TELEOLOGICAL REALISM IN BIOLOGY

A Dissertation

Submitted to the Graduate School

of the University of Notre Dame

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

by

James Alan Barham

_____

Phillip R. Sloan, Director

Graduate Program in the History and Philosophy of Science

Notre Dame, Indiana

July 2011

TELEOLOGICAL REALISM IN BIOLOGY

Abstract

by

James Alan Barham

The concept of purpose, together with the closely related concepts of normativity and agency, stand at the crossroads of three academic disciplines: the philosophy of action, the philosophy of biology, and the nexus of theoretical biology and cognitive science that is concerned with the theory of the "self-organization" of "autonomous agents." Unfortunately, there has been very little cross-fertilization among the literatures of these three disciplines. As a result, the philosophical literature tends to work with a scientifically outdated image of living things as rigid "machines." This results in a picture in which only human beings (or at most the higher animals) can be properly ascribed purposes and agency in the full normative sense. From this perspective, we appear to be faced with an unappealing choice between eliminating teleology and normativity from our picture of nature altogether and understanding these phenomena as they are manifested in our own human form of life as floating free from any grounding in the natural world. The scientific literature, on the other hand, tends to misuse "teleology," "normativity," "agency," and related terms, mistakenly ascribing such concepts to

"autonomous agents" conceived of as subject only to the ordinary laws of physics. From this perspective, the true depth of the difficulty involved in understanding what makes living systems distinctive qua physical systems becomes occluded.

In this dissertation, I investigate the possibility of constructing a realistic view of immanent teleology in biology. I proceed by exploring each of the three literatures in turn, with the goal of finding a middle way between the extremes of eliminativism and dualism. The argument proceeds by analysis of the concepts of teleology and normative agency, by reflection upon the explanatory structure of the theory of natural selection, and by review of some contemporary scientific accounts of "self-organization" and "autonomous agents," as well as of other physical features of living things.

My overall conclusion is that the acceptance of teleological realism in biology is rationally permitted. In other words, teleological realism in biology ought to be viewed as a "live option."

To the memory of my mother

CONTENTS

ACKNOWLEDGMENTS

First, I would like to express my deepest gratitude to Phillip R. Sloan. If I had not had the benefit of his deep learning, great generosity, and near-infinite patience over these many years, this work would not exist. I have also profited greatly from his close reading of numerous drafts.

Next, I would like to thank Lenny Moss, who picked me up, dusted me off, and put me back on my feet. His comments on versions of two chapters were a great help to me, as well.

Finally, I would like to thank Robert Audi and Grant Ramsey for sharing their comments on particular chapters with me.

Needless to say, none of the above-mentioned individuals should be held in any way responsible for my persistence in error.

CHAPTER 1:

INTRODUCTION

1.1 Statement of Principal Aims

In this dissertation I will be exploring the possibility of taking a realistic stance towards teleology in biology—that is, towards teleology in relation to organisms, or to living things as such. By "teleology," I have in mind such words and concepts as "purpose," "end," "goal," "function," "control," and "regulation," as well as the real-world biological phenomena to which these words and concepts refer. This means that the word "teleology" should always be construed here in its internal or "immanent" sense—purposiveness existing in living beings themselves—and never in its external or "transcendent" sense of an overarching cosmic principle.[1]

More specifically, I will be exploring the idea that teleological discourse, both in everyday life and in biological science, is best understood as describing real, objectively existing properties of biological systems, quite apart from human interests and conceptual schemes. This means the claim that—assuming a broadly realistic stance towards the objects of everyday and biological discourse—teleological phenomena should be viewed as being on an ontological par with such well-known biological phenomena as

[1] On this distinction, see Lennox (1992).

locomotion, phagocytosis, photosynthesis, and the like. I will refer to this basic claim as "teleological realism in biology" (**TRB**, for short).

I will attempt to show over the course of this dissertation that **TRB** is rationally defensible. I will not try to show that the evidence in favor of **TRB** is so strong that it is to be positively preferred to other views. In particular, I will not try to prove the superiority of **TRB** to what I will later on be calling the "teleoreductive" view.[2] In other words, I will only claim that **TRB** is rationally permitted, not that it is rationally required. That is, I only intend to show that it ought to be viewed as what is sometimes called a "live option."

The claim that **TRB** is rationally defensible will be defended by means of three subsidiary claims. The first of these is that the family of teleological concepts (purpose, end, goal, function, control, regulation, etc.) can only be properly understood in relation to the broader family of normative concepts, including need, well-being, value, reasons, and the instrumental "ought."[3] Furthermore, I will argue that the concept that binds the teleological and the broader family of normative concepts together is that of agency (the capacity for action). More specifically, I will argue that because the teleological concepts and certain normative concepts (like need and well-being) are clearly applicable to all organisms, or to living things as such, and because teleology only makes sense in connection with agency, the concept of agency in the full normative sense is properly applicable to organisms, or living things as such, as well. Moreover, I will claim that this

[2] In a nutshell, this simply means that the appearances of teleology in biology can be satisfactorily explained through entirely non-teleological means. The concept of "teleoreduction" will be developed in more detail in Chapter 3, below.

[3] This generic use of the term "normative" will be explained and defended at length in Chapter 2.

a priori argument accords with the a posteriori observation that what distinguishes living from non-living systems is that only the former must do work in order to maintain themselves in existence.[4] More specifically, I will claim that normative agency consists in the intrinsic ability of a system to do work for a reason of its own, that all living systems and only living systems can have reasons of their own, and therefore that all and only living systems qualify as normative agents. If this claim can be justified, then normative agency may be viewed as the essential feature of the physical state or condition of being alive—i.e., of life considered as a natural kind. The a priori part of this claim will be defended in Chapter 2, and the a posteriori part in Chapter 4.

