The Process Dynamics of Normative Function

Wayne D. Christensen Mark H. Bickhard

"The hand separated from the body is not a true hand."
Aristotle, *Politics* I.2.1253a20-21.

Introduction

This paper argues for a radically different approach to normative function than the dominant etiological account, an approach with a broader explanatory agenda. Concepts of biological function address two related but distinct issues: understanding biological organisms as complex organized systems, and understanding the way in which parts have come to be present in organisms because of their adaptive value. The etiological account focuses primarily on the latter, treating the complex organization of organisms as a derivative issue solved in passing, as it were, by explanations for the presence of parts within organisms. This way of treating the issues is unsatisfactory, however, since the etiological model of normative function fails to provide the conceptual resources required for understanding the complex process interrelationships characteristic of organisms as organizationally integrated systems. It makes normative functions epiphenomenal, it provides a limited account of dysfunction, and it is unable to adequately characterize multifunctional parts or the organization of densely connected functional complexes. In other words, although explanations for the presence of traits and explanations for the way organisms function as complex integrated systems certainly overlap, they are not conceptually or causally identical issues, and nor can one issue be treated purely as derivative upon the other. Our approach recognizes this and develops an account of normative function that directly focuses on whole-system organization and process interdependence.

Part I: Problems with the Etiological Model of Normative Function

In this section we will outline the conceptual basis of the etiological approach to understanding normative functions. Some of the ground we will cover has been well trodden, but it is worthwhile re-examining since many of the problems of the etiological approach stem from fundamental issues that are often taken for granted.

Normative functions and the design analogy

The basic strategy of the etiological approach is to understand natural selection as the virtual designer that shapes organisms so as to be adaptive and thereby imparts functions to their parts. There are two major aspects of the explanatory structure of the etiological approach: the design analogy account of the *nature* of biological organization, and the natural selection account of the

construction of biological organization. The design analogy conceptualizes the organization of biological systems by analogy with designed artifacts. Artifacts are recognized as being organized in the sense that their parts fit together in an interrelated, globally coherent way. This global coherency is supposed to stem from an organizational scheme or plan in the mind of the designer which specifies the global organization of the system and which has guided its construction. The essence of the design analogy is the claim that saying that a biological system is organized is just to say that the system appears 'as if' it had been designed; it appears as though there is such an intentional scheme which specifies the global organization of the system.[1]

The etiological account explains the appearance of design by appealing to natural selection. Selection favors variations in organisms that are beneficial to the success of the organisms, with the result that, over an extended history of selection, the many traits of a species tend to be co-adapted with each other and with the environment to form organisms that are adaptively organized systems. The functional integration and aptness for its niche of such an adaptive system makes it appear as if it had been purposefully designed. Whilst natural selection involves no literal designer or mental plan, the selection force which differentially favors a particular trait can be considered equivalent to the intentions of a designer in shaping a particular feature of an artifact during the manufacturing process.

Despite its long history and widespread support this alliance between design and selection as way of understanding biological organization has a number of serious problems. Firstly, the design analogy makes all organization quasi-mentalistic. Explaining the properties of organisms in terms of the properties of minds is a strange approach for a supposedly naturalist account. Secondly, and perhaps more fundamentally, the design analogy provides no conception of the *nature* of organization at all. In effect, the design analogy addresses biological organization by simply presupposing a resemblance with a paradigm class of organized systems, namely artifacts. Moreover the explanation of the nature of organization comes to no more than the claim that there is (or it is as if there were) a scheme that specifies global relations for the system. Thus, at most the design analogy explains the *transfer* of organization from the plan in the mind of the designer to the material artifact. It doesn't explain its nature or origins.

Adding natural selection to the picture does not modify this basic problem. Selection itself is, and must be, organized: for selection-driven adaptation to occur there must already be a coherently organized organism and a possible successful 'lifestyle' (i.e., a systematic pattern of interaction). So whilst explanations that appeal to it can explain the transfer of organization from a structured selection force to changes in organism morphology, there remain fundamental organizational issues that are not explained. What distinguishes an organized or coherent selection force from a disorganized collection of influences? What is it that allows selection to create organized organisms?

---

[1] Compare Kitcher 1998, pp.271-2.

One might at this point object that the fundamental organizational presuppositions of selection are minimal (e.g. Darwin's Principles' (Lewontin 1970)) and that the actual organizational complexity within which selection works is the product of prior selection. However the claim that selection is a powerful explanation of the genesis of organization does nothing to provide a model of that organization, and, therefore, nothing to model the interaction between current selection and current organization. To assume otherwise is to assume that organization is redundant relative to selection, which is baldly false. It is the evolutionary equivalent of the behaviorist assumption that internal organization is redundant relative to reinforcement experience. Furthermore, the underlying intuition of this argument appeals to a strong adaptationism that is implausible. It is far from clear how much of the organization that *participates* in the selection process is properly attributable to selection itself. If the *constraints* on variance have coherent patterns, then they may play an important role in shaping the outcomes of selection.[2]

A quite different kind of objection concedes that a strong adaptationist interpretation of evolution is flawed, but instead claims that the organization presupposed by selection is *causal role* organization, which is not the proper explanatory focus of the etiological model. Rather, the etiological model explains *normative* organization. In framing the issue of norms in biology the special explanatory burden is usually taken to be accounting for dysfunction. A part within a biological system (such as a heart) can be dysfunctional in the sense that it can fail to behave in an appropriate way (such as by not pumping blood). The etiological model explains this sense of failure to match a norm as divergence from the form of organization in the organism that was subject to selection. Consequently, it may be argued, although adaptationist explanation does not explain all organization it does distinctively explain normative organization, and since this is presumably what is relevant to understanding adaptiveness the organizational explanation is far from empty.[3]

Whether a conceptual separation between norms and so-called 'causal role' organization is possible depends, however, on what norms are and what their role is in biological explanation. The orthodox view is that norms are quite distinct from the general problem of form, and is based on a distinction between organization as *mere pattern* and norms as concerned with *preferred pattern.* The intuitive basis of the distinction is that, although for any system it will be true that there is a global pattern of relations amongst its parts (it has *some* form of organization) nothing of significance hinges on this pattern. A new pattern will result if the global relations are modified, but there is no reason to prefer one pattern to another; it is simply that the system

---

[2] Cf. Ahouse 1998, Amundson 1989, Griffiths 1996, Sober 1998. As Ahouse puts it nicely, 'Is anyone surprised when they leave a pizza restaurant with a pizza?' (1998, p.54).

