The British Journal for the Philosophy of Science Advance Access published March 10, 2013

Brit. J. Phil. Sci. **0** (2013), 1–20

Functions Must Be Performed at Appropriate Rates in Appropriate Situations

Justin Garson and Gualtiero Piccinini

ABSTRACT

We sketch a novel and improved version of Boorse's biostatistical theory of functions. Roughly, our theory maintains that (i) functions are non-negligible contributions to survival or inclusive fitness (when a trait contributes to survival or inclusive fitness); (ii) situations appropriate for the performance of a function are typical situations in which a trait contributes to survival or inclusive fitness; (iii) appropriate rates of functioning are rates that make adequate contributions to survival or inclusive fitness (in situations appropriate for the performance of that function); and (iv) dysfunction is the inability to perform a function at an appropriate rate in appropriate situations. Based on our theory, we sketch solutions to three problems that have afflicted Boorse's theory of function, namely, Kingma's ([2010]) problem of the situation-specificity of functions, the problem of multi-functional traits, and the problem of how to distinguish between appropriate and inappropriate rates of functioning.

1 Functions Are Situation-Specific

In this article, we briefly introduce a theory of function that builds on Boorse's biostatistical theory (BST) and accounts naturalistically for the fact that functions must be performed at appropriate rates in appropriate situations. As we

doi:10.1093/bjps/axs041

shall argue, our theory meets a challenge raised by Kingma ([2010]) to naturalistic theories of function.[1]

Kingma stresses that functions are 'situation-specific'. The function of the stomach is to digest food, but sometimes the stomach does not digest food. This may be because of a disease such as gastroparesis, which prevents the stomach from emptying its contents normally, or simply because the animal has not eaten recently and there is no food to digest. The function of the myelin sheath is to facilitate the transmission of action potentials, but it does not always do so. This may be because a disorder such as multiple sclerosis has resulted in myelin damage, or simply because the neuron is at rest. The function of the immune system is to fight infection, but sometimes it does not do so. This may be because of a disease such as leukopenia, which reduces white blood cell count, or simply because there is no infection to fight.

These examples highlight the relationship between functions and dispositions. A trait that has function *F* and is capable of performing *F* behaves differently depending on the situation. There are situations in which it is entirely appropriate for the trait to not perform *F*. Thus, being able to perform a function is a disposition—it is the disposition to perform *F* under appropriate circumstances.[2] Like other dispositions, the disposition to perform *F* is manifested under some conditions and not others. By the same token, the mere fact that a trait (with function *F*) is not performing *F* does not imply that the trait is dysfunctional. What matters is whether the trait is capable of performing *F*, and thus whether it would perform *F* under appropriate circumstances. Being dysfunctional is the same as lacking a disposition; it is the lacking of the disposition to perform a function (under appropriate circumstances). For instance, if someone has been fasting, their stomach is not digesting food, but it is not necessarily dysfunctional. In the absence of any long-term impairment (perhaps wrought by excessive fasting), if food were present, the stomach would digest food. As long as the stomach is still able to digest, it is not dysfunctional.

Given that the appropriate performance of a function is relative to a situation, Kingma constructs a dilemma for Boorse's theory of health (for example, Boorse [1975], [1976], [1977], [1997], [2002]). Roughly, according to Boorse's theory, the physiological function of a trait consists in its

[1] In responding to Kingma, we are skirting a substantive and ongoing debate about whether 'function' should be given a unique explication or multiple ones. We are pretty pluralistic about 'function' (see Garson [2011], [2012] for an alternate approach to function). We formulate here a biostatistical theory of function to show that such a view can explicate the notion that functions must be performed at appropriate rates in appropriate situations.

Kingma is interested in theories of health and disease that rely on the notion of biological function; we will set health and disease to one side and simply focus on theories of function.

[2] This is different from the claim that having a function is a disposition. We leave open the possibility that a trait has a function but lacks the disposition to perform that function.

species-typical contribution to survival and reproduction (relative to an appropriate reference class). Disease represents the impairment of function below normal efficiency and health the absence of disease.[3]

As Kingma points out, however, this definition of function does not reflect the situation-specificity of function. Thus, she proposes relativizing functions to situations in the following way: the function of a trait, relative to situation *s*, consists in its statistically typical contribution to survival and reproduction in *s* ([2010], p. 248). The problem is that this way of relativizing functions to situations undermines normal judgments of function and dysfunction. For example, following paracetamol overdose, liver function is extremely low. Nonetheless, that low level of liver function is statistically typical for the situation of paracetamol overdose. Hence, given Kingma's way of relativizing functions to situations, it is not dysfunctional relative to that situation ([2010], p. 251). Kingma argues that BST cannot escape her dilemma. The only way out, she maintains, is to reject naturalistic accounts of function.[4]

Although we agree that functions must be relativized to situations, we disagree with the manner in which Kingma does so, and hence with her conclusion. In order to determine whether a trait is functional or dysfunctional with respect to an activity, we must first identify a set of appropriate situations for the performance of that function. (For example, an appropriate situation for the performance of the stomach's digestive function is one in which there is food in the stomach.) Then we simply ask, if the trait were in that situation, would it be able to carry out its function? If not, then it is dysfunctional—even if its activity is typical for the kind of situation it happens to find itself in.

In a response to Kingma ([2010]), Hausman ([2011]) introduces a distinction between functioning normally with respect to a relevant situation—that is, functioning close to the average rate (over a representative sample of the population) in a given relevant situation—and having a normal functional capacity—that is, being able to function normally in all relevant situations. Consider a trait (for example, the heart) whose rate of functioning is low because of normal functioning (for example, the organism is at rest) versus because of a dysfunction (for example, heart disease). Hausman points out

3 See (Boorse [1977], p. 555; [1997], pp. 7–8). The concept of 'goal-directedness' embedded in his theory of function will be explicated below. When referring to his theory of function in isolation from his theory of health, Boorse ([2002]) refers to it as the 'general goal contribution' theory of function. We find this unhelpful for reasons to be explained below (see Section 2.1).

