

Biological Criteria of Disease: Four Ways of Going Wrong

JOHN MATTHEWSON*

Massey University, Albany, Auckland, New Zealand

PAUL E. GRIFFITHS

University of Sydney, Sydney, New South Wales, Australia

*Address correspondence to: John Matthewson, PhD, Massey University, Atrium, Albany, Auckland 0632, New Zealand. E-mail: j.matthewson@massey.ac.nz

We defend a view of the distinction between the normal and the pathological according to which that distinction has an objective, biological component. We accept that there is a normative component to the concept of disease, especially as applied to human beings. Nevertheless, an organism cannot be in a pathological state unless something has gone wrong for that organism from a purely biological point of view. Biology, we argue, recognises two sources of biological normativity, which jointly generate four “ways of going wrong” from a biological perspective. These findings show why previous attempts to provide objective criteria for pathology have fallen short: Biological science recognizes a broader range of ways in which living things can do better or worse than has previously been recognized in the philosophy of medicine.

Keywords: *biological function, biological normativity, disease, pathology*

1. INTRODUCTION

Analyzing the concept of (primarily human) disease has been a central pre-occupation of the analytic philosophy of medicine since the 1970s.¹ One of the core issues in this literature is the respective roles of the so-called “biological” and “social” facts. Can disease be analyzed solely in terms of human biology, solely in terms of values and social practices, or only with some mixture of the two? If human disease is our sole focus, a powerful case can be mounted that the science of medicine is like the science of criminology. No matter how rigorous and objective our knowledge about the causes and

consequences of murder and adultery, we cannot deduce that only one is a crime. That is a legal matter. In the same way, it can be argued that the biomedical sciences document the causes and consequences of phenotypic variation but do not determine which of these variations are diseases. Many philosophers of medicine go further, claiming that not only is there no *sufficient* biological condition for being a disease, there is no *necessary* biological condition for being a disease (Glackin, 2010). To continue our analogy, this approach corresponds to the view that any action whatsoever might be a criminal act in some society. According to these philosophers, any phenotypic trait could be either normal or pathological, given the right societal context.²

Efforts to refute this strongly social view of disease usually focus on its relativistic consequences. If an exclusively social account of disease is correct, then psychiatrists in the Soviet Union did not make a biological error when they diagnosed political dissidents as suffering from the disease of “sluggish schizophrenia.” Rather, they applied the concept of disease correctly, given the values and social practices of the Union of Soviet Socialist Republics at that time. Their only error was a moral one: They endorsed values that were elsewhere—or have since been—rejected. Critics of purely social accounts claim that this is not enough: Soviet doctors who made this diagnosis were performing bad science as well as endorsing bad values.

We think that the strong social view of disease is mistaken, but not (just) because of worries about social relativism. Disease is a concept with obvious application to nonhuman organisms, including organisms about which no society cares or even knows. Disease is one of the major factors in the evolution of life. Particular evolutionary dynamics result from the interaction of organisms and infectious agents (Haldane, 1949), and the distinction between the normal and the pathological also finds a home in some very mundane processes of biological description. We contend that there is a real tension between these observations and a *purely* social view of the category of disease. The concept of disease as it features in medicine and is applied primarily to human beings or their domesticates certainly appears to have a strong evaluative component. However, that concept is founded on a broader distinction between normal and pathological phenotypes that applies to all organisms and which, we argue, is governed by an objective biological criterion.

Our argument here was considered, although ultimately dismissed, by Georges Canguilhem in one of the founding texts of the philosophy of medicine (Canguilhem, 1991, 220). Canguilhem stated that “[s]tructures or behaviours can be objectively described, but they cannot be called ‘pathological’ on the strength of some purely objective criterion” (Canguilhem, 1991, 226). This is a counterintuitive claim: As Christopher Boorse (2014) points out, when an ant presents with fewer than six legs after a tussle with a predator, it seems that biologists are justified in concluding that this is pathological

rather than a normal variant. But we can do better than appeal to intuition here. It is uncontroversial that interactions between predators and prey or between parasites and hosts produce distinctive evolutionary dynamics. We can be confident that these dynamics played a role in the evolution of many extinct species that left no fossils and will be forever unknown to science. Someone who claims that there are no objective criteria of pathology must either deny this or argue that the processes of predation and parasitism can be described without implying a distinction between normal and pathological states of the prey/host. Neither option seems plausible.

The concept of disease is a complex one, and it is not implausible to argue that there is a distinctive concept of *human* disease. Nevertheless, human and nonhuman disease ought to be variations on a theme, not simply homonyms. Canguilhem suggests that “these notions [*pathology, malady*] are applied to all living beings through sympathetic regression starting from lived human experience” (Canguilhem, 1991, 223). However, there is more to the biology of disease than an analogy with the human experience of suffering. We suggest instead that pathology is a fundamental biological phenomenon that becomes more complex in a species like our own, due at least in part to our unrivalled degree of developmental and ecological niche construction (Sterelny, 2003, 2012). For this reason, our strategy in seeking to understand the disease concept is to identify the biological criteria for pathology, to which social criteria may be added to accommodate the additional complexity of the concept of disease when applied to human beings.

We begin by outlining the two main extant accounts that claim there are at least *necessary* objective biological criteria for disease and some of the core criticisms of these accounts. We then argue that these criticisms are not sufficient to discount the role of biology in this domain. The notions of abnormality and pathology play essential roles in nonmedical biology, and a more careful assessment of these roles shows why previous attempts to restrict disease through “objective” criteria have failed. We identify four ways in which biologists recognize that things are going badly for a living thing and note that, as living things, human beings also face these “ways of going wrong.” These failures of biological normativity do not necessarily constitute disease states in and of themselves, but they can be deployed as part of an account of disease. The article then explores some of the details and potential problems of this view. Previous biological accounts of disease have been criticized as too restrictive in what they allow to count as disease, but our constraint must be prevented from being overly liberal.