[3] It should be noted that this qualification could at least potentially constitute a major retreat from the original explanatory ambitions of the adaptationist program, which was to explain biological form *in general* as intricate purposeful structure produced by design. If the supposed 'non-normative' organization is sufficiently pervasive it undermines the design analogy as a general model of biological organization. See below, and compare Schaffner 1993 ch.8.

changes. On the other hand the problem of norms arises because, in the case of both living organisms and artifacts, there *is* reason to prefer one pattern to another, since only one pattern specifies what the system, including all of its parts, is *supposed* to do. Introducing design intentionality provides a benchmark for preferring one pattern to another by introducing purposes. Thus, norms are a separate type of thing to pattern per se; they specify *of* a pattern that it is the correct one for the system.[4]

This strong type distinction between causal organization and normative organization begs the question relative to our approach because we adopt an alternative approach to norms that undermines precisely this distinction: we treat norms as a(n emergent of a) type of organizational property. That is, norms are intrinsically organizational because they are a feature of the relationship between the parts of the system and the system's global organization. Systems with norms have a type of organization that specifies a set of interrelations, or roles, for the system's parts. Norms are identified with limiting relations on the behaviors of parts: within the limits the global organization is preserved, outside the limits it is disrupted. So norms are really an *aspect* of the organization of the system rather than something separate from it. We will outline our theory of the nature of these functional interrelations in part II, where we will argue that they provide a much more useful account of norms for understanding adaptiveness than approaches which make norms external to organization.

Epiphenomenal functions and inconsistency

If our criticisms of the design analogy approach to understanding functions are correct then the etiological account doesn't fundamentally explain functions at all, not even normative functions. However, whether one is persuaded of this depends on how pressing the weakness of the design analogy appears and how likely the prospects seem for a viable intrinsic account. Before we present such an account we discuss several further problems with the etiological model, which are that it makes normative functions epiphenomenal, it illegitimately presupposes a more fundamental concept of normative function, and it lacks the ability to adequately account for important biological phenomena relevant to function, including the nature of failure for the organism, and multifunctional features. All of these problems point to the fact that a theory of function that directly addresses the problem of organization is required.

Etiological theories ascribe proper functions in virtue of the fact that the system has the right sort of history. Unfortunately, this makes functions epiphenomenal: a system can have no functions even though it is causally identical to one that does have functions, simply because it lacks the required history. The basic point, if not its implications, is clear in Millikan (1989a). She claims that an accidental double that is molecule for molecule identical with an object that has proper

---

[4] Most of the problems of the etiological approach that we discuss below, and in particular the problem of epiphenomenality, stem from this fundamental assumption. As a result, no amount of tweaking of the formula can solve the problem. A solution is only possible by abandoning the assumption that norms are in principle separate from organization.

functions *does not* have functions because it doesn't have the right history (p. 292). Whatever difference it makes to a system to have etiological normative functions, it is not a causal difference.

Is this really a problem, however? A great deal of the intuitive appeal of the etiological approach stems from the fact that, in the case of biological organisms, historical explanation provides a lot of information about current causal properties. The etiological theory of normative function has been taken to be a successful naturalization precisely because the process it appeals to (evolutionary selection) is causal. In this respect the theory is causal, it simply appeals to distal rather than proximal causes. And since many scientific models appeal to distal causes, why not a model of function? From our perspective the problem is not the appeal to distal causes, but the assumption of the historical constitution of normative function. If normative functions are *constituted* by having a certain type of history then they are not causal at all since they have no consequences, even if the circumstances for *applying* the concept of normative function appeals to (distal) causes.

But whether you agree that this is a crucial problem depends on whether you agree that normative functions ought to be construed as causal properties. Perhaps the best way to show that at least some form of causal account of normative function is required is by highlighting a lurking inconsistency in the etiological approach when it is construed as the fundamental account of normative function. The inconsistency lies in the fact that the etiological account actually presupposes a more fundamental concept of normative function, in the sense that a mutation must perform a function that is useful to the organism for it to be favored by selection. But 'useful' is a clearly normative relation: if something is useful then it is good in some way. Consequently, selection *presupposes* normative relations, and, insofar as the etiological account appeals to selection as the original *source* of norms, it is inconsistent.

Problems with modeling failure and functional complexes

In conjunction, the twin problems of epiphenomenal functions and inconsistency suggest that there are serious conceptual weaknesses in the etiological account. However the etiological account has been popular for a reason: it seems to cohere well with neo-Darwinian evolution theory. It is consequently able to appeal to a large body of scientific practice and empirical evidence that identifies or associates functions with the influence of natural selection. If we are right then the effectiveness of this research paradigm lies with the association rather than the identification. Because biological organisms are so heavily structured by history, historical explanation does play a major role in understanding how they work. Our criticisms have so far hinged on the somewhat subtle point that it is a mistake to *identify* the way organisms work with their history. However if the issue was merely one of metaphysics or vocabulary then the significance of this might not be regarded as great enough to warrant any methodological or heuristic change in approach. Our final criticisms focus on empirical weaknesses of the etiological account, in particular for modeling failure and deeply embedded multifunctional complexes.

Failure. It may seem surprising that there is a problem for the etiological approach in explaining failure, given that it has been framed around the problem of accounting for dysfunction. The difficulty centers on the fact that the concept of functional failure employed makes no direct reference to the interaction between the failing part and the rest of the organism. As a result, the account says nothing about the relative *significance* of functional failure or, indeed, anything *at all* about the effects that the failure has for the organism itself (as distinct from its ancestors). However, functional failures can vary quite extensively in their effects, ranging in severity from total heart failure to mild shortsightedness. Thus, whereas heart failure is invariably lethal (other than with massive medical intervention), shortsightedness may pose comparatively minor problems. In the case of a social primate, for instance, shortsightedness might result in a moderate impairment in capacities such as fruit foraging and reading social intentions. This *could* lead to reductions in health and status, or it may not, depending on circumstances such as food availability and whether the animal has high status social allies. There are some circumstances in which shortsightedness could prove lethal, such as through failure to detect a predator, but it is also possible that the animal may find ways of compensating effectively, such as by relying more strongly on hearing and attending more closely to other members of the social group for predator detection.

Understanding the relative significance to the organism of particular types of dysfunction is clearly an important adaptive issue, and arguably a key issue for an adequate account of normative function. There is a significant normative dimension to the difference between a lethal dysfunction and a dysfunction that imposes few deleterious effects, and this needs to be accounted for. The fact that the etiological account says nothing about these issues is a serious explanatory deficit.