4 Boorse ([1997], pp. 83–4) himself addresses this problem. His solution appeals to the notion of a 'normal environment'. A heart rate of 150 bpm is non-pathological if it results from vigorous exercise; it is pathological if it results from the ingestion of belladonna. The difference, he seems to hold, is that the former is a relatively typical situation for the species and the latter is not. This solution seems to presuppose that environmental pathogens must be in some sense 'atypical', and that seems wrong. The malaria parasite *Plasmodium falciparum* is in no sense an 'atypical' part of the human environment, and the contraction of malaria is in no sense an 'atypical' human experience. Kingma ([2010], pp. 255–6) discusses this passage and reaches a similar conclusion.

that in the case of normal functioning, a trait, in its current state, is still able to respond normally to other situations (for example, increased activity by the organism). In other words, the trait still has its normal functional capacity. In contrast, in the case of dysfunction, a trait, in its current state, is unable to respond normally to other situations. It has lost its normal functional capacity. Because of this difference, Hausman argues, BST can distinguish between normal functioning and dysfunction, thereby escaping Kingma's dilemma.

Hausman is on the right track but his account deserves two refinements. First, his account of normal rate of functioning depends on the notion of a relevant situation, which Hausman does not explicitly account for. Because of this, he is willing to concede that an impaired liver is still functioning normally with respect to the situation of paracetamol overdose ([2011], p. 660). The only problem, in his view, is that the liver is unable to function normally in other relevant situations. This response concedes too much. A liver impaired by paracetamol overdose is not functioning normally in any sense of the term. Thus, we respond to Kingma by providing a method for constructing a (relatively constrained) set of appropriate situations for the performance of a function and specifying the rate of functioning appropriate for each situation in the set.

Second, Hausman's account of rate of functioning (which he calls 'level of functioning') is problematic because it does not avoid the classic problem of (near-)pandemic diseases (for example, Neander [1991], p. 182; Boorse [2002], p. 94). Hausman states that a level of functioning is normal in a situation if its contribution to survival and reproduction 'is not too much less than the average efficiency with which this part would function in a large representative sample of human beings if that sample were in situation *S*' ([2011], p. 659). According to this definition, a near-pandemic disease such as dental caries, a global outbreak of H1N1 flu, or an epidemic of anuran declines due to UV radiation would not represent dysfunctions. After all, a trait affected by a pandemic disease does not perform much less than average.

In Section 2, we present our version of BST. Specifically, we explicate the related notions of function, appropriate situation for the performance of a function, appropriate rate of functioning (relative to an appropriate situation), and dysfunction. Our account does justice to the spirit of Boorse's BST and also clarifies, amends, and extends previous versions of BST on some crucial points. In Section 3, we show how the problem Kingma raises can be dissolved in a way that incorporates what is right about Hausman's response.

2 A General Account of Biostatistical Functions

We now introduce the relevant aspects of a theory of function that builds on Boorse's and Kingma's contributions in order to handle the dispositionality of functions. We explicate the concepts of function, appropriate set of situations for the performance of a function, appropriate rate of functioning (relative to an appropriate situation), and dysfunction.

2.1 Function

Boorse's theory of function can be looked at from two main perspectives. In the most general sense of 'function'—the sense appropriate for both biological organs and artifacts—the function of an entity is its contribution to a goal of the system in which it is contained. In turn, the goal of an entity is that toward which it is 'directively organized' in the sense promoted by theorists such as Sommerhoff ([1950]) and Nagel ([1977]). We will not discuss the merits of this account (see Garson [2008], pp. 539–41).

For Boorse, more specific concepts of function are generated by focusing on specific goals within specific contexts. For artifacts, the goals may be determined by the subjective preferences of the artifact's user. In contrast, in the context of physiology, which is the basis of medicine and therefore the appropriate context for developing a theory of function relevant to health, the appropriate or 'apical' goals of the organism are survival and reproduction (for example, Boorse [1977], p. 556; [1976], p. 84; [2002], p. 76). Correspondingly, for Boorse (with revisions), the physiological function of a biological trait is an activity of that trait that typically contributes to the survival or inclusive fitness of organisms within a reference class, usually characterized as a subset of a species partitioned by age and sex (cf. Boorse [1997], pp. 7–8).[5] We will now elaborate on this basic idea.

An apparent shortcoming of the present approach is that there are cases in which something has a function but performing that function does not represent a species-typical contribution to survival or inclusive fitness. A function of the visual word form area (VWFA) in the brain is to recognize words that correspond to the spelling rules of one's language, yet it is not clear that this ability contributes to survival or inclusive fitness.

There are two ways of handling these cases within a BST framework. The first is that the function of the VWFA to recognize words that correspond to the spelling rules of one's language is not a physiological function. The VWFA has that function because recognizing words promotes reading, and reading is

[5] Boorse writes 'survival and reproduction'. But as Schwartz ([2007], p. 372) points out, a trait may contribute to only one among survival and reproduction, as opposed to both. Because of this, we should follow Schwartz and write 'survival *or* reproduction' (inclusive 'or'). Except that, as Corey Maley has pointed out to us, this account fails for traits such as the stingers of bees. The function of stingers appears to be to deliver a sting. Yet far from leading to either survival or reproduction, a sting leads to the bee's death. But a sting may contribute to the survival of other members of the group, that is, to the bee's inclusive fitness. Because of this, we write 'survival or inclusive fitness'.

a subjective goal or preference of the individual, rather than a physiological goal. By the same token, we might say that a function of the nose is to hold up glasses, but it is not a physiological function. Rather, this function derives from the subjective preferences or goals of the individual. The second way to handle these cases is to say that the VWFA has the function of recognizing words that correspond to the spelling rules of one's language but this truth is parasitical on the fact that one function of the brain is to recognize correlations in the environment, and the VWFA's ability to recognize words constitutes one way of recognizing correlations in the environment. Because recognizing words that conform to the spelling rules of one's own language is a way of recognizing correlations in one's environment, it constitutes a function, whether or not recognizing such words *per se* typically contributes to survival or inclusive fitness.