II. TWO ACCOUNTS OF BIOLOGICAL CRITERIA FOR DISEASE

There are two broad camps among the philosophers of medicine who advocate a biological constraint on what can count as a disease. Some argue

that disease involves the failure of some structure or behavior to contribute adequately to biological fitness, while others claim that disease involves the failure of some component to perform the function for which it was selected in the evolutionary past.

In what is surely the most highly cited paper in the philosophy of medicine, Christopher Boorse defined disease as the substandard contribution of a bodily part or process to fitness, in comparison to others in the appropriate reference class (Boorse, 1977, 2014). Boorse called this the “biostatistical” (BST) account of disease.

The other popular “objectivist” account states that dysfunction occurs when a part or process fails to produce the effect that led to the evolution of that part or process by natural selection. Jerome Wakefield is the most prominent recent advocate of this view (Wakefield, 1992; see also Neander, 1983, 1998). Wakefield’s “harmful dysfunction” account of medical disorders requires both that biological dysfunction occurs and that this dysfunction is considered harmful for the individual. In this article, however, we are not concerned with Wakefield’s overall account of medical disorders, only with his biological criterion of dysfunction.

These two accounts correspond to two widely recognized notions of biological function (Griffiths, 1993; Godfrey-Smith, 1993). The former utilizes “causal role” function with the added stipulation that the causal role is survival and reproduction—christened “evolutionary function” by Paul Griffiths (2009). The latter uses the notion of “selected-effect” or “Proper” function (Millikan, 1984). Selected effect function is the same idea that biologists know as “teleonomy” (Pittendrigh, 1958).

Neither of these accounts of disease has been widely accepted, and indeed a large portion of the literature in philosophy of medicine is devoted to criticizing Boorse’s account: “. . . the BST’s influence is hardly due to its multitude of converts” (Boorse, 1997, 4). Both accounts have been inundated by counterexamples, many of which appear to show that they are too restrictive: that they exclude genuine cases of disease. Like many recent authors, we are not convinced that the process of analysis-counterexample-revision is a useful way to explore a topic like the nature of pathology (Schwartz, 2004; Lemoine, 2013; Griffiths and Matthewson, 2016; Walker and Rogers, forthcoming). Nevertheless, even an explication or revisionary account of a concept must keep reasonably close to the intuitive meaning (Carnap, 1950). If an account of disease excludes paradigmatic instances of disease, there is reason to doubt its adequacy.

For example, the biostatistical account of disease will not include diseases found at epidemic levels for long periods of time. There are few if any cockatoos without feather lice, but feather lice are still parasites, and the birds would be healthier without them. Additionally, some paradigmatic cases of disease may serendipitously lead to improved fitness, and then would not be classified as disease by the biostatistical view. For example, Wakefield

discusses the case of an infection that protects against later illnesses, as occurred with the protective effects of cowpox against smallpox infection (Wakefield, 2000).

Conversely, Wakefield's view excludes cases where there is no clear selected effect of the relevant structure or behavior. For example, if we accept that the appendix is a vestigial organ, this seems to entail that the appendix cannot be dysfunctional according to the selected effects account (Murphy and Woolfolk, 2000). Similarly, failure to perform certain abilities that are currently common in the population might seem detrimental in medically relevant ways, regardless of selective history. A putative example of this is a failure to be able to read, in spite of adequate opportunity to learn. Reading is too recent a behavior to have been under natural selection, yet we might think that inability to read in such circumstances is a disease state (Murphy and Woolfolk, 2000; Kingma, 2013).

Proponents of the two views have modified and clarified their positions to address these points (Boorse, 2014; Wakefield, 2000). Nevertheless, these two prominent views do not appear to have convinced the majority of commentators, at least in part due to their apparent excessive restrictiveness. This has led a number of authors to suggest that no adequate objective account of disease is in the offing, and to favor views that place more emphasis on social facts. Elselijn Kingma, for example, concludes that Boorse's biostatistical theory is "the best and only presently existing naturalistic account . . ." and yet is "inadequate, both as a naturalistic account of dysfunction and as a naturalistic account of disease" (Kingma, 2009, 22).

We think this move away from a biological restriction on what can count as disease is premature for two reasons. The first is that there are many resources available to defenders of a "naturalistic" account that have been neglected in the existing literature, as we have argued elsewhere (Griffiths and Matthewson, 2016). The second reason is that it would create a conceptual divorce between human disease and pathology as a biological phenomenon. The focus on human disease in the philosophy of medicine has led to a neglect of the other roles that the distinction between normal and pathological states plays in biology.

Dysfunction and abnormality occur throughout the biological world, including contexts where there are no social facts to underpin judgments of that dysfunction or abnormality. When a biologist encounters a new species of beetle, a core part of his or her work will be determining which specimens are relevant to describing the biology of that species and which are only relevant to the pathologist. For example, beetles squashed during collection, or half-eaten by a predator, should not be included in the description of the new species. Those which differ from their conspecifics because they have a high parasite load need to be distinguished from the healthy specimens, not lumped in to determine an average value for some phenotype. Biologists do not proceed in these ways because of intuitions about the meaning of words

such as “normal” or “disease.” They do so because refusing to distinguish the normal from the pathological in these contexts would be to overlook substantive biological facts.

How exactly these distinctions are made is of course a difficult issue. Evolution produces a complex pattern of biodiversity, and distinct forms are often not separated by clear boundaries; parasitism merges into mutualism and symbiosis, for example, and determining what is physiological or pathological requires substantial scientific work. But few of these problems have any relationship with the prevailing values of human societies. For example, a recent article proposes that locusts with unusually short wings in a Japanese population may be a “short-winged ‘morph’” rather than individuals with deformed wings (Nishide and Tanaka, 2013, 171). Evidence is presented to distinguish between the two hypotheses of polymorphism or pathology. This evidence ranges from genetics to ecology but at no point do the authors appeal to the prevailing values in Japanese (or locust) society.