The second subsidiary claim is that there is good reason to doubt that either of the two most frequently cited schemes for reducing teleology to mechanism succeeds in that task. By "reducing" teleology to mechanism, once again, I merely mean the claim that the reductive scheme in question provides a theoretically and empirically adequate account of the biological phenomena, without having recourse to any teleological or normative concepts. The two purported reductive schemes I have in mind are the interpretation of biological functions (1) as causal contributions to a system and (2) as selected effects. These two schemes will be evaluated in Chapters 3.

By the end of Chapter 3, it will have been established that there are excellent reasons for regarding organisms as normative agents, and that the success of the two best-known purported teleoreductive schemes is open to doubt. These two claims together already provide good grounds for accepting **TRB**. However, it is one thing to give

[4] Note the normative ("must") and the teleological ('in order to") language that unavoidably imbues our most basic descriptions of a living thing. My basic claim is that this is no accident, but rather reflects a deep truth about the real nature of living systems.

conceptual arguments in support of the existence of a phenomenon, and it is something quite different to give such an existence claim empirical support and to show how the phenomenon might begin to be understood in theoretical terms that are consonant with the rest of empirical science. In short, the discussion through Chapter 3 does not yet grapple with the question: How are normative agents possible? This means that, in order for the overall claim of **TRB** to be persuasive, I must provide support for a third and final subsidiary claim: namely, that the objective existence of the teleological phenomena associated with normative agents is, although far from proven, at least clearly conceivable in light of a number of recent empirical discoveries and theoretical advances in biology and physics.

Considering the importance of this empirical claim for my overall argument in support of **TRB**, it should be clear that my project is strongly naturalistic in orientation. I will simply assume without further discussion that it is preferable, where possible, to avoid appealing to a transcendent or supernatural principle in order to explain certain features of the world,[5] and that an account which somehow integrates all of the world's features into a single coherent picture is for that reason preferable to any dualistic or pluralistic account which accepts as a brute fact two or more fundamental principles whose relations to one other remain unknowable even in principle. By "naturalism," then, I have in mind a metaphysical picture of the world as both exhaustive of reality and somehow unified. The question, of course, is whether it is possible to be both naturalistic in this sense and also realistic about teleology in biology (i.e., normative agency), and if

[5] Where all features of human beings, as well as of other living things, are conceived of as features of the world.

so, how. In other words, I must attempt not only to demonstrate that there is good reason to believe that teleological phenomena in biology are objectively real, but also to show how it is at least conceivable that normative agency could be fully at home in the natural world. Discussion of these matters will be the principal concern of Chapter 4.

1.2 What Is the Problem of Teleology?

Before turning directly to these claims and their supporting arguments, I will spend the rest of this introductory chapter making some important preliminary distinctions: explaining in more detail what my project involves and what it does not involve, situating it within the contemporary intellectual landscape, distinguishing it from some other projects with which it might be confused, and generally attempting to explain and motivate the present undertaking. I will begin by distinguishing between a narrow and a broad sense in which teleology poses a problem for philosophical reflection.

In the narrow sense, teleological discourse seems to apply to a range of heterogeneous cases—including human intentions and actions, manmade artifacts, social institutions and practices, and biological phenomena of every conceivable sort—and one would like to be able to say what exactly it is that all these cases have in common. This may already seem like an ambition of generous scope, but I call this the "narrow" sense of the problem of teleology, because analyses pursued in this spirit typically limit themselves to the investigation of our linguistic practices, without inquiring into the underlying metaphysical reasons for those practices. One might also think of this as the "inner" problem of the relations among the various things to which our teleological concepts refer.

I take the broad problem of teleology to be, in contrast, the "outer" problem of accounting for teleology as such, that is, the problem of understanding how such a thing as teleology is possible at all, given the rest of what we know about the world. In a nutshell, the broad or outer problem of teleology is constituted by the following dilemma. On the one hand, we cannot avoid having recourse to teleological concepts in our effort to understand the living world—and by "we," I mean all of us, both ordinary people and biologists. On the other hand, teleological concepts are not sanctioned by the exact sciences that we have developed for understanding, so successfully, the nonliving world. Of course, this is only a dilemma to the extent that we accept two further ideas: (1) the unity of the world; and (2) the rational priority of the concepts of physics and chemistry over those of biology, leading in turn to the idea that living things are "nothing but" physics and chemistry. Given these two widely shared assumptions, then, the broad or outer problem of teleology consists of understanding how the indispensability of teleological concepts for our understanding of the living world can be reconciled with their absence from our understanding of the nonliving world. This is more or less the same thing that Buller (1999a) has recently characterized as the "metaphysical problem of teleology": "Ever since the rise of [the] scientific world view, the metaphysical problem of teleology has been that of explaining whether, and if so how, there can be goal-directed processes in a universe governed solely by efficient causation" (Buller, 1999a; p. 6).

The claim of **TRB**, then, clearly falls under the heading of the broad or outer problem of teleology. That is, my chief concern in this dissertation will be with the question of how the teleological character of living systems may best be understood in

relation to inanimate matter. This does not mean, however, that the narrow or inner problem of teleology will be of no concern here at all. Obviously, our understanding of the relation between teleology in biology and the rest of nature must be determined, at least in part, by our understanding of what teleology in biology actually is—that is, what its nature consists of. In other words—moving from the epistemological to the metaphysical register—the outer relation and the inner relation must both contribute to determining, and be determined by, the nature of a living thing. So, I must be concerned with the inner problem of teleology, as well as the outer one. However, this concern will take a quite different form in this dissertation from the usual one. For example, I will entirely set aside the topic of the teleological character of manmade artifacts. The reason is that manmade artifacts derive their teleological features secondarily from human intentions and actions, and so have little or nothing to tell us about what teleology is as a non-derivative or primary (or original or fundamental) phenomenon. What I will take up instead is the question of the relation between teleology and other normative concepts such as value, need, well-being, reasons, and the like. In short, the focus of this dissertation, insofar as the traditional narrow or inner problem of teleology is concerned, will be on the role that teleology plays in action. This will be the primary topic of Chapter 2. But the reason for posing the inner problem in this way consists entirely in whatever light it may shed on the outer problem—the problem of understanding how the teleological character of living systems makes them different from nonliving systems.