In reply it might be argued that the relative functional significance of a particular character trait is captured by the specific nature of the selection pressure that gave rise to it (e.g. strong selection for heart function in all vertebrates, comparatively weaker selection for vision accuracy in primates). But this is not in fact part of the etiological account. The only criterion specified for determining normative function is whether there has been selection for the trait per se; the nature of the selection pressure is irrelevant. This means that as far as the etiological account is concerned mild shortsightedness is just as dysfunctional as total heart failure. There is no way to differentiate the relative significance of a particular dysfunction. Moreover, even if the account were expanded to include reference to the nature of the original selection pressure, this would still not address the significance of failure for the current organism.

The underlying issue is that for some issues relevant to normative function it is important to understand the range of interaction possibilities for an organism, not just the circumstances of prior selection. Thus, whether a variation in the structure of an organism or its ecological circumstances results in dysfunction *for the organism*, and what the specific significance of the dysfunction may be, is something that depends on whether the variation compromises the interactive viability of the organism in its current circumstances, not on whether it is a departure from the circumstances under which previous selection occurred.

The Hutchinsonian concept of niche is a good example of a biological conception of viability that appeals to interaction potential in this way (Hutchinson 1957, Griesemer 1992). Hutchinson characterized a niche as a state space whose dimensions represent the environmental factors acting on the organisms of a species. With respect to this state space, the *fundamental* niche is the range of environmental factors that will permit the species to persist, whilst the *realized* niche is the region of the fundamental niche which corresponds to the actual environment in which the species is situated (presuming that it is in fact viable). The fundamental niche concept is thus a way of characterizing the interactive potential of the species, in the sense that it defines a space of potential environments in which the species can be viable. Whether the species is actually viable depends on whether this space of possible environments overlaps with the actual environments that are geographically accessible to it.

There are problems with Hutchinson's concept of niche: it is very abstract and static, so it doesn't provide much insight into the processes by which organisms interact with their ecology, or with processes by which niche change and niche construction occur (cf. Sterelny forthcoming). The important point for our purposes is the general idea of characterizing viability in terms of interaction characteristics. This provides a way of understanding norms that is directly causally relevant to adaptive systems. Thus, we can restate our earlier point about dysfunction in Hutchinsonian terms.[5] Whether an ecological or organismic change is dysfunctional for the organism or species depends on whether it results in a departure from (or restriction to) the intersection of the fundamental niche and the current environment, not on whether it is a departure from the realized niche during the period of adaptation.[6]

Deeply embedded features with multiple functions. The etiological account appears strongest for characterizing functionality in situations where there is a recently evolved feature that performs a single straightforward functional role. A typical kind of example would be the evolution of longer beaks in Galapagos finches in response to food scarcity, allowing the finches access to food in small crevices. In this type of case there is a relatively clear connection between the reasons for the presence of the feature, selection, and the functional role that the feature performs. However the approach is much less effective in characterizing the functionality of phylogenetically old features that are deeply embedded and play multiple functional roles. The problem is not that selection plays no role in shaping these complex functional relations. Rather, the difficulty that these cases pose for the etiological account is that the twin questions 'why is it there?' and 'what does it do?' start to come apart. Consequently an analysis of function that

---

[5] Hutchinson's niche concept applies to species, but as noted in the text we are more concerned with the general idea than specific formulations at this point. The concept of autonomy that we will introduce in Part II can be used as the basis for a suitable interactive account of organism viability.

[6] Organism death and species extinction are departures from this intersection between fundamental niche and actual environment. Milder dysfunctions are restrictions to the intersection. Intuitively, the less healthy an animal is the less it is able to do in its environment and consequently the range of circumstances in which it can survive becomes narrower.

focuses just on selection for the presence of a feature will not reveal the full set of functional roles that the feature performs. When there are multiple functions it becomes increasingly unlikely that the trait has been selectively optimized to perform each of them. Indeed, where a trait has been co-opted to perform a new function it may not have been structurally modified at all.

Preston (1998, p.241) comments that such cases 'appear to be rare' in biology, whilst being common in human technology. In fact, with respect to biology this claim is spectacularly wrong. Wimsatt's concept of generative entrenchment provides a general argument as to why multifunctionality can be expected to be widespread in biology (see e.g. Wimsatt and Schank 1987). Empirical examples are not hard to find. Amundson and Lauder point to the vertebrate jaw as an example (1998). In hamsters the medial extended amygdala is known to play roles in the control of sexual behavior, maternal behavior, grooming, and aggression (Newman, 1999). Multifunctionality is highly prevalent in cellular physiology: A great deal of protein evolution occurs through variations in the assembly of complex macromolecules out of pre-formed subunits. In other words, just as a chair may be used either for sitting on or as a platform for reaching a book on a high shelf, a particular protein domain that in one type of enzyme acts to catalyze a digestive reaction may in another enzyme facilitate blood clotting. A similar pattern of multifunctional utilization occurs with signaling molecules: Cyclic AMP, Ca^{2+} and calmodulin act in interrelated intracellular signaling pathways that participate in a huge variety of cellular processes, including secretion and cell proliferation (Alberts et al. 1989, p.700). Likewise, extracellular signaling molecules typically act on multiple target cells and have different effects in each type of target cell. Acetylcholine, for instance, acts in skeletal muscle cells to stimulate contraction, but serves to *decrease* the rate and force of contraction in heart muscle cells (p.684). Steroid hormones such as testosterone also typically act to trigger transcription of different genes in different target cells (p.692). In ethology, Griffiths (1997, p.66) cites Hinde (1966) as claiming that many behaviors that have acquired communicative functions have not undergone any significant modification. Thus, structures that perform multiple functions are probably no less and possibly far *more* common in biology than in culture.

As we noted above, this poses the problem that optimality assumptions become increasingly difficult to sustain in the face of dense multifunctionality, and in certain instances it may be not be the case that some of the functions that are performed actually make a selective difference to the trait. If one of the many cellular processes that are triggered by testosterone were lost, for example, it is unlikely that this alone would make a significant difference to testosterone production.

Perhaps even more significantly, though, as functional interconnections become denser the material, organizational and historical constraints of the structures and processes play an increasingly large role in shaping the organization of the system because the constraints on both variance and success are tighter. Raff (1996) discusses at some length the ways that internal organization systematically constrains evolutionary possibilities, by limiting transformation and imposing extensive patterning on possible outcomes (see especially pp.296-9). He characterizes

the evolutionary effect of what we have been calling dense multifunctionality in terms of 'integration constraints', or constraints imposed by "the resistance of existing integrated organizations to modification or reorganization" (p.300). He says that integration constraints are "the category of constraints that most of us think have exercised the primary effect on metazoan evolution" (p.300).