To reiterate, for Boorse, the function (in the physiological sense) of a biological trait is an activity of that trait that typically contributes to the survival or inclusive fitness of organisms within a reference class. What is crucial is that the term 'typically' denotes a conditional rather than an unconditional probability. The function of a trait is neither what it typically does nor what it typically does that happens to benefit the organism. The function of a trait is its typical contribution to survival or inclusive fitness on those occasions in which it does contribute to the survival or inclusive fitness of the organism (Boorse [2002], p. 93).[6]

To use a common example, take sperm (for example, Millikan [1984], p. 29). The (or at least one) function of fertile sperm is to fertilize the ova. But fertile sperm typically swim about in the testes and do not contribute to survival or inclusive fitness in any direct way. This is not a problem for BST. Even though the unconditional probability of a sperm's fertilizing an ovum is low, the conditional probability that a sperm fertilizes an ovum, given that it is contributing to survival or inclusive fitness, is extremely high. That is, on those occasions in which the sperm contribute to survival or inclusive fitness, it is very likely that they do so by fertilizing the ova.

Formally, we can define the function of a trait in terms of a conditional probability. If X is a trait, Y is an activity of that trait, and RC is the reference class, then:

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.

[6] Below, we amend this thesis by identifying the function of a trait with its non-negligible contribution to survival or inclusive fitness on those occasions on which it does so contribute.

Why does condition (i) say that the probability that the trait contributes to survival or inclusive fitness is non-negligible? There are cases in which a trait may contribute to survival or inclusive fitness but it does so rarely enough that it would be considered a fortuitous benefit rather than a function. Why does condition (ii) specify that this conditional probability must be non-negligible, rather than, say, that this probability must be high? In paradigmatic cases of function, such as the heart's circulating blood or the sperm's fertilizing ova, this probability is very high. In other cases, a trait has a function but this probability is low.[7] This gives rise to the 'problem of multi-functional traits'.

The medulla oblongata in the brainstem has several functions that are commonly performed on those occasions in which it contributes to survival, such as regulating breathing, circulation, and blood pressure. It also has several functions that are rarely performed on those occasions in which it contributes to survival, but which are crucial to the well-being of the organism. They include initiating vomiting and the gag reflex (pharyngeal reflex). On those occasions in which the medulla contributes to survival, it is overwhelmingly probable that it does so by contributing to breathing. Yet the probability that the medulla's initiating gagging contributes to survival, given that it is doing something to contribute to survival, is non-negligible albeit low. Thus, it must be accorded a function.[8] Our definition also avoids the problem of pandemic diseases because it does not entail that the activity by which a trait contributes to survival or inclusive fitness need be, in any sense, typical—more on this below.[9]

The notion of a trait's making a contribution to, for example, survival by performing a specific activity is intended to discriminate between those activities that causally contribute to survival and those activities that represent a byproduct of that contribution. The function of hemoglobin is to transport oxygen, because that is an activity that contributes non-negligibly to survival

[7] This observation is due to Karen Neander (cf. Boorse [2002], pp. 92–3).

[8] This solution avoids the problem of 'fortuitous benefits': how low may this probability be before the activity in question would no longer constitute a function? As an anonymous referee asks, does the tiny pebble in a mouse's foot that reflects light that distracts a cat and saves the mouse's life have the function of distracting cats? No. Suppose that in a large population of mice ($N =$ 10,000), each mouse, on at least one hundred occasions, uses its feet to scurry away from a predator, but one mouse in that population, on one occasion, uses its foot to catch a stone that reflects light into a predator's eyes, thus avoiding predation. Then, the probability that the foot contributes to survival or inclusive fitness by catching a stone, given that it contributes to survival or inclusive fitness, is 1/1,000,001 or 0.0000009, which we would consider negligible. There does not seem to be a principled way to draw a line between an activity that constitutes a function of a trait and a fortuitous benefit of that trait, however, and we are content to leave this boundary vague (cf. Boorse [2002], p. 71).

[9] Someone might worry that the very idea of a trait, *X*, making a contribution to survival presupposes an alternate set of traits in relation to which *X*'s contribution is superior. But all that it presupposes is that *X*'s doing *Y* plays a prominent role in the organism's ability to survive (cf. Buller [1999], pp. 285–6). The concept of an 'appropriate rate of functioning' is relative to an alternate set of traits, however, as will be discussed below.

or inclusive fitness. A byproduct of this contribution is that hemoglobin makes the blood red. Making the blood red does not itself contribute to survival. Thus, only the transportation of oxygen is a function of hemoglobin.[10]

2.2 Appropriate situations for the performance of a function

As we've seen, just because a trait is not performing its function, it doesn't follow that the trait is dysfunctional. The performance of the function may simply be uncalled for. If my stomach is not digesting food solely because today I haven't eaten anything yet, my stomach is still perfectly capable of digesting food and hence it is not dysfunctional.

When a trait is dysfunctional at *t*, not only is the trait not performing its function at time *t*, but also the trait is incapable of performing its function at *t*, so even if the performance of that function were appropriate, the trait would not perform the function. Because of this, an account of dysfunction requires an account of which situations are appropriate and which are inappropriate for the performance of a function. Though Boorse recognizes that functions have appropriate conditions for their performance, he does not elaborate this point or develop criteria for identifying these situations (Boorse [1997], p. 79).