The key point is that biologists recognize and utilize *biological normativity*. Biology itself distinguishes between normal variants and those that are abnormal, and between organisms that are flourishing and those that are not. If biologists omitted these evaluative claims, they would be omitting an important part of what is going on in the systems and processes that biology seeks to understand.

We suggest there is strong *prima facie* reason to think the distinction between health and disease in human beings is closely connected with these biological facts. Given that all other living things can be in various states of pathology or biological abnormality, we would need very strong reasons to think anything different of human beings. And, if human beings can be in states of pathology or biological abnormality, we would need very strong reasons to think that the distinction between health and disease is not intimately associated with this. It may well be that disease in human beings is not merely the same thing as pathology in nature generally, but a complete disconnection between the two does not seem plausible.

Once we examine the biological sciences from this perspective, it turns out there are at least four different types of situation that constitute a failure of biological normativity, and both of the popular naturalistic accounts of human disease miss at least some of these. By recognizing more ways of going wrong, we make it less likely that a biological constraint on what counts as a disease will be too restrictive.

III. FOUR WAYS OF GOING WRONG

We now turn to the four distinct senses in which a biologist might say that things have “gone wrong” for some organism. These are not necessarily diseases, and perhaps not even all pathological states, but they are all situations

where one organism does worse than another from a biological point of view. Following this, we will see how these ideas may be developed into a criterion for disease.

One Way of Going Wrong

Consider the db/db mouse. This strain of mouse has a genetic mutation that leads to faulty receptors for the hormone leptin. Leptin regulates a number of bodily processes, importantly including mediation of hunger and fat metabolism. Leptin action generates feelings of satiety, and so mice with faulty receptors eat to the point of gross obesity (Friedman, 2002).

In this example, some kind of evaluative or normative judgment seems apt. Something has gone wrong for these mice—they are not the way they ought to be. Exactly what underwrites such judgments is not always obvious in biology but in this case it is quite straightforward. A mechanism within these mice is broken: The leptin receptor does not have the right parts arranged in the right way. The mechanism is not able to do what it ought to do.

This way of going wrong is familiar from the literature on “selected effect” or “Proper function,” most prominently discussed by Ruth Millikan (1984; Neander, 1983, 1991). A biological structure fails to perform its function if it is unable to fulfill the causal role for which it has been selected in the recent evolutionary past (Godfrey-Smith, 1994). In the case of the db/db mouse, the leptin receptors lack the key property for which they have been selected, namely, responding to leptin. Many people would describe this situation as one of *dysfunction*. However, because we need to distinguish between different kinds of dysfunction, we will use the more specific term *mechanism failure*. Regardless of one’s position concerning the analysis of the concept(s) of function, mechanism failure is one important way in which things can go wrong from a biological point of view. But it is not the only way.

A Second Way of Going Wrong

Consider a European glow-worm (*Lampyris noctiluca*) living in urbanized areas, where there is a substantial amount of ambient light in the evening. Male glow-worms usually locate females by their light signature. However, it has been demonstrated that even quite dim light impairs their ability to discern female glow-worms signaling their availability. Instead, the male glow-worms simply miss these mating opportunities, oblivious to the presence of their conspecifics (Bird and Parker, 2014).

Once again, something seems to have gone wrong here: Male glow-worms are failing to locate prospective mates. However, no biological mechanism is faulty in this case. There is nothing wrong with either the mechanisms that produce light in the female abdomen or the male eye. The respective

mechanisms are working exactly as designed by natural selection, and this must therefore represent a different way of going wrong to the first example.

Following Millikan (1984) again, the issue here appears to be one of an *abNormal environment*: The mechanism is operating in accordance with its design but outside the operating parameters for that design. The mechanism that allows glow-worms to find mates was not designed to operate in environments with this level of ambient light. How are we to delineate the design parameters of a biological mechanism? Once again, Millikan supplies us with a useful (but not necessarily the only viable) framework with which to understand this idea. Given that a biological mechanism ought to do what it was selected to do in the evolutionary past, a “Normal” environment for the mechanism is one of the range of environments in which it was selected (Millikan, 1984). The ability of male glow-worms to detect the light signature of females was selected in a low-ambient light environment, making that the Normal environment for that trait. Since a glow-worm surrounded by light pollution is in an environment different to its Normal environment, it is operating outside its design parameters, and we can objectively say that something has gone wrong.

Therefore, as pointed out by Millikan and others, there are at least two ways in which the evolutionary past can ground normative judgments about phenotypes: Failure to operate as designed and failure to be in the Normal environment. These options might seem to exhaust the possible ways in which strictly natural norms can be founded on biological facts. However, this is not the case.

A Third Way of Going Wrong

Consider a common monkey flower (*Mimulus guttatus*) growing in its Normal environment. This species has evolved traits to deal with the variable circumstances that regularly arise in this environment. One such trait is early flowering. If the plant is growing in a poor enough location with respect to soil nutrients and water availability, its life span and growth potential will be limited, and the optimal age and size for reproduction is thereby reduced. This is not a good situation for the plant to be in. It will produce less seed than a conspecific that grows longer and larger before flowering. But for the less fortunate plant, that strategy would incur too great a risk of not reproducing at all. The plant’s developmental mechanisms are designed to make the most of a bad situation (Galloway, 1995).