For this purpose, it will suffice to look at two broad ranges of phenomena: the teleological features of human action, in particular, and those of living systems, in general. The former must be considered because many of our teleological concepts have

their clearest application in the case of human action, but the tendency of this study throughout will always be to advance to the general case of teleology in biology, that is to say, in living systems, or organisms, as such.

1.3 Teleology at the Intersection of Three Disciplines

There are three principal intellectual disciplines, with their concomitant literatures, which bear upon the subject of this dissertation. One is the philosophy of action; another is the philosophy of biology; and the third is a recent trend within empirical science that itself lies at a disciplinary intersection (of theoretical biology and cognitive science) and which consists of such relatively new methodological approaches and programs as the dynamical-systems modeling of physiology and behavior, ecological psychology, embodied cognition, enactivism, situated robotics, and the like. Teleology is a subject of interest in all three of these disciplinary areas, and, significantly, in the first and last of them the concept of agency has recently begun to assume a central role (agency has of course always been the central focus of the philosophy of action). Let us examine each of these in turn, albeit in the briefest of overviews.

For ease of exposition, I begin with the philosophy of biology literature, although my discussion of that literature (in Chapter 3) will follow my discussion of the philosophy of action literature (in Chapte 4). Needless to say, teleology has often been at the center of philosophical reflection upon the phenomenon of life from Aristotle onward (Gilson, 2009; Nordenskiöld, 1928; Pichot, 2004). Over the past half century or so, and especially since the institutionalization of philosophy of biology as a sub-discipline of academic philosophy within the English-speaking world during the 1960s and 1970s, the

problem has tended to be discussed most often in connection with the concept of "function" (Allen et al., 1998; Ariew et al., 2002; Buller, 1999b; Wouters, 2003, 2005). More specifically, for several decades now various analyses of biological function have been proposed and debated, with the aim of showing how the biological phenomena to which our teleological concepts refer can be accounted for in purely non-teleological terms. As is well known, two main positions have attained a dominant position within the literature. The first position, stemming from a seminal article by Cummins (1975), views being a function fundamentally as making a causal contribution (in the efficient-causal sense) to the maintenance of a larger system of which the function in question is a component part. The other position, adumbrated by Wright (1973) and clearly articulated by Millikan (1984), takes a present trait's being a function to be equivalent to its having been naturally selected due to the fitness advantage conferred on an organism by the physical effects of the ancestral trait of the same type from which the present trait-token is descended. It has also been suggested that it may be necessary to combine elements of both analyses (e.g., Kitcher, 1993). I will be examining both analyses in Chapter 3. Here, I would just like to point out that within this literature, until very recently, the problem of how to analyze functions had come to be seen as tractable, if not fully solved. For example, Buller (1999a; p. 1) writes: "Within the past decade a near-consensus has emerged among philosophers concerning how to understand teleological concepts in biology." He goes on to explain that while disagreements do of course persist, "broad agreement about certain fundamental commitments can be discerned in the recent literature" (ibid.). Needless to say, the consensus in question was a consensus in support of a reductive analysis of function (where, to repeat, a reductive analysis of a teleological

feature of a biological entity is one that provides an adequate account of the feature in efficient-causal terms without appealing to or tacitly assuming any teleological concepts). In other words, the consensus that Buller reports on was a consensus around the projected future success of some version of either the causal-contribution or the selected-effects reductive analysis, or some combination of the two.

It turns out, however, that Buller's announcement of a reductive consensus with respect to teleology in the philosophy of biology literature was premature. Ironically, it was articulated at the very moment when the consensus began to show serious signs of strain. In a series of important articles and books over the past decade or so, Bedau (1990, 1992a, 1992b, 1993), Cameron (2004), Christensen & Bickhard (2002), Jacobs (1986), Manning (1997), Maund (2000), McLaughlin (2001, 2009), Mossio et al. (2009), Mundale & Bechtel (1996), Nanay (2010), Nissen (1997), Perovic (2007), Walsh (2006), and Zammito (2006) have cast grave doubt on the coherence of any reductive analysis of function. Some of these authors (e.g., Jacobs, Maund, Zammito) call explicitly for a reconsideration of the possibility that teleological phenomena in biology might be both objectively real and irreducible. Some others (Cameron, Christensen & Bickhard, McLaughlin, Mossio et al., Perovic, Walsh) go so far as to suggest an alternative interpretation of teleology in biology in terms of concepts borrowed from dynamical systems theory.[6] The positive proposals that I will be making regarding teleology in this dissertation have a definite kinship with the work of this last group of thinkers, several of whose ideas will be explored in Chapter 4. Some, but not all, of the arguments of the

[6] There is some overlap between this last group of philosophers and the empirical scientists comprising the third disciplinary group I will be discussing in a moment; Bickhard, Moreno, and colleagues, in particular, have a foot in each camp.

other authors, which are naturally quite varied, will be discussed as well, mainly in Chapter 3. For now, the only point I wish to make is that the reductionist consensus in the philosophy of biology is no longer so monolithic as it appeared to Buller to be to a decade ago. The present project takes its place naturally within the recent and growing anti-reductionist trend that is calling that consensus into question.