To identify situations as appropriate or inappropriate for the exercise of a function, we can use conditional probability again. 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 (i.e. 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. 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: $P(X \text{ is in } S|X\text{'s doing } Y \text{ contributes to survival or inclusive fitness}) \approx 1$.[11]

(ii) Specificity: there is no S' which is a proper subset of *S* such that (i) is true of S'.

For example, suppose that *X* is the medulla oblongata, *Y* is the activity of initiating the gag reflex, and *S* is the situation-type asphyxiation due to an obstruction in the throat. On any given occasion on which the medulla's initiating the gag reflex contributes to survival or inclusive fitness, the

[10] We thank Hayley Clatterbuck for drawing our attention to this potential problem.

[11] We use '$\approx$' rather than '=' to mean that the probability is very close to 1. This is meant to exclude 'fortuitous benefits' of *X* doing *Y*, which may occur in situations other than *S* but are so rare as to be negligible. Such fortuitously beneficial situations are excluded from consideration when evaluating both (i) and (ii) (cf. Footnote 8). Note also that the expression '*X* is in *S*' does not refer to set membership but refers to the fact that *X* happens to find itself in that situation.

probability that the medulla is in the situation of asphyxiation due to an obstructed throat is equal or very close to one. Moreover, there is no proper subset of *S* of which that is true. Therefore, *S* consists of the situation of asphyxiation due to an obstructed throat.

The rationale for conditions (i) and (ii) is to avoid two related but opposite errors that might arise in describing the set of appropriate situations, namely, that in trying to identify the set of appropriate situations, *S*, one may mistakenly identify an irrelevant (proper) subset of *S* or an irrelevant (proper) superset of *S*. Intuitively, the situation of asphyxiation due to an obstructed throat is an appropriate one for initiating of the gag reflex. Any proper subset of *S*—for example, *S′*, the situation of asphyxiation on a weekday—would leave out some relevant situations.[12] Condition (i) excludes inappropriately narrow subsets such as *S′* because the probability that the medulla is in *S′*, given that the medulla's initiating gagging contributes to survival or inclusive fitness, is only 5/7 or 0.71—assuming there is no correlation between one's chances of asphyxiating and the day of the week. Condition (i) ensures that we attempt to identify the most inclusive set of situations that are appropriate for the gag reflex.

Condition (i) alone is not enough because it allows for the opposite error, namely, the problem of irrelevant (proper) supersets. Let *S* be a set of situations that satisfies (i). Then, trivially, the set described by the union of *S* and *not-S* (the complement of *S*) also satisfies (i), because $P(X \text{ is in } S \cup not\text{-}S) = 1 \geq P(X \text{ is in } S)$. Condition (ii) disqualifies such irrelevant supersets, as long as *S* is a proper subset of $S \cup not\text{-}S$. While the set $S \cup not\text{-}S$ satisfies (i), it does not satisfy (ii) because there is a more specific subset of $S \cup not\text{-}S$ for which (i) is true. Thus, identifying the appropriate set of situations for the performance of a function requires identifying the set that is both inclusive as well as specific.

By the same reasoning, condition (ii) resolves what may be called the 'problem of irrelevant disjunctions', which is a version of the problem of irrelevant supersets.[13] By (i), if *S* is an appropriate situation for *X*'s doing *Y*, then *S* or *T* is an appropriate situation for *X*'s doing *Y*, for any situation-type *T*. For example, if asphyxiation is an appropriate situation for the medulla's initiating the gag reflex, then the situation of asphyxiation or seeing a meteor would be an appropriate situation. But if *S* is a proper subset of $S \cup T$, the situation-type *S or T* would violate condition (ii). Thus, (ii) rules out irrelevant disjunctions.

[12] A similar example would be that a situation appropriate for the processing of visual information by the eyes is a situation involving the presence of light (appropriate situation), rather than a situation of being, say, in bright daylight (irrelevant subset). We are grateful to an anonymous referee for bringing this to our attention.

[13] We thank an anonymous referee for drawing our attention to this.

2.3 Appropriate rates of functioning

The above account would be adequate if a trait were always either fully functional or completely dysfunctional. On the contrary, many traits have different rates of functioning, some of which are at the lower end of their normal spectrum, while others are dysfunctional. As Schwartz ([2007]) points out, function theorists have not traditionally drawn a principled line between rates of functioning that fall in the normal range and rates of functioning that are dysfunctional. Schwartz, however, discusses the 'line-drawing problem' without explicitly considering the situation-specificity of functions. An adequate account needs to draw the line between functional and dysfunctional rates of functioning relative to a situation.

It is important to understand what a solution to the line-drawing problem does and does not require. It does not require that a line be drawn at one precise rate of functioning for every situation. To require that would be like requiring the exact number of hairs below which someone is bald. The boundary between appropriate and inappropriate rates of functioning is vague. Accordingly, there will be a gray area between clear cases of appropriate rates and clear cases of inappropriate rates. What a solution to the line-drawing problem does require, however, is a principled way to classify clear cases into appropriate versus inappropriate rates of functioning in a manner that can guide us in biological or medical contexts. In the spirit of BST, our solution explicates appropriate versus inappropriate rates of functioning in terms of contributions to fitness.

Let *RC* be a reference class, *X* a trait that is distributed over that reference class, *Y* a function of that trait, and *S* the set of situations appropriate for the performance of *Y* by *X*:

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

(i) If an organism in *RC* possessing *X* is in s and $s \notin S$, then *X* performs function *Y* at a rate of zero (or close to zero).[14]

(ii) If an organism in *RC* possessing *X* is in s and $s \in S$, 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.