Once again, something has gone wrong for this plant. However, this case does not fit either of our previous categories of natural normative failure. The plant’s systems are doing what they are designed to do, and they are doing it in exactly the setting that selected for those designs. The reason the plant has the facility to flower early is that the facility was selected in the past by this kind of environment. So, here we have an example of a selected

function deployed in its Normal setting. Something *else* has gone wrong for the plant. Although this normative assessment is not captured by the selected effects account, this does not mean that it is merely conventional. Rather, the assessment is still based on biological facts: The plant's fitness has been adversely affected. From a biological point of view, all else being equal, it is *better* for a plant if it is not forced to flower early.

This further normative category arises because "Normal environment" does not necessarily mean "hospitable environment." The biological world is a hostile place, and whether a species encountered some environment in its evolution is largely orthogonal to whether that environment is highly conducive to reproductive success.

So, "something can go wrong" not only due to mechanistic failure or abnormal environment but also due to a Normal but *inhospitable environment*. Inhospitability is clearly a matter of degree. However, the important point is that we are able to judge, on purely biological grounds, that things have gone *worse* or *better* for an organism relative to conspecifics and depending on its environmental circumstances. This is the case even when all the organisms are performing their selected function in some Normal environment. We have now distinguished three ways in which things can go wrong. However, there is yet another.

A Fourth Way of Going Wrong

Consider a water flea (*Daphnia cucullata*) born in an environment that was until recently full of predators. This organism has a special defensive phenotype: Water fleas can produce a tail spike and "helmet" that makes them more difficult to ingest as prey. However, spikes and helmets take extra resources to produce. If the flea does not need them, it is better off to not produce them. For this reason, *Daphnia* have evolved developmental plasticity in response to chemicals released by predators. Additionally, this effect is transgenerational. If a water flea's mother experienced high predator rates in her lifetime, her offspring will be born with readymade spikes and helmets (Agrawal, Laforsch, and Tollrian, 1999).

This makes excellent evolutionary sense. If the helmet phenotype could only be altered by genetic change, then fluctuations in levels of predation could not be accommodated except over evolutionary time scales. On the other hand, development requires appropriate timing, where a commitment to a particular phenotype must be made at the right moment. The solution, therefore, is to balance these two concerns: Obtaining information regarding the likely future and then committing to a certain developmental trajectory on that basis. Developmental switches of this kind are common in biology.

Things can go wrong, however, because the organism must initiate a particular developmental pathway based on less than perfect information.

Daphnia born during a transition from an environment with high numbers of predators to low numbers will be unnecessarily well equipped to deal with predators. It is always better to be in a situation of low predation than not, but nevertheless, this failure of the heuristic decision method will mean the offspring have paid a developmental cost they did not need to pay. These water fleas would have been better off if they had not prepared for predation.

In this case, we have a mechanism that is discharging the function for which it was designed, in just the type of environment that selected for that predictive mechanism. Furthermore, the environment in which the mechanism is operating is a benign one—the mother flea was in a Normal but inhospitable (predator-rich) setting, but her offspring is in a Normal and relatively hospitable (predator-free) setting. So, it appears that we have a new way in which things can go wrong for an organism: Not because of a faulty mechanism, or abnormality, or a harsh environment, but simply because sometimes developmental trajectories must be initiated in the setting of imperfect information. The best choice, given that information, may still turn out to be the wrong option. We call this way of going wrong *heuristic failure*.

The distinction between our third and fourth ways in which things can go wrong might not be totally apparent from this one example. Here is an alternative way to describe the difference. In the case of inhospitable environment, the relevant structure is operating perfectly well and according to its design, but the environment itself limits how well things can possibly go. So, the *realized* fitness of the organism may be maximized, *given the circumstances*: It is just that the circumstances are poor. In the case of heuristic failure, the relevant structure is also operating perfectly well and according to its design. However, while this succeeds in maximizing *expected* fitness, it fails to maximize *realized* fitness because even an optimal decision, given the available information, can sometimes lead to a bad actual outcome. This conclusion does not require the available information to be misleading; it merely has to be incomplete. Furthermore, in cases like this, the environment might only vary between thoroughly benign states. Two equally good developmental strategies may be available, each of which produces the same fitness if—but only if—the organism adopts the right strategy for its circumstances.

It is important to note that these two latter types of normative failure are not at all unusual. It is common for the appropriate behavior in a Normal setting to still result in a poor outcome. Being in an inhospitable patch of the environment is unfortunately quite common for many organisms. Similarly, the occasional predictive failure is all but inevitable in changeable environments: Many locusts must have adopted a phenotype suited for population-dense environments, only to find that no one else is preparing to swarm; human children born into a cold environment do not have the

optimal number of active sweat glands when they move to a hot climate (Bateson et al., 2004; Gluckman, Hanson, and Spencer, 2005). In these cases and many others, a “good bet” was made, given the information available, but it nevertheless turned out to be the wrong option.

Four Ways for Human Beings to Go Wrong

We have outlined four distinct ways in which doing justice to the biological facts seems to require a judgment that something has gone badly for an organism. Unsurprisingly, each of these can occur in human beings. Human beings can have faulty mechanisms. For example, improperly formed chloride ion gates mean that secretory cells cannot dilute mucus or certain digestive enzymes appropriately, leading to the distinctive symptoms of cystic fibrosis. Human beings can also be in situations for which they were not selected. All respiratory mechanisms are working as designed in the setting of carbon monoxide poisoning. It is just that human hemoglobin was never selected to effectively manage high levels of this gas, and so the selected function of oxygen delivery cannot be properly performed.

Human beings can be in Normal but inhospitable environments. It is part of our evolved history that nutrition is sometimes limited, and the classic pattern of fetal development when maternal food intake is poor—reduced growth in many organs, but spared cerebral growth—shows this situation has been common enough that we have evolved developmental plasticity to survive such hardship. Nevertheless, children born with intrauterine growth restriction and head-sparing are not doing as well as children who developed in a setting of good nutrition. These children have suffered from an inhospitable environment.