The etiological account has been adjusted in a variety of ways to cope with the complexities of selection influence. A 'modern history' approach deals with the problem of changing function by identifying proper function with the function that has been mostly recently subject to selection (Godfrey-Smith 1994). This is reasonable as a refinement of the basic etiological formulation but it doesn't address the problem of multifunctionality. Millikan (1999) attempts to deal with the problem of co-option for new function with retention of original function by incorporating secondary adaptation into the account. Thus, she suggests that an effect of a trait that is utilized by another part of the system can be considered a proper function of the trait when the part that performs the utilization has been selected for. In other words, although the primary trait hasn't been selected to have that particular effect, the presence of the secondary traits that utilize the effect *is* the result of selection for the effect. Consequently, the effect itself can now be considered a proper function. However, if the trait that utilizes the effect is itself multifunctional this strategy breaks down, because it may not be possible to clearly assign utilization of that particular effect as the reason for the secondary trait's presence. In a densely connected functional complex this is likely to be a significant issue.

Functional complexes also pose another kind of problem. Not only do they limit the potential for selective optimization of parts, but also the way they work is not effectively characterized by the strategy of localizing functions to parts. Deep vein thrombosis provides a topical illustration of the point.[7] It is paradigmatic that the function of the heart is to pump blood, but adequate blood circulation in fact depends on a complex ensemble of processes beyond the heart, including bodily movement. Consequently, when individuals are immobile circulation in the limbs becomes poor, and if this persists for a long period (as it does on long airline flights when people are confined to cramped seats) blood clots can form. Characterizing massively interconnected functional processes like blood circulation is deeply problematic for the etiological approach because these processes are holistically structured and they tend to exploit many features of the system, whether or not those features have been selectively optimized to contribute to the process. Is it, for example, the function of the calf muscles to pump blood? This is a perplexing question for an etiological analysis, but there is no doubt that by contributing to leg movement the calf muscles play an important role in *helping* blood circulation.

The etiological account is an inherently single factor type of explanation because it localizes functions to parts and it seeks to identify selection as the sole factor conferring functions on parts. In this respect it really is committed to a naïve form of adaptationism because it looks for selective optimization for each and every function it is willing to countenance, whether on the

---

[7] We thank Bill Herfel for this example.

output or the uptake side. As the process relations become increasingly complex the localization strategy becomes less effective as a way of understanding functional organization, and as selection effects become increasingly indirect the plausibility of individual optimization is reduced (cf. Wimsatt 1997). It is important that this diminishing explanatory power is squarely recognized, because as theoretical qualifications are added to cope with complicating factors there is a danger of sliding from a reasonable attempt to increase the accuracy of the account into a scholastic a priori Darwinism immune to refutation (cf. Ahouse 1998). In particular, at some point a Zen-like question arises: when is a massively qualified single-factor explanation really a multifactorial explanation? In other words, when should non-selection factors cease to be treated as qualifiers to adaptationist explanation and instead be considered as primary causal factors in their own right?

One response to the problem of accounting for the complicated influences on biological form is to retreat from the assumption that adaptation is ubiquitous to a recognition that there are many counterinstances, but that adaptationism is nonetheless the most useful working hypothesis (cf. Sober 1998. Thus, Godfrey Smith (in press) distinguishes empirical, methodological, and explanatory versions of adaptationism, where empirical adaptationism is the assumption that adaptation is ubiquitous, methodological adaptationism is the assumption that it is the most useful working hypothesis, and explanatory adaptationism is the assumption that selection explains special and important features of life, such as diversity and complex adaptation (see Sterelny and Griffiths 1999, ch.10). In those terms the retreat we describe is from empirical to methodological adaptationism.

However another response is to cease attempting increasingly qualified versions of adaptationism and instead develop a full-blooded multifactorial approach that attempts to systematically identify all of the significant factors that influence biological patterns and model their interaction. Miklos (1993) argues trenchantly that not only is adaptationism empirically false it can also be methodologically deeply unhelpful. He advocates model system research as a more empirically adequate methodology. Raff (1996) can also be viewed as arguing against methodological adaptationism, but he doesn't see model systems research as a panacea. Instead he argues for a sophisticated integrative approach that combines a plurality of theories, models and investigative techniques to form an integrated evolutionary account. The central thrust of this integrative approach is to treat biological systems as being shaped by a complex mixture of factors, including development, ecology, and population genetics. Each type of factor makes its own distinctive contribution to biological organization, and adequate biological explanation depends on understanding their interaction.

For our purposes one of the key issues arising from this is that the etiological account is overly committed to a particular mechanism for explaining biological organization. If a pluralistic integrative approach proves more empirically successful than methodological adaptationism, and our bet is that it will, then a functional approach that insists on single-factor explanation will have a rather limited usefulness. An integrative approach can in fact be more historically oriented than adaptationism because it allows greater scope for the recognition of historical constraints (cf.

Griffiths 1996). Perhaps surprisingly, though, it also presents a need for a non-historical approach to modeling function, because there is a need to model biological systems as *integrated* systems in a way that is at least partly independent of the mechanisms responsible for producing the organization. By modeling biological integration directly it then becomes possible to examine the relative effects that various factors can have to induce changes to the system, thus making possible a more sophisticated account of evolutionary mechanisms.

Thus, we believe that there are in fact strong empirical motivations for developing an account of normative function that is independent of selection. Such an account is not inconsistent with a suitably framed explanatory adaptationism,[8] but it does insist that if we are to learn what it is that adaptation distinctively explains then explanatory adaptation has to form part of a broader pluralist integrative approach. In this respect, Sterelny and Griffiths (1999) make an important if somewhat cryptic point: if explanatory adaptationism is to have force it cannot afford to identify adaptation with the outcome of natural selection:

> We cannot at the same time define adaptation as whatever natural selection causes and promote natural selection on the grounds that it is the explanation of a particularly puzzling phenomenon, namely, adaptation. If the theory of explanatory adaptation is to mean something substantial, then adaptation, especially complex adaptation, must be characterized independently of its putative explanation, natural selection (p.228).

Our reasons for developing an approach to function that is independent of selection have been based on concerns other than defending explanatory adaptationism per se, but we share the sentiment.