To determine whether a given rate of functioning is adequate, we compare it with the rates of functioning that are physiologically possible for that trait in

[14] Why do we say that, in non-appropriate situations, the appropriate rate of functioning is (near-)zero? After all, it is at least conceivable that a trait performs a certain activity at a moderate to high rate even when it is not called upon to perform the function in question. If there are such cases, our clause (i) should be weakened to allow them. But such cases appear unlikely given the metabolic cost of carrying out an activity that does not serve a function.

the relevant situation. Then, we determine the contribution that each rate of functioning makes to some measure of fitness. On that basis, we partition the possible rates of functioning into two sets: those which provide comparatively better contributions to that measure of fitness and those that provide comparatively worse contributions. Whether our initial rate of functioning provides an adequate contribution depends on which of these sets it belongs to.

In male humans, for example, a resting heart rate in the range of 65–74 beats per minute is adequate relative to a resting heart rate >85, which is inadequate (Cook *et al*. [2006]). This is because the latter, compared with the former, makes a substantially worse contribution to fitness on various measures such as likelihood of heart failure or sudden cardiac death. Similarly, a rate of gastric emptying of 50% of stomach contents by one hundred minutes following consumption of solid food is adequate, relative to a much lower rate, such as 50% after four hours (see Section 3), because the latter rate is associated with much higher mortality rates. Clearly, which measure of survival or inclusive fitness we select will be influenced by various contextual factors. If we are talking about a disease that only affects elderly people, then we will not use number of offspring as the relevant measure, but we might use life expectancy.

Our approach distinguishes between appropriate and inappropriate rates of functioning in terms of the relative severity of negative consequences associated with the performance of the function at that rate, where negative consequences are construed in terms of relative contributions to some measure of fitness. But Schwartz argues that we cannot define 'negative consequences' simply on the basis of some measure of fitness, because someone may have a disease that does not affect their fitness. Thus he holds that 'any measure of "negative consequences" needs to be more fine-grained and contextual than just looking at effects on survival and reproduction' ([2007], p. 379). We can avoid this problem, however, not by expanding the definition of 'negative consequences', but simply by clarifying that we are measuring the expected contribution to fitness relative to an appropriate situation. For example, achromatopsia (a condition characterized by poor to non-existent colour vision) does not necessarily affect one's actual fitness. However, if that individual were in a situation requiring the processing of colour information, that person's fitness may be negatively impacted. This is all that matters for assessing the appropriateness of the rate of colour processing.

As noted above, this account improves on Hausman's definition of 'level of functioning' as our definition—unlike Hausman's—allows for near-pandemic diseases (or what Schwartz calls 'common diseases') to be dysfunctional. For example, one explanation of the rapid decline of the Cascades frog, *Rana cascadae*, in Oregon appeals to the recent increase in ultraviolet-B radiation, which appears to affect immune response and hence render the larvae more susceptible to a regional fungus (Sarkar [1996]). Regardless of whether this immune response becomes pandemic, the frog's lowered immune response does not provide an adequate contribution to survival or inclusive fitness when it leads to premature death of larvae. This is so because the lowered immune response leads to a substantial fitness decrease relative to what we know to be physiologically possible for the frog.

The notion of adequate contribution to survival or inclusive fitness is appropriately somewhat vague because the notion it explicates—appropriate rate of functioning—is somewhat vague to begin with. The important point is that the vague boundaries of the two notions—the explicandum and the explicans—correspond to one another. And they do: if a trait that is contributing to survival or inclusive fitness is functioning at an inappropriate rate, its contribution is inadequate; if its contribution is inadequate, its rate of functioning is inappropriate.

Moreover, our account explicates one term—appropriate rate of function—that is not clearly applicable to empirical cases without relying on intuition, with another term—adequate contribution to survival or inclusive fitness—that has clear empirical criteria of application. What we have to do to apply the latter term is enumerate the possible rates of functioning and measure the contribution to survival or inclusive fitness associated with each.

Within the gray area that lies between clear cases of appropriate functioning and clear cases of inappropriate functioning, the range of factors and situations that are considered in judging whether a rate of functioning is appropriate are somewhat dependent on the interests of the observers. For example, is a resting heart rate that falls between 75 and 84 beats per minute adequate or not? Different experts may answer this question differently. This is acceptable because the boundary between appropriate and inappropriate rates of functioning is vague (cf. Boorse [2002], p. 71; Schwartz [2007], p. 383). For example, different people may choose to draw the line between the lower end of normal thyroid production and hypothyroidism (a dysfunction) in slightly different places, although they will all agree on when the thyroid is clearly performing well and when it is clearly performing poorly. This is, in part, a conventional decision to be resolved by the epistemic and pragmatic context. The values of observers (for example, the level of risk of death a subject considers acceptable) may even play a role in deciding where to place the boundary. From an ontological point of view, what matters is that the distinction between adequate and inadequate rates of functioning is grounded on the objective contribution by a trait to survival or inclusive fitness.[15]

[15] Someone might object that the possible influence of value judgments in setting the exact boundaries of the term makes 'function' (or 'appropriate rate of function') into an evaluative term. But that value judgments may be invoked to settle borderline cases doesn't make the term 'function' evaluative—or if it does, it does so in a relatively trivial way. The boundaries of most terms, even paradigmatically descriptive terms such as 'tall' or 'red', can vary based on speakers' values. For

The absence of an exact boundary between appropriate and inappropriate rates of functioning does not matter for another reason. In principle, the method we just outlined may be used to replace the qualitative distinction between appropriate and inappropriate functioning with an objective, precise, quantitative scale. Any rate of functioning may be given a score relative to a situation based on the contribution it makes to the survival or inclusive fitness of an organism. If a quantitative scale replaces the qualitative distinction between appropriate and inappropriate rates, the boundary issue disappears entirely and therefore becomes irrelevant because there is no longer any need to 'draw a line'.[16] For present purposes, however, we may continue to speak of appropriate versus inappropriate rates of functioning.