Finally, we know that human beings do make predictive phenotypic choices, and these can be subverted. We have already mentioned the case of sweat gland number, and an important possible instance of predictive failure features in the current medical literature regarding metabolic disease. Consider the famous case of people born in the Netherlands in the immediate aftermath of the 1944–1945 famine. Although their mothers experienced severe calorie restriction for a period during gestation, the children found themselves in a good nutritional environment. Unfortunately, as they grew into adulthood, an unusual proportion of these children developed obesity, diabetes, and cardiovascular disease. This has been described by some researchers as the possible result of a *predictive adaptive response*, or PAR (Gluckman, Beedle, and Hanson, 2009, 84–88; Low, Gluckman, and Hanson, 2012).

According to the PAR hypothesis, if it “appears” to a human fetus that its mother is not receiving adequate nutrition, its metabolism develops to be suited for future nutritional hardship. Alternatively, if it seems there is plenty

of food about during this early phase of life, human children develop a metabolism better suited to nutritional abundance. Once again, this makes good sense from a biological point of view. However, it has the result that children born during a transition from low nutrition to good nutrition will have a metabolism suited to a nutritionally poor environment but find themselves surrounded by calories. It is always better to be in a situation of adequate nutrition than not, but nevertheless, this failure of a usually very effective heuristic can lead to a number of metabolic problems. These children would have been better off if they had not prepared for famine, so this is a case of heuristic failure.

Only Four Ways?

We have shown that there are at least four ways in which a normative judgment is appropriate from a biological point of view and that each of these can occur in human beings. We have not given any reason so far to think that we have exhausted all of the options. However, we *can* say that these categories are not exclusive. For example, there is no reason why mechanism failure cannot occur in an abNormal environment. It is also certainly possible for heuristic failure to occur in an inhospitable environment. *Daphnia* offspring in an environment only recently colonized by predators will develop without helmets when they most need them.

Furthermore, it will often be illuminating to differentiate the various sources of normative failure in instances where things have gone wrong in more than one way. Consider a skin injury such as a cut. Let us assume for the moment that the body responds to this injury just as it was designed to, with the formation of a clot, the arrival of various inflammatory mediators, and so on. However, in spite of these response mechanisms functioning normally, it is almost inevitable that a scar will be left after such an injury. Scar tissue is not as effective as ordinary skin in discharging the functions of this organ, and so some residual mechanism failure will result. The worse the scarring, the worse the failure. What was the initial cause of this mechanism failure? Cuts are ubiquitous in our evolutionary history, so this is unlikely to be a case of abNormal environment. Rather, some patches of the environment are inhospitable—they contain more claws or thorns than other patches.

Here we have a case where one kind of normative failure leads to another kind downstream. Both of these are important, and it would be a mistake either to miss one due to a particular focus on the other or to lump them together. Additionally, it is not essential to the case that the person's skin was functioning correctly. There are ailments that lead to improper wound healing, such as poor vascular supply to the affected area. In that case, different ways of going wrong (inhospitable environment and mechanism failure) would be combined simultaneously as well as sequentially.

Two Sources of Biological Normativity

It is interesting to recognize how many different types of normative judgment are made in biology. However, we can also ask whether anything more general or systematic can be said about what underwrites these judgments, and our four examples show that there are at least two sources of natural normativity. One arises from the perspective of what has occurred in the evolutionary past. Natural selection has designed many biological traits to perform certain tasks in certain settings. From this perspective, failure to perform the task in the Normal environment or attempting to perform the task in an abNormal environment means that things are not going as they were designed to go. So, this source of normativity gives us two different ways in which things might go wrong.

However, sometimes biological traits function in accordance with their design in exactly the environment for which they were designed, and yet the outcome is a poor one. Significantly, we do not need to import any normative claims from outside biology to make this point, because evolution also has a forward-looking criterion of success—representation in future generations. As we have seen earlier, functioning as designed in the environment for which you were designed by no means guarantees this success. Sometimes the world is a harsh place, and sometimes organisms have bad luck in spite of utilizing the appropriate developmental heuristic.

Importantly, note that considerations of fitness alone would also be incomplete. We could have presented our examples in the other direction, starting with a forward-looking perspective and pointing to cases that would then be missed. For example, fitness can be high even in light of a structure failing to perform as it should. To use an old example, being severely short sighted is an abnormal state in virtue of mechanism failure, even when it prevents one from being conscripted in wartime.

These two sources of normativity are clearly closely related to (but not exactly the same as) the two views of dysfunction we discussed in Section 2 earlier. Wakefield uses a notion of Proper function in his account, while Boorse's view and its relatives are based on contribution to fitness. However, while these two "naturalist" views have standardly been considered exclusive options, we have shown that biological normativity combines both design and fitness considerations. Given this, it is no wonder the previous accounts have been too restrictive; each excludes an important source of how things might go wrong in biology.

Additionally, and essentially for our purposes, it is entirely possible to utilize both sources of naturalized normativity simultaneously, as the views are not contrary to one another. Indeed, given that both of them arise through the concerns of biological science, it would be artificial and simply unmotivated to restrict ourselves to just one. Either of these sources of normativity might underwrite a particular judgment regarding disease, but each is always important, and indeed as we have seen earlier, things might go wrong for two or more reasons at the same time.

For this reason, we claim that the necessary biological criterion for pathology is that the phenotype must constitute a failure of biological normativity, where this is understood as either a failure to discharge a selected effect or a lowering of fitness (or both).

IV. BIOLOGICAL NORMATIVITY AND DISEASE

Can the notion of biological normativity be usefully applied as a restriction on instances of human disease? It can, but it is not as simple as merely deploying the above-mentioned findings.

