.

Funding

This research was supported under the Australian Research Council's *Discovery Projects* funding scheme (grant no. DP130101774).

Acknowledgements

We would like to acknowledge the generous help of Justin Garson, Stefan Gawronski, Elselijn Kingma, Pierre-Olivier Méthot, Dominic Murphy, Gualtiero Piccinini, and three anonymous reviewers for this journal.

Paul E. Griffiths
Department of Philosophy and Charles Perkins Centre
University of Sydney
NSW, Australia
paul.griffiths@sydney.edu.au

John Matthewson
School of Humanities
Massey University
Albany, New Zealand
J.Matthewson@massey.ac.nz

References

Boorse, C. [1976]: 'Wright on Functions', *The Philosophical Review*, **85**, pp. 70–86.

Boorse, C. [1977]: 'Health as a Theoretical Concept', *Philosophy of Science*, **44**, pp. 542–73.

Boorse, C. [1997]: 'A Rebuttal on Health', in J. M. Humber and R. F. Almeder (*eds*), *What Is Disease*? Totowa, NJ: Humana Press, pp. 1–134.

Boorse, C. [2002]: 'A Rebuttal on Functions', in R. Cummins, A. Ariew and M. Perlman (*eds*), *Functions: New Essays in the Philosophy of Psychology and Biology*, Oxford: Oxford University Press, pp. 63–112.

Boorse, C. [2014]: 'A Second Rebuttal on Health', *Journal of Medicine and Philosophy*, **39**, pp. 683–724.

Brandon, R. [1990]: *Adaptation and Environment*, Princeton, NJ: Princeton University Press.

Carnap, R. [1950]: *Logical Foundations of Probability*, Chicago, IL: University of Chicago Press.

Cooper, R. [2002]: 'Disease', *Studies in the History and Philosophy of Biology and Biomedical Sciences*, **33**, pp. 263–82.

Cummins, R. [1975]: 'Functional Analysis', *Journal of Philosophy*, **72**, pp. 741–65.

Diamond, J. M. [1997]: *Guns, Germs, and Steel: The Fates of Human Societies*, New York: W. W. Norton.

Garson, J. [2013]: 'The Functional Sense of Mechanism', *Philosophy of Science*, **80**, pp. 317–33.

Garson, J. and Piccinini, G. [2014]: 'Functions Must Be Performed at Appropriate Rates in Appropriate Situations', *British Journal for the Philosophy of Science*, **65**, pp. 1–20.

Glass, D., Lens, M., Swaminathan, R., Spector, T. D. and Bataille, V. [2009]: 'Pigmentation and Vitamin D Metabolism in Caucasians: Low Vitamin D Serum Levels in Fair Skin Types in the UK', *PLoS ONE*, **4**, p. e6477.

Gluckman, P., Beedle, A. and Hanson, M. [2009]: *Principles of Evolutionary Medicine*, New York: Oxford University Press.

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

Godfrey-Smith, P. [1994]: 'A Modern History Theory of Functions', *Noûs*, **28**, pp. 344–62.

Goode, R. and Griffiths, P. E. [1995]: 'The Misuse of Sober's Selection for/Selection of Distinction', *Biology and Philosophy*, **10**, pp. 99–108.

Griffiths, P. E. [1993]: 'Functional Analysis and Proper Function', *British Journal for Philosophy of Science*, **44**, pp. 409–22.

Griffiths, P. E. [2009]: 'In What Sense Does "Nothing Make Sense Except in the Light of Evolution"?', *Acta Biotheoretica*, **57**, pp. 11–32.

Hausman, D. [2011]: 'How Can Irregular Causal Generalizations Guide Practice?', *Preventive Medicine*, **53**, pp. 229–31.

Hausman, D. [2012]: 'Health, Naturalism, and Functional Efficiency', *Philosophy of Science*, **79**, pp. 519–41.

Hausman, D. [2014]: 'Health and Functional Efficiency', *Journal of Medicine and Philosophy*, **39**, pp. 634–47.

Hawkes, K. [2003]: 'Grandmothers and the Evolution of Human Longevity', *American Journal of Human Biology*, **15**, pp. 380–400.

Jablonski, N. G. and Chaplin, G. [2000]: 'The Evolution of Human Skin Coloration', *Journal of Human Evolution*, **39**, pp. 57–106.

Kingma, E. [2007]: 'What Is It to Be Healthy?', *Analysis*, **67**, pp. 128–33.

Kingma, E. [2010]: 'Paracetamol, Poison, and Polio: Why Boorse's Account of Function Fails to Distinguish Health and Disease', *British Journal for the Philosophy of Science*, **61**, pp. 241–64.

Kingma, E. [2013]: 'Naturalist Accounts of Mental Disorder', in K. W. M. Fulford (*ed.*), *The Oxford Handbook of Philosophy and Psychiatry*, Oxford: Oxford University Press, p. 363.

Kingma, E. [2014]: 'Naturalism about Health and Disease: Adding Nuance for Progress', *Journal of Medicine and Philosophy*, **39**, pp. 590–608.