2.4 Dysfunction

We are finally in a position to explicate '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, with $s \in S$, at the rate that is appropriate in s.

Roughly, to say that X is dysfunctional with respect to Y means that either X is in an appropriate situation for the performance of Y but does not perform Y at the appropriate rate, or there is at least one situation that calls for the performance of Y such that if X were in that situation, X would not perform Y at its appropriate rate.[17]

One consequence of this account is that a trait may be dysfunctional even though it does not reduce the fitness of the individual in his or her current environment. For example, having an amputated leg does not necessarily decrease one's contribution to survival or inclusive fitness. If there are appropriate situations, however, within which having an amputated leg would reduce one's contribution to survival or inclusive fitness—and in which

example, if people must be tall in order to go on certain rides at an amusement park, this might provoke a disagreement about whether some people are tall, and that disagreement may hinge on values. But that does not make 'tall' an evaluative term. The same holds for 'function' or 'appropriate rate of function'.

[16] Hausman ([2012]) argues for a similar conclusion by suggesting that the qualitative distinction between 'health' and 'pathology' can be replaced by a quantitative measure of functional efficiency.

[17] *Ceteris paribus*, that is barring finks (Martin [1994]), antidotes (Bird [1998]), and similar defeating conditions (Fara [2009]).

amputees would not have alternate means of achieving the same goal—it is dysfunctional.[18]

Another consequence of our definition is that it introduces a kind of parity into judgements of function and dysfunction. Having flat feet, for example, may be dysfunctional with respect to rapid locomotion on treacherous terrain. The same trait may be functional with respect to avoiding the draft. Any particular trait-variant in a given individual may be functional with respect to some functions and dysfunctional with respect to others.[19] That our theory allows a trait to be functional with respect to function F_1 and dysfunctional with respect to function F_2 is a strength, as it is consistent with the possibility that complex biological traits embody trade-offs between competing selection pressures.

3 Performing Functions at Appropriate Rates in Appropriate Situations

Kingma accurately points out that in order to determine whether a trait is dysfunctional, BST functions must be indexed to situations. Kingma's observation that functions must be indexed to situations leads to her main critique of Boorse, namely, that BST cannot distinguish between when a trait is functional and when it is dysfunctional. To illustrate that point, she imagines the digestive tract in four different situations. The first three situations (*S*1, *S*2, and *S*3) are fairly common; the fourth (*S*4), in which the individual has ingested a poison that has paralyzed the digestive tract, is uncommon ([2010], p. 251).

In *S1*, the individual is resting after a meal. Here, the function of the digestive tract is digesting food, because that is its statistically-typical contribution to survival or inclusive fitness in that situation. In *S2*, the individual has been fasting. Here, the stomach is 'doing very little' ([2010], p. 251). Yet, in *S2*, according to Kingma, 'doing very little' is not dysfunctional, because that is its statistically-typical contribution to survival or inclusive fitness in that situation. Presumably, the idea here is that it would be bad for the person if the stomach attempted to perform its digestive function in this situation (for example, by releasing digestive acids in the absence of food). In *S3*, the individual has been exercising heavily. Here, the digestive tract 'does not perform its

[18] Note that this definition would also allow us to evaluate how well (or poorly) an organ is functioning even when it has been artificially removed from the body. For example, if a neuron is not in its normal environment (say, the neuron is in a dish), then whether it is dysfunctional depends on whether, if the neuron were in its normal environment (embedded in the brain of an organism), it would release neurotransmitters when adequately depolarized. This can be tested by depolarizing the isolated neuron and determining whether or not it releases neurotransmitters. We thank an anonymous reviewer for helping us clarify this point.

[19] We thank Dan Hausman for bringing this issue to our attention.

normal function [of digestion]' ([2010], p. 251). However, in *S3*, according to Kingma, not performing its normal function is the appropriate response of the digestive system ([2010], p. 251). This is presumably true for the same reason that the stomach's 'doing very little' is appropriate in *S2*.

In *S4*, the individual has ingested a poison that paralyzes the digestive tract. (This, among other things, would cause a condition known as gastroparesis, which is characterized by an inability of the stomach to push out food.) Here, the digestive tract 'does not perform its function [...] or only performs this function to a very small degree' ([2010], p. 251). In this situation, however, the absence or decrement of function is 'the normal contribution to survival and reproduction [survival or inclusive fitness, in our formulation] of digestive tracts on the occasion of poison taking' ([2010], p. 251). Consequently, and contrary to our intuition, we must judge that the poisoned digestive tract is not dysfunctional.

Kingma uses another example to illustrate this problem, namely, the situation of paracetamol overdose which severely reduces liver function. As she puts it, '[a]nybody who overdoses on paracetamol will induce a severe and irreversible case of liver failure. Liver failure, which is a reduced level of liver function, is the situation-specific quantitative normal species function in the situation of overdosing on paracetamol' ([2010], p. 251). In her view, BST cannot deem paracetamol-induced liver failure dysfunctional, because liver failure is the statistically normal response to paracetamol overdose.[20]

Once BST is modified as suggested above, however, we have a principled basis for evaluating whether a trait is functional or dysfunctional. We will use Kingma's example of digestion to illustrate our theory. In order to determine whether the digestive system is functional on these four occasions, *S1–S4*, we must first determine what the relevant function of the digestive tract is. Here, we can rely on the observation that the function of a trait is defined as a conditional probability. Under those circumstances in which the digestive system contributes to survival or inclusive fitness, the probability that it does so by digesting food is non-negligible. Consequently, one function of the digestive system is to digest food. (This does not preclude the possibility that the digestive system has other functions as well, such as protecting the inner organs from harmful ingested substances).