Remarkably, a similar story may be told with respect to the second intellectual discipline and its literature: namely, the philosophy of action. Here, too, a reductive consensus had held sway for some time, and then rather recently began to falter. In a nutshell, Davidson (1963) had argued that normative reasons for action must ultimately be understood as a species of "cause" (in the efficient-causal sense) on pain of rendering unintelligible the determinative force of reasons. As he put it (Davidson, 1963; p. 691): "Central to the relation between a reason and an action it explains is the idea that the agent performed the action because he had the reason" (original emphasis).[7] This simple argument had an enormous impact on the English-language literature on action and allied notions such as practical rationality, value, and normativity. As Alvarez (2010) has recently noted, "Davidson's conception of reasons, or something close to it, became the orthodoxy and remains so to this day" (Alvarez, 2010; p. 2). But while the idea that reasons are a species of cause may remain the orthodoxy, it is an orthodoxy that finds itself under increasing assault. Among the many recent authors who question the intelligibility of interpreting normative reasons for action in this way, and who insist on the irreducibly teleological character of reasons, one may mention Alvarez (2010), Boyle

[7] Hereafter, all emphasis in cited passages should be assumed to be original, unless otherwise stated.

& Lavin (2010), Burge (2010), Foot (2001), Hacker (2007), Hanna & Maiese (2009), O'Connor (2000), Okrent (2007), Parfit (forthcoming), Rundle (1997), Schueler (2003), Sehon (2005), M. Thompson (2008), Wedgwood (2007), and G.M. Wilson (1989).[8]

Once again, of course, we have to do with a wide spectrum of nuanced positions supported by a wealth of varied arguments. While I will be examining some of these arguments in detail in Chapter 2, I must pass over the majority of them in silence. Here, I wish only, as before, to emphasize the fact that a previously existing reductive near-consensus in a philosophical discipline has begun to disintegrate in recent years. The upshot is that drawing on teleological concepts related to action in support of **TRB** is not now as radical a move as it might have seemed just ten years ago. Needless to say, I will have to do a great deal of work in order to justify what still can only seem to be a very ambitious move. Nor would very many, if any, of the authors listed wish to follow me all the way down my path. Nevertheless, it is striking that this sea change has come about. In a sense, I will only be exploring a little farther down a path that has already been opened up by others. In particular, I will be attempting to respond in my own way to the challenge facing any sort of teleological realism, whether biological or otherwise, that Davidson (1963) placed in such high relief:

> One way we can explain an event is by placing it in the context of its cause; cause and effect form the sort of pattern that explains the effect, in a sense of 'explain' that we understand as well as any. If reason and action illustrate a different pattern of explanation, that pattern must be identified. (ibid.; p. 692)

I take this challenge to heart throughout this dissertation and I attempt to meet it full-on in Chapter 4, below.

[8] This list includes only major monographs. There are a number of important journal articles arguing not only for teleological realism in human action, but also for the propriety of ascribing at least some action concepts to at least some non-human animals. I will be reviewing this literature in Chapter 2.

Finally, the third disciplinary area impinging on our topic is a bit harder to define. It is itself constituted by interdisciplinary work originating from a number of traditional disciplines within the empirical sciences, including theoretical biology, experimental psychology, and cognitive science. The body of work I have in mind goes under a variety of sub-disciplinary names, including "ecological psychology," "embodied cognition," "enactivism," and "situated robotics," but a unifying factor among these approaches is that all of them take it as part of their brief to try to understand two things in particular: the phenomenon of "self-organization," and what it is for a physical system to constitute an "autonomous agent." This is not to say that these schools all approach this question in the same way; rather, each views the problem of autonomous agency from its own particular angle. But it does seem that the idea of the objective reality and irreducibility of autonomous agency is taking hold in some empirical-science disciplines, a development that potentially constitutes a challenge to the reigning reductionist consensus in the natural sciences that is every bit as weighty as the challenges posed to the similar consensuses in the philosophy of biology and the philosophy of action.

The empirical-science reductionist consensus with respect to teleology has been expressed pithily by Crick (1966; p. 10): "The ultimate aim of the modern movement in biology is in fact to explain all biology in terms of physics and chemistry"—where "physics and chemistry" are of course understood as comprising exclusively efficient causes. The notion that biological systems are subject to a sui generis teleological causal principle is often derided as "vitalism," and it is undoubtedly fair to say, as Mayr (1982; p. 52) has put it, that "for biologists vitalism has been a dead issue for more than fifty years."

For the sake of clarity on this point, it might be helpful to distinguish two senses of "vitalism." In a narrow sense, a position would count as "vitalist" if it distinguished sharply between at least some of the properties of living things and those of non-living things. In a broad sense, a position would count as "vitalist" if it both accepted the narrow position and further claimed that the distinctive properties of living things were not "grounded" in the microstructure of organisms.[9] According to this distinction, **TRB** is clearly "vitalist" in the narrow sense, but clearly not "vitalist" in the borad sense.

After all, the idea that the apparently teleological character of organisms can be explained by the real existence of a causal principle different from ordinary causation and unique to living systems would only be "vitalist" in a broad sense if the principle in question were viewed as beyond the reach of empirical investigation. However, to maintain that such a sui generis causal principle really exists and that empirical science must and can expand its conceptual boundaries in order to accommodate that reality—such a claim is not vitalism—at least not in the broad sense, which I believe is the traditional, pejorative sense of the word—and there is no good reason for it to be dismissed out of hand. Of course, he who would espouse "vitalism" even in the narrow sense still has a heavy burden of proof. But at bottom, the claim is an empirical one, and should be evaluated like any other such claim on the basis of evidence and argument.