Part II: An Alternative Process Model of Normative Function

Thus far we have argued that the etiological account faces serious problems as a fundamental theory of normative function. The design analogy fails to adequately explain functional organization, the etiological account makes functions epiphenomenal and inconsistently presupposes a more basic normative concept of functional usefulness, and it faces empirical difficulties modeling failure and functional complexes that have wide ranging implications for understanding biological organization. In this section we will present an approach to function that avoids these problems by focusing directly on functional integration and the issue of usefulness. Our account tackles these issues by developing a theory of viability in terms of autonomous or self-generating systems, which provides the basic perspective or benchmark for characterizing normativity, and with an account of functional organization in terms of process interdependence. Before describing our approach, however, it is important to identify some of the reasons why it is usually thought to be impossible. This helps to clarify the criteria against which it should be evaluated.

---

[8] Amongst other things it must be an adaptationism that doesn't claim 1 function = 1 adaptation.

The disregard of usefulness

On the face of it the fact that the issue of usefulness has been resolutely eschewed by the bulk of the functionality literature is highly puzzling, since it would seem to be a relevant phenomenon. Indeed, selection works by favoring traits that are useful! However there are a number of factors that have conspired to make usefulness an unattractive topic. One of these is an historical emphasis on assigning functions to parts, which motivates drawing a strong function as/function of distinction. In this context usefulness appears to be a secondary phenomenon regarded largely in contrast to *having* a function. Thus, Wright (1973) distinguishes functions from accidental benefits, and identifies functions with the reasons for a parts presence. Much of the subsequent literature has followed this trend. The rare attempts to associate functions with usefulness, such as Bigelow and Pargetter (1987), are often accused of confusing the function as/function of distinction or not satisfying the explanatory agenda of function theory (e.g. Millikan 1989b, Mitchell 1995). This way of framing the problem has been reinforced by a similar practice in biology of distinguishing adaptations from features that merely happen to benefit the organism (cf. Godfrey-Smith 1998).

Another factor that discourages interest in the issue of usefulness has been the difficulty of providing a non-circular account of viability, health, goal states, or even a principled specification of the relevant systems, that might ground usefulness. Schaffner (1993, pp.366-367) points out that the cybernetic approach pioneered by Rosenbleuth, Weiner and Bigelow (1943), which tried to explain purposiveness in terms of negative feedback control systems, failed to provide a non-circular analysis of goal states. Goal-directed behavior is explained in terms of feedback mechanisms, which are in turn explained in terms of regulation for goal states. Similarly, explaining normative functions as contributions to health or flourishing suffers from the problem that these are themselves ostensibly normative concepts.

Perhaps the most prominent reason, however, is the sheer variability of usefulness. The diversity of life and some creative thought experiments suggest that almost anything could prove useful in some circumstances or other, which makes it seem impossible to provide a principled account of what usefulness consists of. Consequently, it is argued, the only systematic way to understand usefulness is retrospectively, in terms of what selection *has* found to be useful. Thus, Millikan (1999) argues that ‘system functions’ simply aren’t well defined. In contrast with these difficulties the etiological approach at least seems to provide clear physical principles for applying normative concepts, whatever other problems it may have.

Understanding usefulness certainly isn’t straightforward, but nor is it as hopeless as it is often regarded. Arguing that usefulness is so variable as to defy any principled account is a dangerous tactic for a proponent of selection because if usefulness really were so variable that it lacked any coherent patterns then selection wouldn’t be possible. And whilst it is true that our account of normative function isn’t directly framed around the agenda of assigning functions to parts, that is because we believe that the agenda itself is problematic, at least as a way of providing a basic account of normative function. In this respect our account of the fundamental nature of

normativity is nonstandard, because we don't treat it as the failure of a part to perform a function it is supposed to perform. Instead, we take the system as primary, model functional organization in terms of the process interactions that generate the system, and characterize usefulness in terms of contribution to system-generation.[9]

Autonomy and functional norms

In one respect we agree with the etiological approach: an account of normative function must provide a perspective or standard for distinguishing successful, irrelevant and dysfunctional behavior. Because it doesn't supply such a standard the method of functional analysis characterized by Cummins (1984) cannot by itself model normative function. However there is an alternative to using selection history as the benchmark for specifying norms; namely, using the system as the benchmark for norms. As noted above, there are difficulties with doing this, but a large part of the problem has stemmed from not finding an appropriate way to characterize the relevant systems. Our solution to this is a theory of autonomous systems.[10]

Autonomous systems are self-generating (or 'self-governed') in the sense that they interactively create the conditions required for their existence. There are two key aspects to this: firstly, autonomous systems are cohesive in the sense that they interact with the environment as a causally integrated whole, secondly, the conditions required for the cohesion of the system are, at least partly, generated by the system itself through the capacity to perform work. Rocks are cohesive but they are not autonomous systems in our sense because the cohesion conditions (a stable molecular lattice) are sustained after the processes that generated the cohesion conditions of the rock (such as magma cooling) have ceased. Consequently, if a rock is damaged it won't

---

[9] Any temptation to assume that this type of approach must be incorrect because parts are ontologically prior to systems, and hence are the appropriate locus of naturalistic explanation, should be tempered by the fact that selection operates directly on organisms, not parts. One of us (Bickhard) has argued at length elsewhere that quantum physics and considerations of general naturalism strongly support the view that process is ontologically prior to structure. Whether or not one follows the issues all the way through to physics and general naturalism, however, there are strong empirical reasons for treating process as prior to structure in biology. Traditional Chinese Medicine provides an interesting non-Western example of an ontology that places primary emphasis on process (see Gao, Herfel, and Rodriguez, forthcoming).

[10] The development of these ideas by the authors initially occurred independently, and we have subsequently combined our efforts to produce a unified approach. One of the legacies of this history is some difference in terminology. Bickhard developed a concept of self-maintenant systems (see Bickhard 1993), whilst Christensen developed a concept of autonomous systems that was subsequently refined in collaborative work with C.A. Hooker and J.D. Collier (see Christensen, Collier and Hooker 2000). Rosen's (1985) account of metabolic-repair systems, Ulanowicz's (1986) ecological concept of autonomy, and Collier's (1988) concept of cohesion influenced the formulation of Christensen's account of autonomy. There are minor differences between self-maintenance and autonomy that don't affect the discussion here; the choice of autonomy as the covering term in this paper is purely a matter of convenience.

perform work to reform itself. Moreover, if a rock is split at an arbitrary point the result is two smaller rocks with exactly the same type of cohesion conditions.