To begin with the positive features of this more inclusive view, it avoids many of the standard objections to earlier objective accounts of disease. This is simply because we are not concerned with asserting a particular kind of biological harm as a criterion rather than whatever kinds of harm are recognized in biology. So, for example, this criterion accommodates the case of long-standing epidemic disease in just the same way as standard selected effect accounts, while it also accommodates cases where nonselected structures are the focus of disease as long as this lowers fitness. Even if all human adults have atherosclerosis, this will still qualify as a case of biological harm through mechanism failure on our account. Alternatively, appendicitis meets our criterion regardless of concerns about the presence of dysfunction, as an inflamed appendix certainly affects one's fitness. So, the problem of false negatives—of failing to capture instances of genuine disease—looks to be mitigated.

What about cases of false positives? There would seem to be many examples of things going wrong from a purely biological point of view that are not considered disease states. For example, mechanism failure is the very point of elective vasectomy, but vasectomy would not normally be considered a disease state. This is a more complex issue. In the first instance, addressing cases like this may not appear to be such an important objective for us, since we are explicitly not attempting to supply an account of what is sufficient for disease to occur. If there are cases of biological normative failure in human beings that are not considered to be diseases, these may simply be cases where other important conditions are not met. Following Wakefield, one might adopt the view that for disease to occur there must be both a failure of biological normativity and some other more personally or socially determined harm. In cases where Proper dysfunction or loss of fitness is not seen as harmful by or to the individual, such an account would imply that these instances of dysfunction do not constitute disease states.

However, things are not necessarily as easy as that. If our biological constraint is to do real work in demarcating normal physiology from pathology, then it cannot be too weak. The biology must exclude a certain proportion of such potential cases or there would be little point in using a biological criterion at all.

As things stand, a constraint that incorporates mere fitness cost appears to set the bar too low. For instance, if a plant initiates flowering early due to drought, things have gone badly for it, but it is not necessarily in a pathological state. Similarly, human beings can have slightly reduced fitness for all sorts of reasons. A fetus receiving slightly less than optimal amounts of nutrition during gestation will have reduced fitness to some extent, but small babies are not (always) in a state of pathology. Things can go poorly for an organism from a biological stance, without that organism suffering from disease.

This is a concern for any account of pathology that incorporates fitness considerations, and there is a long history of attempts to address the problem. So, one option here is to simply adopt one of these previous attempted solutions. For example, Boorse's account of disease discussed earlier is based on fitness costs but restricts cases of disease to statistical outliers. It is only once fitness is restricted to a sufficient extent—once the baby is sufficiently small, for example—that we might consider a case to meet the threshold.

Approaching the problem in this way has a lot to recommend it. For starters, rather than needing to develop an entire new framework of disease-by-fitness-reduction, we can utilize a previously developed and well-respected account. At the same time, we do not need to import its difficulties, such as lacunae where a disease is statistically normal in a particular age-group, or the insistence that social normative facts cannot feature in the full account of disease. Perhaps there are better fitness-lowering accounts of pathology out there—some possibilities would be [Hausman \(2012\)](#) and [Garson and Piccinini \(2014\)](#). This is also fine by us: We are not wedded to any of these accounts in particular as long as they are appropriately restrictive.

A related issue here is the “line-drawing” problem. Loss of fitness comes in degrees, and it seems there is no obvious principled biological criterion to delineate how costly a state must be in order to qualify as pathology. For example, it is not clear exactly how severe a restriction on cardiac output needs to be before it is considered pathological ([Schwartz, 2007](#)). This would appear to cause difficulties for the idea that fitness cost generates a purely objective criterion for disease, simply dictated by the biological world.

There are a number of points to be made here. First, the need to impose a cutoff regarding a graded variable does not necessarily make a criterion merely evaluative, unless we wish to call all properties that exhibit borderline cases evaluative. We know lowered fitness is a quantity that determines pathology, and we know that cases of very severe fitness loss will be pathological. The fact that there may be some vague boundaries where pathological and nonpathological fade into one another does not automatically make the distinction a socially determined one ([Garson and Piccinini, 2014](#), footnote 15).

Additionally, we should not be at all surprised to find that a biological account of harm will generate such middling cases. Biology is a domain

that usually does not lend itself to binary distinctions, and even categories such as “species” or “predator” involve vagueness at their borders. So, the point that such vagueness occurs regarding the boundaries of pathology is merely a reflection of the biological reality. If the use of established biological notions precludes a sharp binary distinction between disease and health, so much the worse for the idea that disease has a sharp binary distinction, even if this entails some revision of the concept.

We close this section with two problem cases, both of which build on the possible concern that our restriction is not restrictive enough. A desirable output we might hope to gain from establishing a necessary biological criterion for disease is the distinction between the so-called risk factors and genuine diseases. Moderately raised cholesterol is not standardly considered a disease, but its potential downstream results are (Walker and Rogers, forthcoming). Precancerous change is a risk for later malignant disease. However, if we include any state that lowers prospective fitness as a candidate disease, raised cholesterol and carcinoma in situ may make the grade, as they do restrict fitness in virtue of increasing one’s probability of developing other conditions. People in these states therefore have worse fitness prospects than they would otherwise have.

There is a variety of replies that might be made to this. Here, we outline some possible responses without necessarily endorsing any of them. First, one or more of the additional restrictions discussed earlier could be employed. For example, most risk factors are likely to lower fitness to a relatively minor extent: Less than the diseases they are a risk for, at least. In this case, a requirement that fitness costs are substantial may track the standard distinction between risk factor and disease reasonably closely. Alternatively, we could rely on further restrictions through a different criterion, such as a judgment of harm. People with raised cholesterol are asymptomatic until this leads to some other state that causes them recognizable harm, at which point we would judge them to have developed a disease.