A number of working scientists have come to the conclusion in recent years that the evidence favors this claim. The situation is not yet so striking as in the two philosophical fields already canvassed; it would be premature to speak yet of the

[9] The concept of "grounding" will be discussed in Chapter 4, below. In a nutshell, it just means that "Its being the case that S consists in nothing more than its being the case that T, U, . . ." (Fine, 2002; p. 23).

reductionist consensus in biology as under severe strain, sociologically speaking. Nevertheless, it is striking that a range of investigators in a variety of disciplines have come to view teleology as objectively real and irreducible. Sometimes, their views are couched in terms of "autonomous agency" or "normativity," but the implication for teleology itself is clear. And often enough, it is quite explicit. I have in mind the work of such thinkers as Barandiaran & Moreno (2008), Barandiaran et al. (2009), R. Campbell (2009), Chemero (2009), Christensen & Bickhard (2002), Christensen & Hooker (2001), Di Paolo (2005, 2009), Di Paolo & Iizuka (2008), Di Paolo et al. (2010), R.D. Ellis & Newton (2010), Freeman (2001), Hanna & Maiese (2009), Hooker (2009a, 2009b), Juarrero (1999), Kauffman (2000, 2004), Kauffman & Clayton (2006), Kauffman et al. (2008), Kelso (2008), Moreno & Ruiz-Mirazo (1999), Mossio et al. (2009), Ruiz-Mirazo et al. (2010), Skewes & Hooker (2009), E. Thompson (2007), and A. Weber & Varela (2002). I will be referring to these authors collectively as "Autonomous Agency Theorists." As before, it is necessary to acknowledge a considerable diversity of approaches among them. But what they have in common is a preoccupation with using the concept of "self-organization" to explain the phenomenon of "autonomous agents"—that is to say, organisms conceived of as endowed with autonomy and normative agency.

To help orient the reader toward this literature, it might be useful to mention that the Autonomous Agency Theorists have a good deal in common with the contemporary movement known as "systems biology." For example, here is an authoritative description of the goals of systems biology by an eminent participant in that endeavor:

> The essence of life must lie somewhere between molecule and autonomously living unicellular organism. Modern biology generally views organisms as beads along the necklace of lineage; it attempts to explain life from an evolutionary viewpoint, with reproduction (of cells) and replication (of DNA) as defining

> phenomena. Systems biology, however, studies each bead per se as an autonomous entity. I suggest that, for systems biology, the defining difference between a living organism and any nonliving object should be that an organism is a system of material components that are organised in such a way that the system can autonomously and continuously fabricate itself, i.e. it can live longer than the lifetimes of all its individual components. Systems biology, therefore, goes beyond the properties of individual biomolecules, taking seriously their organisation into a living whole. (Hofmeyr, 2007; p. 217)

Here is a similar viewpoint, this time expressed by an eminent philosopher:

> My contention is that recognizing organization does not require a rupture with the tradition of mechanistic science. Mechanism has the resources to identify and incorporate the forms of organization critical in living systems. Moreover, attempts to focus on organization independently of the matter and energy of actual systems are likely to fail, as the organization required to maintain autonomy is an organization that is suited to the matter and energy available to the system. It is in this context that the notion of basic autonomy reveals its importance: it provides a framework for relating organization tightly to the matter and energy of the system as the organization of interest is one which, given the energy and material to be utilized, is able to be built and maintained by the living system. (Bechtel, 2007; p. 297)

This characteristic note of concern for global or system properties, tightly coupled to an assumption of naturalism, if not mechanism, is typical both of contemporary systems biology, in general, and also of the Autonomous Agency Theorists, in particular. The only real difference is that systems biology in the ordinary sense usually restricts itself to more tractable problems, such as the modeling of particular metabolic circuits, whereas the Autonomous Agency Theorists have taken for their object of study the deepest and most difficult problem in biology—the problem of penetrating to the physical principles underpinning the normative agency of living things as such.

I will be paying close attention to some of the arguments of some of the Autonomous Agency Theorists, in Chapter 4. The main reason I will do so is to show that **TRB** is no longer unthinkable, even among empirical scientists themselves.

The problem of teleology in biology is a peculiar one for a number of reasons. Not the least of these is that it is located at a point of intersection among the three intellectual disciplines we have been examining in this section. This point of intersection itself has been largely overlooked up until now. I know of no study that draws on all three literatures in the way that I will be doing here. And yet I am convinced that there is potentially great utility in doing so, and that each of the three perspectives may be significantly enhanced by being put into contact with the other two. Moreover, I find it encouraging that in all three disciplines, a movement in the direction of **TRB** is already clearly discernible. My hope is by that synthesizing insights from all three fields in a new way, I may be able to lay the foundation for a real advance in understanding teleology, normativity, and agency as objectively existing phenomena.

1.4 The Phenomenology of Biological Function and Function-Discourse

In any study conducted at a considerable level of generality and abstraction, it is a good idea to anchor the discussion in some concrete examples. I will begin this section, accordingly, by describing a fundamental biological function, for use as a suitable touchstone against which our intuitions and other sorts of evidence concerning functions may be tested. I will then examine some of the basic ways we talk about functions, both in everyday speech and in biological practice, in order to show both the continuity of our function discourse between both domains and its patently teleological character in each of them. More detailed probing of the issues raised in this section will follow in the each of the chapters to follow; the purpose of this section is simply to establish a preliminary case for taking the teleological character of function-talk in biology seriously.

Since my project aspires to complete generality with respect to the domain of living systems, the choice of a biological function to use as an example for the sake of concreteness ought to be guided by this aspiration. For this reason, I choose one of the most fundamental capacities of the most primitive sort of living system known—namely, bacterial chemotaxis. While this capacity is unique in its details to certain bacteria, it incorporates several basic features of biological organization that are universal in extent, or nearly so.