We must also determine the set of situations in which the exercise of that function is appropriate. Using the considerations presented above (Section 2), *S* can be identified as the set of situations in which food is present in the digestive system and in which the trait's physiological and ecological context

[20] Kingma ([2010], p. 259) clarifies that 'liver failure' here refers not to the complete cessation of liver function—since this would not constitute a contribution to survival or inclusive fitness at all—but to severely reduced liver function, which still makes a contribution, albeit low, to survival or inclusive fitness.

is conducive to the performance of the function (for example, the food is non-toxic, the metabolic costs of digestion are not too high, and so on). This is because the probability that the digestive system is in this situation-type, given that it successfully contributes to survival or inclusive fitness, is close to one, and there is no proper subset of that situation-type for which that is true.

Third, we must determine the appropriate rate of digestive function in situations in which it is appropriate for the digestive system to digest. One measure would be the rate of gastric emptying, which is often measured as gastric half-emptying time, the amount of time taken for the stomach to push out half of its contents following consumption of a standardized meal. We can look at the rates of gastric emptying that are physiologically possible for the stomach and partition this set into those rates that adequately contribute to survival or inclusive fitness and those that do not, on an appropriate measure of fitness. For example, the average gastric half-emptying time for healthy subjects is about one hundred minutes following consumption of solid foods (Lee *et al.* [2000]). Rates of gastric emptying that are substantially lower than this are associated with increased mortality (Jung *et al.* [2009]). On the basis of this information, we can identify the conditions under which the trait is dysfunctional: those that render the stomach unable to digest at the appropriate rate in *S*.

Now we can briefly evaluate the functional status of the digestive system in each of the four situations. In *S1* (the individual is resting after a meal), the situation is appropriate for digestion and digestion should occur at a relatively high rate. By hypothesis, digestion is occurring, so the stomach is not dysfunctional. In *S2* (the individual has been fasting), the digestive system is not digesting food. This is as it should be because, as noted in Section 2, if the organism is in s ($s \notin S$), the rate of functioning of zero (or close to zero) is appropriate. Furthermore, by hypothesis, the stomach is still able to digest should the situation become appropriate (for example, by the ingestion of food). Thus, the stomach is not dysfunctional. Situation *S3* (the organism has been exercising heavily and hence is digesting at a low rate) is similar to *S2*, except that the situation is such that the appropriate rate of functioning is now low rather than nil.

In contrast, in *S4* (the individual has ingested a poison which paralyzes the digestive tract), the stomach is dysfunctional because it cannot digest food. That is, there is at least one s ($s \in S$) such that even if the individual were placed in s, the digestive system would not digest food. Because of this, the stomach is dysfunctional.[21]

[21] What if the stomach were temporarily paralyzed due to heavy exercise (cf. Hausman [2011], pp. 665–7)? In that case, the stomach may be (temporarily) dysfunctional with respect to digestion.

The same account applies to the paracetamol case. Paracetamol-induced liver failure is dysfunctional because paracetamol renders the liver unable to perform its function at the appropriate rate in situations appropriate for the liver to perform its function. It makes no difference that liver failure is the statistically normal liver response to paracetamol overdose, because on our account the appropriate rate of liver functioning is determined by looking at the entire set, *S*, of situations appropriate for the liver's exercise of its function, and the rates of functioning appropriate to each situation, $s \in S$. In the case of paracetamol overdose, there is at least one $s \in S$ such that the liver would not perform its function at the appropriate rate; consequently, it is dysfunctional. It is dysfunctional because it cannot perform its function in the appropriate situation for the performance of the function, namely, one in which it is called upon to contribute to metabolism or the processing of toxins in the blood.

Kingma knows that, in order to determine whether the digestive system is dysfunctional in *S4*, we must assess not only whether it is currently digesting food, but also whether it is disposed to perform its function of digesting food under appropriate circumstances. She does not deny that functions have a dispositional character. In fact, she builds this dispositional character into her basic analysis ([2010], pp. 246–7). But she asserts that in the relevant counterfactual situation, the digestive system would digest food:

> [. . .] although the digestive tract does not perform its function (digesting food) on the occasion of having taken the poison or only performs this function to a very small degree, that is the normal contribution to survival and reproduction [. . .] on the occasion of poison taking. *Moreover, the digestive tract is still disposed to behave normally on all other occasions (such as resting after eating a meal which did not include the poison, or resting after fasting)*. ([2010], p. 252; emphasis ours)

Here, Kingma misdescribes the relevant counterfactual situation. The way to construct the relevant counterfactual situation is to freeze the trait in its current state and imaginatively transport it into the counterfactual context (cf. Hausman [2011], p. 663). In the counterfactual context (one in which the organism is resting after a meal), the digestive tract should be described as paralyzed and the stomach as having gastroparesis. Hence, the digestive system would not be able to perform its function in that situation. This method for constructing counterfactuals can be justified simply by reflecting on the usual purpose of such counterfactuals in biomedical contexts, as will be explained below.

One important question that arises when a trait is not performing its function is the following: is the trait, as it were, intrinsically dysfunctional or is it merely within an environment that is not conducive to the performance of the function? Suppose that a person is submerged in water. Clearly, the lungs cannot perform their function of distributing oxygen. However, at least initially (a short time after submergence), the lungs are not dysfunctional. This is because if they were transposed to the normal environment for their functioning (an environment replete with air), they would perform their function. Alternatively, imagine that someone's lungs are unable to perform their function of distributing oxygen because that person's chest was crushed in a stampede and the lungs suffered gross structural damage. If that person's lungs (given the structural damage), were placed within the normal environment for their functioning (one without a stampede and without chest damage), they would not be able to distribute oxygen. It is for this reason that, in the case of lung damage following being trampled, the lungs are dysfunctional. They are not merely unable to perform their function due to an environment that is not conducive to performance.