Furthermore, if it turns out that some putative risk factor directly causes harm or imposes high fitness costs (for example, perhaps extremely raised cholesterol), it becomes less likely that medicine will distinguish such “risk factors” from “genuine diseases.” If it is the case that severe hypercholesterolemia lowers prospective fitness very severely, this represents a significant failure of biological normativity, and one might wonder what purpose it would serve to distinguish this from a disease state.

The second problem case we wish to consider is one where a socially sanctioned phenotype results—solely via that sanction—in lowered fitness. A classic thought experiment we can use to illustrate the idea (standardly used to question the attribution of heredity to genes or environment; see Jencks, 1972) is a social setting that marginalizes those with red hair. Let us imagine this is because society has decided that having red hair is a disease. This is precisely the kind of case that cries out for help from a biological criterion

to avoid seemingly arbitrary socially designated disease states. However, if society treats red-haired people poorly enough on that basis, this might result in their having worse biological prospects than they would otherwise have had, and so this would meet the criterion of a failure of biological normativity.³ Such an interaction between trait and environment may even make this “disease” seem genetically heritable (Lynch and Kemp, 2014).

This returns us to the core question we mentioned at the start of this article: The respective roles of social and biological considerations in determining disease states. In fact, here we may see at least part of the reason why discussion of these issues has proven to be so intractable. It is true that there is a purely biological notion of normativity which applies to all organisms, including us. However, it is also true that sometimes biological outcomes are affected by social considerations. We are not the only species whose environment is affected by social and constructed elements, but this occurs to an extreme and unusual extent in human beings. It appears likely that human manipulation and construction of our social and environmental niche is a (or perhaps even the) defining feature of the human life course (Sterelny, 2003, 2012).

Further, negotiating this social environment will certainly affect one’s biological prospects. If one is unable to pick up on the subtleties of social mores, one is less likely to do as well in human society, and this may manifest in ways visible to biology. Given this, even if considerations of our biological evolution (both forward- and backward-looking) are essential to understanding human disease, this does not necessarily exclude the effects of our social environment. In our species, the biological and social are not easily separable, because the social environment is a key feature of our biological fitness.

However, it is important not to overstate this point. It is true that we are unusual in the way in which our biological outcomes are shaped by social factors. From this, we see that discussion of human disease is complicated in ways that would be missed, were we to consider ourselves merely from a narrow biological point of view. But this acknowledgement does not refute the idea that biology places boundaries on what gets to count as a disease in human beings. We cannot arbitrarily stipulate that some state is a disease regardless of biological considerations. Human disease requires some failure of a structure to perform its Proper function or for the person to suffer some cost in fitness. To deny this would be tantamount to declaring human beings unlike all other living things with respect to pathology.

V. CONCLUSION

We have argued that attribution of disease ought to be restricted to cases where there is a failure of biological normativity. This is not to say that

such a criterion is sufficient for disease to occur; just that disease cannot be solely a matter of social convention. Our argument for this is based on the fact that notions of normality, abnormality, pathology, and physiology are essential to understanding the biology of living things, even in cases where human values play no role whatsoever. This can be seen through the consideration of situations where things have gone wrong for an organism, for either of two different reasons. A mechanism can fail to perform the function for which it has been selected (due to mechanism failure or because of an abnormal environment) or biological fitness can be reduced (due to an inhospitable environment or the failure of a heuristic). Although the literature regarding disease concepts and biological functions in general has often considered these notions to be in tension, we claim that both are important and do not compete. This explains why previous accounts of biological normativity that utilize only one of these considerations have proven to be too restrictive.

As well as establishing a criterion that tracks biological normativity in a way that is consonant with nonmedical biology, this approach may also help cast light on some further issues in the philosophy of medicine. For example, it addresses concerns regarding a lack of a sharp distinction between what is normal and pathological, since this is merely a reflection of how these categories intersect in the biological world. Additionally, it may provide further resources to help develop an account of why human health seems a unique case, in light of the extent to which our social and biological worlds are entangled. In the future, this approach may add to discussion regarding the roles society and biology play in human health and disease while avoiding a naive dichotomy between these two domains.

NOTES

1. For an introduction to this extensive topic, see [Murphy \(2015\)](#). A recent special issue of this journal gives a good picture of the current state of play ([Schramme, 2014](#)).
2. An anonymous referee has pointed out that disease might be defined by deontological norms. For example, a Kantian might hold that diseases are states whose medicalisation an agent could consistently will as a universal maxim. We do not explore this view here because Glackin, Canguilhem, and other normativists do not hold this view, and we aim to counter their criticisms of “naturalist” accounts of disease.
3. Thank you to Bill Fish and Steven Chadwick for raising this type of case with us.

ACKNOWLEDGEMENTS

Sincere thanks to Mary Walker and Wendy Rogers for their support and patience. Thanks also for feedback on an earlier draft from Brian Hedden, James Maclaurin, and Nicholas Shea. Versions of this paper were presented at: The Macquarie “Defining the Boundaries of Disease” workshop, PBDB 8, NZAP 2015 and the Obesity Australia 5th Annual Summit 2016. Audience discussion was very helpful to us in all cases. This research was supported under Australian Research Council’s Discovery Projects funding scheme (project number DP130101774).