Like all living systems, bacteria must take in certain specific types of molecular materials from their environment in order to meet their basic metabolic demands.[10] Furthermore, some chemical features of the environment (e.g., too-high or too-low pH-level) may be deleterious to bacteria. This means that in order to survive, bacteria must find and ingest certain types of molecules and avoid others. In other words, like all living things, bacteria have certain requirements or needs that must be met. These needs are met by "feeding" and "fleeing" behaviors—that is, movement towards food sources ("attractants") and away from chemicals posing danger ("repellants"). The organismic subsystem by means of which these behaviors are executed is called "bacterial chemotaxis." The following is how bacterial chemotaxis works, in broad outline.[11]

The first requirement of bacterial chemotaxis is to distinguish attractants from repellants, and thus to discern which spatial directions have positive and negative valence

[10] I take it that this description is uncontroversial. Notice, however, that as soon as one begins to describe how a living thing works, one is immediately plunged into a universe of teleological discourse — "must," "in order to," "demands" — that has no counterpart in chemistry or physics.

[11] For brief descriptions and interpretative discussion, see Shimizu & Bray (2003); Wadham & Armitage (2004); and Webre et al. (2003); for full technical details, see Stock & Surettte (1996).

for the organism. That is to say, perception is logically prior to response.[12] Since bacteria are motile organisms that live in fluid environments in which most molecules of interest diffuse in gradients, the problem of perception translates into the problem of evaluating gradients as positive or negative.[13] In the case of bacterial chemotaxis, this problem takes the form of sampling the environment repeatedly and comparing samples taken at different times. If the comparison shows that the organism is moving up a positive gradient (i.e., the concentration of attractant is increasing with time) or down a negative gradient (the concentration of the repellant is decreasing with time), then the appropriate response is to stay the course. In the converse case, the appropriate response is to change course, the new direction being determined randomly. Finally, when the environment is neutral (presenting neither attractant nor repellant gradients), the problem for a bacterium is to find a positive gradient. Under these circumstances, the animal also moves about in a random fashion. Therefore, with respect to motility bacteria have two requirements: maintaining a given course by moving in rectilinear fashion and changing course by changing direction in a random manner.

Accordingly, bacteria are equipped with a means of locomotion, namely, a set of proteinaceous fibers called "flagella" that extend from the outer membrane into the extracellular medium and are capable of rotatory motion. The flagella have, as it were, two gears that lead to two motile states: steady-as-she-goes and picking a new heading by chance. These two states of motion are achieved by counterclockwise and clockwise

[12] Though what counts as appropriate perception is of course partly dependent upon what counts as appropriate response—more on all of this in Chapter 2, below.

[13] Stuart Kauffman glosses this bivalent partitioning of the environment as the originary axiological categories of "yum" and "yuck" (Kauffman, 2000; Kauffman & Clayton, 2006; Kauffman et al., 2008).

rotation of the flagella, respectively. When rotating counterclockwise, the flagella form a coherent bundle that direct the bacterium in rectilinear fashion (called "running"). When rotating clockwise, the flagella flail about incoherently, causing the bacterium to move in a random direction (called "tumbling"). When the environment is sense to be neutral, short runs are frequently punctuated by tumbles. When an increasing attractant gradient (or decreasing repellant gradient) is encountered, tumbling is suppressed, and rectilinear running is extended in time.

So far, I have been describing the functional logic of bacterial chemotaxis quite freely in terms of such normative locutions and concepts as "must," "in order to," "demand," "need," "survival," and so on. Needless to say, the particular workings of this behavioral subsystem can be described in quite other terms. Several generations of persistent research have yielded a wealth of information about the "mechanisms" involved in the perception and response systems. For instance, we know that perception is mediated by a special class of protein "receptors" embedded in the outer membrane. We know the composition of the "motor" which rotates the flagella in both directions. We even know in great detail the "wiring diagram" connecting the sensory system to the effector system by means of numerous classes of "second messenger" molecules.

One the main questions to be investigated in this dissertation is whether this second sort of description, which I shall call a "mechanistic" description, is capable by itself of constituting an adequate explanation of a biological phenomenon like bacterial chemotaxis, or whether the first set of normative descriptors is in some way

indispensable.[14] However, my purpose at present is merely to provide a concrete example of the kind of natural phenomenon we will be investigating herein.

Let bacterial chemotaxis, then, be our paradigm case of a biological function. If it can be established that our teleological descriptions correspond to something objectively real in this case, then we would have excellent reason to believe that teleology is an objective fact in all cases of biological functions whatsoever. However, it might be objected that bacterial chemotaxis is above all the province of the professional biologist, while the teleological discourse I have been using so freely is merely a holdover from our pre-scientific way of talking about organisms. Perhaps the teleological descriptions are merely a matter of convenience, and the mechanistic description is the only one to which any kind of objective reality ought to be attributed. In order to meet this objection, I will now examine how we talk about some more familiar, intuitively obvious examples. Then, I will show that there is no principled way to distinguish between the familiar examples and the cases that clearly lie within the domain of biological science.

Consider the following series of questions and answers:

Q. (Pointing to an eye) What is this?
A. It's an eye?
Q. What's it for?
A. It's for seeing.

Q. (Holding up a hand) What is this?
A. It's a hand.
Q. What's it for?
A. It's for grasping.

Q. (Lifting a leg) What is this?
A. It's a leg.

[14] Where a "mechanistic" interaction is simply one involving only "efficient" causation, i.e., one not involving teleology.