Kingma, E. [2016]: 'Situation-Specific Disease and Dispositional Function: Response to Hausman', *British Journal for the Philosophy of Science*, **67**, pp. 391–404.

Kraemer, D. M. [2013]: 'Statistical Theories of Functions and the Problem of Epidemic Disease', *Biology and Philosophy*, **28**, pp. 423–38.

Lack, D. [1947]: 'The Significance of Clutch-Size', *Ibis*, **89**, pp. 302–52.

Le Couteur, D. and Simpson, S. J. [2011]: 'Adaptive Senectitude: The Prolongevity Effects of Aging', *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, **66**, pp. 179–82.

Lemoine, M. [2013]: 'Defining Disease beyond Conceptual Analysis: An Analysis of Conceptual Analysis in Philosophy of Medicine', *Theoretical Medicine and Bioethics*, **34**, pp. 309–25.

Lewens, T. [2015]: *The Biological Foundations of Bioethics*, Oxford: Oxford University Press.

Millikan, R. G. [1984]: *Language, Thought, and Other Biological Categories: New Foundations for Realism*, Cambridge, MA: MIT Press.

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

Murphy, D. [2009]: 'Concepts of Disease and Health', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, available at < plato.stanford.edu/archives/sum2009/entries/health-disease/>.

Murphy, D. and Woolfolk, R. L. [2000a]: 'Conceptual Analysis versus Scientific Understanding: An Assessment of Wakefield's Folk Psychiatry', *Philosophy, Psychiatry, and Psychology*, **7**, pp. 271–93.

Murphy, D. and Woolfolk, R. L. [2000b]: 'The Harmful Dysfunction Analysis of Mental Disorder', *Philosophy, Psychiatry, and Psychology*, **7**, pp. 241–52.

Neander, K. [1998]: 'Mental Illness, Concept Of', in E. Craig (*ed.*), *Routledge Encyclopedia of Philosophy*, London: Routledge, <rep.routledge.com/articles/mental-illness-concept-of/v-1/>.

Neander, K. [1991]: 'Functions as Selected Effect: The Conceptual Analyst's Defense', *Philosophy of Science*, **58**, pp. 168–84.

Nesse, R. M. [2001]: 'On the Difficulty of Defining Disease: A Darwinian Perspective', *Medicine, Health Care, and Philosophy*, **4**, pp. 37–46.

Nesse, R. M. [2007]: 'Evolution Is the Scientific Foundation for Diagnosis: Psychiatry Should Use It', *World Psychiatry*, **6**, pp. 160–1.

Papineau, D. [1987]: *Reality and Representation*, New York: Blackwell.

Phillips, B. L., Kelehear, C., Pizzatto, L., Brown, G. P., Barton, D. and Shine, R. [2010]: 'Parasites and Pathogens Lag behind Their Host during Periods of Host Range Advance', *Ecology*, **91**, pp. 872–81.

Roff, D. A. [1992]: *Evolution of Life Histories: Theory and Analysis*, New York: Chapman and Hall.

Roff, D. A. [2002]: *Life History Evolution*, New York: W. H. Freeman.

Schaffner, K. F. [1993]: *Discovery and Explanation in Biology and Medicine*, Chicago, IL: University of Chicago Press.

Schwartz, P. H. [2007]: 'Defining Dysfunction: Natural Selection, Design, and Drawing a Line', *Philosophy of Science*, **74**, pp. 364–85.

Schwartz, P. H. [2014]: 'Reframing the Disease Debate and Defending the Biostatistical Theory', *Journal of Medicine and Philosophy*, **39**, pp. 572–89.

Sober, E. [1984]: *The Nature of Selection: Evolutionary Theory in Philosophical Focus*, Chicago, IL: University of Chicago Press.

Stearns, S. C. [1992]: *The Evolution of Life Histories*, Oxford: Oxford University Press.

Wakefield, J. [1992]: 'The Concept of Mental Disorder: On the Boundary between Biological Facts and Social Values', *American Psychologist*, **47**, pp. 373–88.

Wakefield, J. [2000]: 'Spandrels, Vestigial Organs, and Such: Reply to Murphy and Woolfolk's "The Harmful Dysfunction Analysis of Mental Disorder"', *Philosophy, Psychiatry, and Psychology*, **7**, pp. 253–69.

Wakefield, J. [2007]: 'The Concept of Mental Disorder: Diagnostic Implications of the Harmful Dysfunction Analysis', *World Psychiatry*, **6**, pp. 149–56.

Williams, G. C. [1957]: 'Pleiotropy, Natural Selection, and the Evolution of Senescence', *Evolution*, **11**, pp. 398–41.

Wilkins, J. S. and Griffiths, P. E. [2013]: 'Evolutionary Debunking Arguments in Three Domains: Fact, Value, and Religion', in G. Dawes and J. Maclaurin (*eds*), *A New Science of Religion*, New York: Routledge, pp. 133–46.

Wright, L. [1973]: 'Functions', *The Philosophical Review*, **82**, pp. 139–68.