What is crucial to distinguishing between these two situations is the ability to construct counterfactual situations that probe how the trait, in its current intrinsic condition, would behave in an altered environment. The main purpose of constructing such counterfactuals would be undermined if one were allowed to imagine what the trait would do were it restored to its original intrinsic condition (i.e. prior to the damage). But this is what Kingma does when she claims that the paralyzed digestive system is not dysfunctional. After all, she notes, it would be capable of performing its function if it were in a situation in which, among other things, it were not paralyzed. But if we accepted this method of constructing counterfactuals, then nothing could ever be dysfunctional!

This response—that Kingma misdescribes the relevant counterfactual situation—could have been formulated without relying on our revised version of BST, which is what Hausman ([2011]) does. But one virtue of our revised formulation is that it provides a more principled response to Kingma's dilemma by allowing us to determine what the relevant situations are for evaluating whether a trait is functional or dysfunctional. In other words, Hausman shows that even if one accepts Kingma's method for relativizing functions to situations, and thus accepts that a liver impaired by paracetamol is functioning normally relative to that situation, one can still show that the impaired liver is dysfunctional. This response concedes too much. Having improved on Kingma's method for relativizing functions to situations, we conclude that a liver impaired by paracetamol overdose is not functioning normally in any sense of the term.

4 Conclusion

We have sketched a novel and improved version of Boorse's biostatistical theory of function. Roughly, our theory maintains that: (i) functions are non-negligible contributions to survival or inclusive fitness (when a trait contributes to survival or inclusive fitness); (ii) situations appropriate for the performance of a function are typical situations in which a trait contributes to survival or inclusive fitness; (iii) appropriate rates of functioning are rates that make adequate contributions to survival or inclusive fitness (in situations appropriate for the performance of that function); and (iv) dysfunction is the inability to perform a function at an appropriate rate in appropriate situations. Contra Kingma ([2010]), we have shown that our improved theory is able to distinguish situations in which a trait does not perform its normal function because it is dysfunctional from situations in which it does not perform its normal function without being dysfunctional. It also resolves two other problems that have afflicted Boorse's theory of function, namely, the problem of multi-functional traits and the line-drawing problem.

Acknowledgements

Thanks to our audiences at the 2011 Central States Philosophical Association Meeting and 2012 APA Central Division Meeting, and to our respective commentators, Eric Kraemer and Dan Hausman. Special thanks to Dan Hausman for several incisive comments on previous drafts. Finally, thanks to Elselijn Kingma and the anonymous referees for this journal.

Justin Garson
Department of Philosophy
Hunter College/City University of New York
695 Park Ave
New York, NY, 10065, USA
jgarson@hunter.cuny.edu

Gualtiero Piccinini
Department of Philosophy, Center for Neurodynamics
University of Missouri – St. Louis
599 Lucas Hall (MC 73)
1 University Blvd.
St. Louis, MO 63121-4400, USA
piccininig@umsl.edu

References

Bird, A. [1998]: 'Dispositions and Antidotes', *The Philosophical Quarterly*, **48**, pp. 227–34.

Boorse, C. [1975]: 'On the Distinction Between Disease and Illness', *Philosophy and Public Affairs*, **5**, pp. 49–68.

Boorse, C. [1976]: 'Wright on Functions', *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. 3–134.

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

Buller, D. J. (*ed.*) [1999]: 'Etiological Theories of Function: A Geographical Survey', in *Function, Selection, and Design*, Albany: SUNY Press, pp. 281–306.

Cook, S., Togni, M., Schaub, M. C., Wenaweser, P. and Hess, O. M. [2006]: 'High Heart Rate: A Cardiovascular Risk Factor?', *European Heart Journal*, **27**, pp. 2387–93.

Fara, M. [2009]: 'Dispositions', in E. Zalta (*ed.*), *Stanford Encyclopedia of Philosophy*; <plato.stanford.edu/archives/sum2009/entries/dispositions/>.

Garson, J. [2008]: 'Function and Teleology', in A. Plutynski and S. Sarkar (*eds*), *Companion to Philosophy of Biology*, Malden, MA: Blackwell, pp. 525–49.

Garson, J. [2011]: 'Selected Effects and Causal Role Functions in the Brain: The Case for an Etiological Approach to Neuroscience', *Biology and Philosophy*, **26**, pp. 547–65.

Garson, J. [2012]: 'Function, Selection, and Construction in the Brain', *Synthese*, **189**, pp. 451–81.

Hausman, D. [2011]: 'Is an Overdose of Paracetamol Bad for One's Health?', *British Journal for the Philosophy of Science*, **62**, pp. 657–68.

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

Jung, H., Choung, R. S., Locke, G. R., Schleck, C. D., Zinsmeister, A. R., Szarka, L. A., Mullan, B. and Talley, N. J. [2009]: 'The Incidence, Prevalence, and Outcomes of Patients with Gastroparesis in Olmsted County, Minnesota from 1996 to 2006', *Gastroenterology*, **136**, pp. 1225–33.

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.

Lee, J. S., Camilleri, M., Zinsmeister, A., Burton, D., Kost, L. and Klein, P. [2000]: 'A Valid, Accurate, Office-Based Non-radioactive Test for Gastric Emptying of Solids', *Gut*, **46**, pp. 768–73.

Martin, C. B. [1994]: 'Dispositions and Conditionals', *Philosophical Quarterly*, **44**, pp. 1–8.

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

Nagel, E. [1977]: 'Teleology Revisited: Goal-Directed Processes in Biology and Functional Explanation in Biology', *Journal of Philosophy*, **74**, pp. 261–301.

Neander, K. [1991]: 'Functions as Selected Effects: The Conceptual Analyst's Defense', *Philosophy of Science*, **58**, pp. 161–84.

Sarkar, S. [1996]: 'Ecological Theory and Anuran Declines', *Bioscience*, **46**, pp. 199–207.

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

Sommerhoff, G. [1950]: *Analytical Biology*, London: Oxford University Press.