Now contrast a living cell, which is a paradigm autonomous system. Cells are far-from-equilibrium systems that require energy input from the environment, moreover they perform work to acquire this energy. In this respect the processes that generate the cohesion conditions of the cell are an ongoing part of the cell itself. The result is that cells *do* perform work to repair themselves. On the other hand, if the processes that perform the work are sufficiently disrupted the cell will die. This means that, in contrast with a rock, if a cell is split at an arbitrary point it usually won't result in two new cells. The processes that generate cellular cohesion stop and the cell collapses into its biochemical constituents. Autonomy involves positive feedback, but it is considerably more complex than simple feedback processes of the kind that plague PA systems. Positive feedback in a PA contributes to its own existence, but almost all of the work comes from the amplifier. In contrast, cells contain infrastructure that allows them to convert energy into useful forms and respond flexibly to perturbations. It is this infrastructure that makes them self-generating in a much more interesting way than PA feedback.

In addition to cells, other examples of autonomous systems include multicellular organisms, species, and cities. It may seem surprising to regard both organisms and species as autonomous systems, but species have the essential property of being self-generating systems. In this case the process that generates species cohesion is reproduction. There is much more to be said about the nature of autonomy. For detailed analyses of the central issues see Bickhard (1993) and Christensen, Collier and Hooker (2000). With respect to our current purposes the key point is that autonomous systems are the relevant class of systems for understanding normative functionality. Living organisms generally are autonomous systems, and autonomy is essentially a dynamical way of characterizing viability. Autonomy thus serves to provide the fundamental perspective for determining normativity: processes that contribute to the autonomy of the organism are useful to it.

As such, contributions to autonomy are the basic instances of *serving a function.* This model of serving a function is, first, not epiphenomenal: it makes a definite causal difference whether or not a particular far from equilibrium autonomous system continues to exist, or goes to equilibrium and thereby ceases to exist. Second, it is normative: such a contribution can succeed or fail in supporting the system, and this makes a distinct difference to the system, and to the world. The asymmetry of functional/dysfunctional is derived in this model from the fundamental physical asymmetry between equilibrium and far from equilibrium systems.

Autonomy and functional organization

Autonomous systems are composed of webs of interdependent processes. Modeling functional organization in our approach is a matter of characterizing these process interdependencies. To illustrate what this involves we shall consider a simplified case in which an autonomous system is composed of three processes: A, B and C. Because the system is autonomous, these three

process are involved in generating the conditions for the system's cohesion. Effectively this means that the conjoint operation of A, B and C is self-sustaining. Typically, A, B and C will be *organizationally interdependent* in the sense that each process requires output from the other two. This dependency will generally be tolerant within a range of variation, so a way of describing this interdependence is to say that for A to operate, B must be operating within a range of variation B_v and C must be operating within a range of variation C_v. Likewise for B and C. All of this, of course, will also depend on a set of environmental conditions $E_{(1...n)}$ with their own characteristic ranges. The parallels with Hutchinsonian niche theory should be clear, though it should also be apparent that this is a more detailed form of analysis that explicitly addresses the process dynamics of the system.

To apply this form of analysis to a real biological example, consider the process interrelationships between heart activity, cellular metabolism, and motor activity (including breathing, eating and excreting). As everybody knows, the function of the heart is to pump blood, or more accurately to pump blood as a contribution to an ensemble of activities that result in blood circulation. The function that this serves, however, is to provide fluid transport for delivering nutrients to cells and removing metabolic end products. In this respect heart activity and cellular metabolism are interdependent processes. Without heart activity fluid transport stops, and with it cellular metabolism. And if cellular metabolism ceases then heart activity also ceases, and subsequently fluid transport. In addition to heartbeat, cellular activity also produces other motor action that contributes to interaction processes such as breathing, food acquisition, eating and excreting. In turn these processes provide the resources required for cellular metabolism and expel waste products, thus contributing to the cellular processes that subserve them.

In our view these patterns of process interdependence in biological systems are the central issue for understanding normative function. Functions are essentially relations, and these process interdependency relations are what determine the nature of organisms as viable (cohesive) systems. Individual parts and processes serve normative functions within autonomous systems because of the way they satisfy the requirements of other processes within the system. This approach is able to account for the issues discussed above that are problematic for the etiological account. Firstly, it deals naturally with the issue of the relative significance of functional failure. In characterizing process interdependence we have for the sake of simplicity focussed on the situation where there is maximum interdependence, but patterns of interdependence will generally be complex and sometimes asymmetrical. Not every functional process is maximally connected with every other process. The relative functional significance of a process can therefore be analyzed by determining what other processes depend on it. Heart failure is lethal because it stops fluid transport and cellular metabolism, and *every* biological process depends on cellular metabolism. Mild vision impairment, on the other hand, has an immediate effect only on sensorimotor coordination. If the impairment isn't strong enough to compromise crucial sensorimotor processes such as food acquisition and predator detection, or if the animal can find ways to compensate, then there will be only limited flow-on effects to other processes.

Likewise, multifunctionality and functional complexes can be modeled in an informative way. The overall requirement of a complex process such as blood circulation is that it supplies the fluid transport requirements of cellular metabolism, and it can utilize inputs from many sources to achieve this, whether from systems that are specialized for a role in circulation, such as the heart, or from systems that aren't specialized for contributing to circulation, such as motor activity of the limbs. The central issue for modeling functional organization is the nature of the overall process; identifying parts that are specialized to play a role in the process, such as the heart, is a derived issue. The former does *not* collapse to the latter! Moreover, the fact that calf muscles are multifunctional by contributing to both walking and blood circulation poses no special problems. There is no requirement in our account that a part should play a role only in functions that it is specialized for. Indeed, it is to be expected in densely integrated functional complexes that parts will play roles in multiple processes.

However this account of normative function does not furnish an intentional-design concept of proper function of the kind sought by proponents of the etiological account. In this respect our account of dysfunction is highly unusual in the sense that the basic form of dysfunction is a process requirement that is not met. Thus, the fundamental dysfunction that occurs if the heart stops beating is that the fluid transport requirements for cellular metabolism are not met. Whether the heart is doing or failing to do what it was selected for is a secondary issue. For reasons of space we will not address the issue of what it is for a part to have a function in this paper. In this area we think there is a role for an etiological account, although we also believe that the issue involves more factors than current etiological theories allow for.[11] Our criticisms have thus not been aimed at rejecting all forms of etiological analysis, but are rather concerned with using etiology as the fundamental basis for understanding normative function. Nevertheless the issue of having a function shouldn't be overemphasized. Treating it as the main question leads to focus on parts at the expense of understanding the interrelations between the parts and the rest of system. At the start we cited Aristotle's claim that "The hand separated from the body is not a true hand", [12] and this has some genuine force on our account.