REFERENCES

- Agrawal, A. A., C. Laforsch, and R. Tollrian. 1999. Transgenerational induction of defences in animals and plants. *Nature* 401:60–3.
- Bateson, P., D. Barker, T. Clutton-Brock, D. Deb, B. D'Udine, R. A. Foley, P. Gluckman, et al. 2004. Developmental plasticity and human health. *Nature* 430:419–21.
- Bird, S., and J. Parker. 2014. Low levels of light pollution may block the ability of male glow-worms (*Lampyrus noctiluca* L.) to locate females. *Journal of Insect Conservation* 18:737–43.
- Boorse, C. 1977. Health as a theoretical concept. *Philosophy of Science* 44:542–74.
- . 1997. A rebuttal on health. In *What Is Disease?*, eds. J. M. Humber and R. F. Almeder, 1–134. Totowa, NJ: Humana Press.
- . 2014. A second rebuttal on health. *Journal of Medicine and Philosophy* 39:683–724.
- Canguilhem, G. 1991. *The Normal and the Pathological*. Brooklyn, NY: Zone Books.
- Carnap, R. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Friedman, J. M. 2002. The function of leptin in nutrition, weight, and physiology. *Nutrition Reviews* 60:S1–14.
- Galloway, L. F. 1995. Response to natural environmental heterogeneity: Maternal effects and selection on life-history characters and plasticities in *Mimulus guttatus*. *Evolution* 46:1095–107.
- Garson, J., and G. Piccinini. 2014. Functions must be performed at appropriate rates in appropriate situations. *British Journal for the Philosophy of Science* 65:1–20.
- Glackin, S. N. 2010. Tolerance and illness: The politics of medical and psychiatric classification. *Journal of Medicine and Philosophy* 35:449–65.
- Gluckman, P., A. Beedle, and M. Hanson. 2009. *Principles of Evolutionary Medicine*. New York: Oxford University Press.
- Gluckman, P. D., M. A. Hanson, and H. G. Spencer. 2005. Predictive adaptive responses and human evolution. *Trends in Ecology & Evolution* 20:527–33.
- Godfrey-Smith, Peter. 1993. Functions: Consensus without unity. *Pacific Philosophical Quarterly* 74:196–208.
- . 1994. A modern history theory of functions. *Noûs* 28:344–62.
- Griffiths, P. E. 1993. Functional analysis and proper function. *British Journal for the Philosophy of Science* 44:409–22.
- . 2009. In what sense does 'nothing in biology make sense except in the light of evolution'? *Acta Biotheoretica* 57:11–32.
- Griffiths, P. E., and J. Matthewson. 2016. Evolution, dysfunction and disease: A reappraisal. *British Journal for the Philosophy of Science* [On-line]. Available: <https://doi.org/10.1093/bjps/axw021> (accessed February 5, 2017).
- Haldane, J. B. S. 1949. Disease and evolution. In *Evolution*, ed. M. Ridley, 41–7. Oxford, United Kingdom: Oxford University Press.
- Hausman, D. 2012. Health, naturalism, and functional efficiency. *Philosophy of Science* 79:519–41.
- Jencks, C. 1972. *Inequality: A Reassessment of the Effect of Family and Schooling in America*. New York: Basic Books.
- Kingma, Elselij. 2009. Paracetamol, poison, and polio: Why Boorse's account of function fails to distinguish health and disease. *British Journal for the Philosophy of Science* 61:241–64.
- . 2013. Naturalist accounts of mental disorder. In *The Oxford Handbook of Philosophy and Psychiatry*, ed K. W. M. Fulford. Oxford, United Kingdom: Oxford University Press.

- Lemoine, M. 2013. Defining disease beyond conceptual analysis: An analysis of conceptual analysis in philosophy of medicine. *Theoretical Medicine and Bioethics* 34:309–25.
- Low, F. M., P. D. Gluckman, and M. A. Hanson. 2012. Developmental plasticity, epigenetics and human health. *Evolutionary Biology* 39:650–655.
- Lynch, K. E., and D. J. Kemp. 2014. Nature-via-nurture and unravelling causality in evolutionary genetics. *Trends in Ecology & Evolution* 29:2–4.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- Murphy, D. 2015. Concepts of disease and health. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta [On-line]. Available: <http://plato.stanford.edu/archives/spr2015/entries/health-disease/> (accessed February 5, 2017).
- Murphy, D., and R. L. Woolfolk. 2000. The harmful dysfunction analysis of mental disorder. *Philosophy, Psychiatry, & Psychology* 7:241–52.
- Neander, K. 1983. *Abnormal Psychobiology. Doctoral dissertation*, La Trobe University, Bundoora, Australia.
- . 1991. Functions as selected effects: The conceptual analyst's defense. *Philosophy of Science* 58:168–84.
- . 1998. Mental illness, concept of. *Routledge Encyclopedia of Philosophy* [On-line]. Available: <https://www.rep.routledge.com/articles/thematic/mental-illness-concept-of/v-1> (accessed February 5, 2017).
- Nishide, Y., and S. Tanaka. 2013. Wing dimorphism in the migratory locust, *Locusta migratoria*: Differentiation of wing morph and phase polyphenism. *Entomological Science* 4:421–31.
- Pittendrigh, C. S. 1958. Adaptation, natural selection and behavior. In *Behavior and Evolution*, eds. A. Roe and G. G. Simpson, 390–416. New Haven, CT: Yale University Press.
- Schramme, T. 2014. Christopher Boorse and the philosophy of medicine. *Journal of Medicine and Philosophy* 39:565–71.
- Schwartz, P. H. 2004. An alternative to conceptual analysis in the function debate: Introduction. *The Monist* 87:136–53.
- . 2007. Defining dysfunction: Natural selection, design, and drawing a line. *Philosophy of Science* 74:364–85.
- Sterelny, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Malden, MA: Blackwell Publishing.
- . 2012. *The Evolved Apprentice*. Cambridge, MA: MIT Press.
- Wakefield, J. 1992. The concept of mental disorder: On the boundary between biological facts and social values. *American Psychologist* 47:373–88.
- . 2000. Spandrels, vestigial organs, and such: Reply to Murphy and Woolfolk's 'The Harmful Dysfunction Analysis of Mental Disorder.' *Philosophy, Psychiatry, & Psychology* 7:253–69.
- Walker, M., and W. Rogers. 2017. Introduction: The Boundaries of Disease. *Journal of Medicine and Philosophy* 42:343–9.