Q. What's it for?
A. It's for walking.

These seem to be perfectly intelligible things for anyone to say. A tad artificial, perhaps, because taken for granted by all adult language speakers. But we can easily imagine such a give and take occurring. Such questions and answers might even occur naturally in the context of practicing vocabulary with a young child or perhaps teaching it to a foreigner. What one cannot imagine is anyone being puzzled by the questions, simply not knowing what was being asked of them.

These observations show several things. One thing they show is that the teleological character of body parts seems to be analytical. When we speak of an "eye," we do not just mean a sphere located in the front of our skulls. We do mean that, of course, but we also mean the seat of the faculty of sight. That is, it is part of our concept of an eye that it is for seeing. And similarly for hands and legs and all the other commonly named, external parts of the body. The functions of the objects to which these names refer constitute an important part of the meaning of those names. This can also be seen from the way in which we use the function of a human body part to guide us in applying terms analogically to animals with very different kinds of bodies, to machines, and so forth. Thus, we speak of a horseshoe crab's "eyes," a robot's "hands," a spider's "legs," etc. It is the function, not the morphology, of the human body part that determines how the name of the part gets applied analogically.

Another thing that the above reflections show is the following. Since the request for identification of a body part by ostension seems to elicit functional definitions just as readily as morphological ones, the morphological and the functional properties of body

parts would seem to be on an epistemic par. If that is right, then an eye's function of seeing is a property that is just as intuitively obvious as its property of being a sphere embedded in the front of the skull. A hand's function of grasping is a property that is just as intuitively obvious as its property of having a span of, say, eight inches, or its property of having five fingers. And so on. Since the epistemic basis for the function ascription is similar to the epistemic basis for the ascription of the morphological properties, if we are going to be realists about the latter, it seems that we ought to be realists about the former. In short, at least insofar as everyday life is concerned, there seems to be little reason to deny an objective reality to functional properties that we instinctively grant to extensive properties like location, size, shape, and number. This sort of parity of reasoning forms the basis for what one might call the 'biological-function realism of everyday life."

It might be objected that we seldom use the word "function" in everyday speech, and that it is in fact a strictly scientific concept. Therefore, the so-called "biological-function realism of everyday life" does not really extend to the concept of biological function at all, or at least cannot be assumed to do so without further argument. Now, it is true, I think, that in everyday life the word "function" is mostly reserved for cases in which we need to ask for clarification about the purpose of something, most commonly a manmade artifact of some sort. For example, if two persons came across an object of unidentifiable shape and material while walking in the woods, one might ask the other "What do you think it's for?," and the other might reply: "Do you think it even has a function?" I submit that this readily imaginable bit of dialogue shows that we have a concept of a function in everyday life that is quite close to that of the biologist, even if we

seldom use the biologist's favored word to express it.[15] This can also be demonstrated by translating part of the preceding series of questions and answers into a slightly different form, explicitly using the word "function."

Q. (Pointing to an eye): What is this?
A. It's an eye.
Q. What's the function of an eye?
A. The function of an eye is to see with.

Q. (Holding up a hand) What is this?
A. It's a hand.
Q. What's the function of a hand?
A. The function of a hand is to grasp with.

. . . and so on . . .

While there is no doubt that it is more natural in ordinary conversation to ask of something "what's it for?" than to ask "what's its function," nevertheless, I contend that the two forms are for all practical purposes interchangeable, and that the casting of the series of questions and answers into the form explicitly referencing function, while admittedly awkward, would nevertheless be perfectly intelligible to any English speaker. Therefore, it seems that there really is an instinctive "biological-function realism of everyday life."

Still, it is perhaps not obvious than any of this settles the question whether or not there is a significant difference between the biological-function attributions that we routinely encounter in everyday life and those to be found a highly technical, biological and other scientific discourses. How can we be sure that the biological-function realism of everyday life will carry over to biological function as it is understood by biological

[15] Actually, the word "function" is more highly favored by philosophers than it is by working biologists. I will address this point in a moment.

science? As it happens, it is a noteworthy and perhaps surprising fact that everyday speech and sophisticated scientific discourse are in striking accord on this point. Indeed, the one shades imperceptibly into the other. Consider, for example, the following sequence of function ascriptions.

> The function of the heart is to circulate the blood.[16]
>
> One of the functions of the circulation of the blood is to carry necessary metabolites, including oxygen, to all the cells of the body.[17]
>
> One of the main functions of oxygen in cells is to participate in oxidative phosphorylation in the mitochondria.
>
> One of the main functions of oxidative phosphorylation in the mitochondria is ATP synthesis.
>
> One of the main functions of ATP is to undergo conversion to ADP by means of hydrolytic dephosphorylation, an exergonic process.
>
> One of the main functions of the conversion of ATP to ADP is for the process to be thermodynamically coupled to endergonic metabolic processes throughout the body.

We started out with a statement of function that almost any layperson would spontaneously make. We ended up with a statement of function that few besides a trained biologist or physician would be capable of making. Perhaps, the layperson would get lost rather quickly as we go down the list. However, most educated people would, I think, be capable of naming the function of the circulation as the transport of metabolites, including oxygen. And many would have a vague idea that the function of oxygen is to

[16] This statement may conceal difficulties relating to fineness of discrimination of functions. For example, one might claim that the primary function of the heart is to beat, and that beating is instrumental to the further function of circulating the blood. But this does not mean that the notion of function is either impossibly vague or subjective; it just means that careful, empirical discriminations must be made. But that is hardly surprising; after all, no one, I think, wishes to claim that functions are knowable a priori. The crucial point is just that correct ascription of biological function is a matter for empirical investigation.