Finally, we note that our account avoids the standard problems associated with system-oriented approaches. Firstly, the account isn't prone to the circularity problems that affected cybernetic models. As noted above, it is not a simple feedback model because autonomy is a form of self-generation that employs internal infrastructure. One might worry, however, that it possesses the same kind of latent circularity as cybernetic accounts by characterizing norms in terms of processes that contribute to autonomy, and autonomy in terms of processes of self-generation of the system (which is normative). Thus, using autonomy as the explanation for norms is a

---

[11] The most comprehensive and biologically informed philosophical analysis of selection based function is Wimsatt (1972), and see also Wimsatt (1997). One suggestion for a non-selection factor in having function is degree of entrenchment In other words, if an organism has a part that has outputs which supply one or more crucial process requirements then the functioning of those processes *presupposes* the presence of the part, and hence it 'has' that function. For an elaboration on such a notion of functional presupposition, see Bickhard (in press).

[12] Also cited in Depew and Weber 1995, p.35.

concealed version of the fallacy of explaining norms in terms of goals, health or flourishing. However, the close association is grounded in the dynamical nature of autonomous systems and is not a definitional circularity. Autonomy essentially involves the co-occurrence of two dynamical conditions: (i) processes that perform work, and (ii) the overall cohesion of the system. Condition (ii) serves as the reference point or normative perspective against which the processes of condition (i) are evaluated. Autonomous systems are unique because they involve the co-occurrence of these dynamical conditions, but they are not definitionally circular because each dynamical condition can occur separately. Rocks are cohesive but they don't contain processes that perform work which maintains the cohesion of the rock, and there are plenty of processes that perform work but which don't contribute to their own cohesion.

Secondly, although our approach bears some strong similarities with Cummins (1984) account of functional analysis, it is a more constrained approach and these constraints are what allows it to model normative function. The key steps missing from Cummins are a principled specification of the relevant systems, thus providing the overall perspective for determining norms, and an account of process interdependence, thus restricting the analysis of dispositions to those that matter to the system. This rules out trivial dispositions such as the fact that hearts contribute to body mass (Sober 1993).

In this context it is worth examining the connection between our account and the defense of Cummins-style 'causal role' functional analysis mounted by Amundson and Lauder (1998). These authors argue compellingly that many kinds of biological research, including morphology and functional anatomy, depend on non-etiological forms of analysis. In this we agree with them strongly. We disagree with them in two respects: firstly, with their claim that the primary limitations on the application of etiological functional analyses are epistemic, and secondly, with their assumption that all forms of causal role analysis are non-normative. As our criticisms above are intended to show, the etiological account has fundamental limitations as a general model biological organization. Amundson and Lauder themselves point to pleiotropy and the fact that selection acts directly on performance parameters rather than parts as factors that limit etiological functional analysis. These problems are more than just epistemic.

More significantly for our account, however, whilst we agree that all of the forms of causal role functional analysis they characterize are non-purposive, in the sense that they don't appeal to design, we would classify some but not all of it as normative. Some of the research they describe is clearly not concerned with normative function on either our account or the etiological account, such as where a functional anatomist studies the properties of a muscle at 40 per cent of its rest length despite the fact that it never reaches this condition in vivo (p.235). However other forms of 'causal role' functional research *are* normative by our standards, or at least potentially so. In particular the study in functional morphology of 'functional integration', meaning "the interconnectedness of structures and their CR [causal role] functions" (p.252). Thus, they point out that a functional morphologist might study the interconnectedness of the components of the jaw, such as the way that in a particular species it may be the case that modification to a single muscle has deleterious effects on mouth opening (pp.252-3). This kind of analysis doesn't quite

involve modeling the full sense of process interdependence we characterized above, but it comes close. From the standpoints of theory and research methodology we neither wish to assimilate all causal role analysis to our account, or eliminate it. We simply suggest that the kind of analysis we have described is an important component within a suite of research approaches to understanding biological systems.

Amundson and Lauder also highlight an issue that we take to be a great strength of the kind of approach we (and they) advocate; namely, its ability to examine the significance of interconnected constraints for understanding patterns in phylogeny. They point to Lewontin's (1969) contrast between equilibrium and transformational explanatory approaches in biology. Equilibrium approaches study optimization, whereas transformational approaches study the effects of particular constraints on directions in evolutionary transformation.[13] As such, a transformational approach might examine questions such as whether a segmented body plan imposes constraints on subsequent evolution, or, in Raff's case, what the effects of developmental constraints are on evolutionary modification of early, middle and late development. With respect to these kinds of issues, the problem with the etiological account is that it is not historical enough.

Conclusion

There are two features of our account of normative function which make it highly unorthodox. Firstly, we abandon the design analogy as the general model for understanding biological organisms, and secondly, we shift from assigning functions to parts to understanding process organization. Despite its venerable history in natural theology and subsequent co-option by adaptationist neo-Darwinism, the design analogy is starting to outwear its usefulness. Although it can be a useful heuristic for some kinds of research it can also be highly misleading. Molecular developmental and evolution of development research is drawing a picture of organisms as complexly intertwined processes and constraints quite unlike anything humans have designed to this point (who knows what God would design), and certainly nothing like the pristine structure of a watch. Shifting the emphasis from parts to process also runs against the grain in the Western tradition, to the point where it looks like confusing the explanandum with the explanans. However there are strong reasons why doing so results in a more coherent and powerful account of normative function. Biological organisms are constituted by processes, and they persist because they are being continually remade by processes. These processes have a flexible holistic organization that eludes the localization approach of the etiological account, and they are typically shaped by a complex mixture of factors that cannot be captured by the single factor form of adaptationist explanation employed by the etiological account. If they are to be adequately understood then process organization must be modeled directly rather than be treated as the additive consequence of many parts that individually possess functions. Characterizing functional organization independently of selection in the way that we suggest does not make the

---

[13] Raff 1996 is a good example of a transformational approach. For philosophical discussion of the issue see Griffiths 1996, Sterelny and Griffiths 1999, ch.10

approach opposed to selection, it simply treats functional analysis as a legitimate component of a pluralist integrative approach to evolution rather than as a highly derivative method for describing the effects of selection.

Bibliography

Ahouse, J.C. (1998). The tragedy of apriori selectionism: Dennett and Gould on adaptationism. *Biology and Philosophy* **13**(3): 359-91.

Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., Watson, J.D. (1989). *The Molecular Biology of the Cell* (2nd ed.). NY: Garland.

Amundson, R. (1989). The Trials and Tribulations of Selectionist Explanations. In Hahlweg and Hooker (eds.) *Issues in Evolutionary Epistemology*. Albany: State University of New York Press, 413-432.

Amundson, R. and Lauder, G.V. (1998). Function Without Purpose: The Uses of Causal Role Function in Evolutionary Biology. In Hull, D. and Ruse, M. (eds.) *The Philosophy of Biology*. Oxford: Oxford University Press.

Bickhard, M.H. (1993). Representational Content in Humans and Machines. *Journal of Experimental and Theoretical Artificial Inteligence* **5**: 285-333.

Bickhard, M. H. (in press). Autonomy, Function, and Representation. *Communication and Cognition*.

Bigelow, J. and Pargetter, R. (1987). Functions. *The Journal of Philosophy* **84**: 181-196.

Christensen, W., Collier, J.D., and Hooker, C.A., (2000). Autonomy. In Christensen, *An Interactivist-Constructivist Approach to Adaptive Intelligence and Agency.* PhD thesis. Available at: http//www.newcastle.edu.au/department/pl/staff/WayneChristensen/dissertation.htm

Collier, J.D. (1988). Supervenience and reduction in biological hierarchies. *Canadian Journal of Philosophy*, supplementary volume: Biology and Philosophy, Calgary: University of Calgary Press.

Cummins, R. (1984). Functional Analysis. In E. Sober (ed.), *Conceptual Issues in Evolutionary Biology.* Cambridge: MIT Press.

Depew, D. and Weber, B. (1995). *Darwinism Evolving: Systems Dynamics and the Genealogy of Natural Selection.* Cambridge MA: MIT Press.

Gao, Y., Herfel, W., Rodrigues, D. (forthcoming). Making Sense in Western Terms: Traditional Chinese Medicine from a Nonlinear Dynamical Perspective.

Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Nous* **28**:344-362.

Godfrey-Smith, P. (1998). Functions: Consensus Without Unity. In Hull, D. and Ruse, M. (eds.) *The Philosophy of Biology.* Oxford: Oxford University Press.

Godfrey-Smith, P. (in press). Three kinds of adaptationism. In Orzack, S., Sober, E. (eds.) *Optimality and Adaptationism.* Cambridge: Cambridge University Press.

Griesemer, J.R. (1992). Niche: historical perspectives. In E.F. Keller & E.A. Lloyd (eds.), *Keywords in Evolutionary Biology*, Harvard.

Griffiths, P. (1996). The historical turn in the study of adaptation. *British Journal for the Philosophy of Science* **47**: 511-32.

Griffiths, P. (1997). *What Emotions Really Are.* Chicago: University of Chicago Press.

Hinde, R.A. (1966). Ritualisation and communication in rhesus monkeys. *Philosophical Transactions of the Royal Society* **251**: 285-94.

Hutchinson, G.E. (1957). Concluding Remarks. *Cold Spring Harbor Symposium on Quantitative Biology* **22**: 415-427.

Kitcher, P. (1998). Function and Design. In Hull, D. and Ruse, M. (eds.) *The Philosophy of Biology.* Oxford: Oxford University Press.

Lewontin, R.C. (1969). The Bases of Conflict in Biological Explanation. *Journal of the History of Biology* **2**: 35-45.

Lewontin, R.C.: 1970, 'The units of selection', *Annual Review of Ecology and Systematics* **1**: 1-14.

Miklos, G.L.G. (1993) Emergence of organisational complexities during metazoan evolution: Perspectives from molecular biology, palaeontology and neo-Darwinism. *Memoranda of the Australasian Association of Palaeontologists*, **15**, 7-41.

Millikan, R.G. (1989a). In Defense of Proper Functions. *Philosophy of Science* **56**: 288-302.

Millikan, R.G. (1989b). An ambiguity in the notion "function". *Biology and Philosophy* **4**:172-176.

Millikan, R.G. (1999). Wings, Spoons, Pills and Quills. *Journal of Philosophy* **96**(4): 191-218.

Mitchell, S. (1993). Dispositions or Etiologies? A Comment on Bigelow and Pargetter. *Journal of Philosophy* **90**(5): 249-259.

Newman, S. W. (1999). The Medial Extended Amygdala in Male Reproductive Behavior: A Node in the Mammalian Social Behavior Network. *Advancing from the Ventral Striatum to the Extended Amygdala*. *Annals of the New York Academy of Sciences*. Vol. *877*, 242-257.

Preston, B. (1998). Why is a wing like a spoon? A pluralist theory of function. *The Journal of Philosophy* **XCV**(5): 215-254.

Raff, R. (1996). *The Shape of Life: Genes, Development, and the Evolution of Animal Form*, Chicago: University of Chicago Press.

Rosen, R. (1985). Organisms as Causal Systems Which Are Not Mechanisms: An Essay into the Nature of Complexity. In R. Rosen (ed.), *Theoretical Biology and Complexity*, London: Academic Press, Inc.

Rosenbleuth, A., Weiner, N., and Bigelow, J. (1943). Behavior, Purpose, and Teleology. *Philosophy of Science* **10**:18-24.

Schaffner, K. (1993). *Discovery and Explanation in Biology and Medicine.* Chicago: University of Chicago Press.

Sober, E. (1993). *Philosophy of Biology*. Boulder: Westview Press.

Sober, E. (1998). Six Sayings About Adaptationism. In Hull, D. and Ruse, M. (eds.) *The Philosophy of Biology*. Oxford: Oxford University Press.

Sterelny, K. (forthcoming). Darwin's Tangled Bank.

Sterelny, K. and Griffiths, P. (1999). *Sex and death: An introduction to philosophy of biology*. Chicago: University of Chicago Press.

Ulanowicz, R. (1986). *Growth and development: ecosystems phenomenology*. New York: Springer.

Wimsatt, W.C. (1972). Teleology and the Logical Structure of Function Statements. *Studies in the History and Philosophy of Science* **3**: 1-80.

Wimsatt, W.C. (1997). Functional Organization, Functional Analogy, and Functional Inference. *Evolution and Cognition* **3**(2): 102-132.

Wimsatt, W. C. and J. C. Schank (1987). Generative entrenchment and evolution. In Fine, A. and Machamer, P. (eds), *PSA 1986.* East Lansing, Mich.: Philosophy of Science Association.

Wright, L. (1973). Functions. *Philosophical Review* **82**: 139-68.