[17] Complicating matters further is the fact that multiple functions are the norm in living things. Nevertheless, discriminations can be made, and, as already noted, making such discriminations is one of the main preoccupations of experimental biology.

participate, somehow or other, in metabolism. By the time we arrive at the conversion of ATP to ADP and thermodynamic coupling, of course, we have clearly passed from the sort of function attribution that we might expect from the man in the street to the sort that requires a fair amount of technical expertise. But note that that expert knowledge has mainly to do with what the microscopic bearers of function in living things are, how such function bearers are interrelated, which physical principles underlie them, and so forth. Such expertise does not appear to involve the concept of function as such, or the way in which functions in general are identified. At no point do we get the sense that we have passed from talking about functions in one sort of way to talking about them in another, qualitatively different sort of way.

In short, these examples seem to show that the layperson and the highly trained scientist employ the concept of function in much the same way. But if that is so, then the kind of evidence that is adduced in biological function attributions in everyday life will be the same kind of evidence that is adduced in biological function attributions tout court. In other words, the claim that the function of the heart is to circulate the blood and the claim that one of the functions of ATP is to undergo conversion to ADP are on an epistemological par qua function attributions.[18] Since the kinds of justifications that one might adduce to support a function attribution would seem to be very similar, if not identical, in both cases, it appears that there is no reason specific to the nature of functions for skepticism to be greater in the one case than in the other. That is to say, there is no special, or extra, reason for skepticism about biological function in biological

[18] This is not, of course, to deny that the specific evidence for the former attribution is far simpler and more understandable to the man in the street than the specific evidence for the latter; it is just to say that the evidence for function attribution as such operates in the same way in each case.

science, beyond whatever reasons there may be for skepticism about biological function tout court.[19]

Several objections to this line of reasoning come to mind. For one thing, as was noted above, the heart's function of circulating the blood is already a fact that is learned in school. Therefore, it might be thought that it is not on an epistemological par with intuitively obvious function ascriptions like the hand's function of holding. In other words, it might be denied that there is any such thing as an everyday use of the concept of biological function—apart from a few primitive cases like the eye, the hand, the legs, and such—and hence no basis for a biological-function realism of everyday life, after all, or at least, not one that can be extrapolated to scientific contexts.

There is of course no doubt that the nature of the functions of the various internal organs of the body was a matter of speculation until modern times—indeed, in some cases, until today.[20] But, of course, it is one thing to know that something must have a function; it is something else to know what that function is. Aristotle may have been wrong in thinking that the heart was the seat of sentience, but he was not wrong in believing that the heart had a function. What modern science has added to our knowledge since the time of Aristotle is immensely greater insight into the true nature of the functions of the component parts of living things. But the idea of function itself is no insight of modern science. The very fact that the correct function of the heart is now

[19] Needless to say, the case against wholesale skepticism about biological function, encompassing both biological science and everyday life, still remains to be made out, but that is precisely the aim of this dissertation as a whole.

[20] In humans, for example, the function of the thymus gland in protecting against autoimmunity was only discovered in the 1960s. Speculation is still ongoing concerning the possible function of the appendix.

known to every schoolchild shows how intuitively obvious the notion of function really is. It might even be argued that cases like the heart's function of circulating the blood that are the result of historical discovery, but are now known to all, actually form an intermediate link between everyday functions like those of the eyes, the hands, and the legs, and scientific functions like those of the conversion of ATP to ADP and of thermodynamic coupling. From this perspective, the case of the heart provides additional evidence for an everyday concept of function, and for a biological-function realism of everyday life based upon it.

Another objection that might be raised to the general line of argument above is that the term "function" does not appear very frequently in the most sophisticated descriptions of biological phenomena, such as those found in textbooks. This fact might then be taken as a reason to deny the relevance of the biological-function realism of everyday life to the question of the ontological status of the phenomena under investigation by biological science.

It is true, for the most part, that textbooks do not make much explicit use of the term "function" in their detailed physical descriptions of the myriad structures and processes that comprise living things. However, there are at least two other facts which offset this one, and which show that the concept of function implicitly permeates all of biological science, at whatever level of sophistication it is conducted.

The first of these countervailing facts is that the discovery of a novel biological structure or process immediately initiates a hunt for its function. No characterization of a biological entity is deemed complete, no matter how exhaustive it may be in physico-

chemical terms, so long as the entity's function remains unknown. For example, there is a nuclear protein in humans and other mammals named Ki-67. Ki-67 is expressed during cell proliferation, and so is widely used as a marker in cancer research and clinical work. However, though this protein was discovered in 1983, and has long been structurally well characterized, its function was entirely unknown until recently, and remains under active investigation. Scholzen & Gerdes (2000) provide an interesting account of the earlier efforts to solve this riddle, which makes palpable the researchers' sense of frustration. It is now thought that Ki-67 may be implicated in the synthesis of ribosomal RNA outside the nucleus, but the details remain elusive (Bullwinkel et al., 2006). This example clearly shows that in practice biologists treat function on an epistemological par with physico-chemical structure.

The second countervailing fact is that, if the term "function" is scarce in biology textbooks, other terms with a distinctly teleological or intentional character are not. In fact, teleological language in a profusion of forms is ubiquitous in biology, as may be easily verified by referring to almost any page of any textbook.[21] Since my main concern in this dissertation is not with the use of the term "function" per se, but rather with its teleological character, and with the natural phenomenon of teleology more broadly, the fact that textbooks happen to eschew that particular term is of little consequence for the argument limned above.

[21] For example, opening Alberts et al. (2007) at random, I find on p. 415 the following statement: "Thus a cell can control the proteins it makes by (1) controlling when and how often a given gene is transcribed (transcription control); . . ." (Five other types of control systems are listed.) "Control" implies the existence of a distinguishable preferred state of a system, the maintenance of which is the purpose of the control systems mentioned. "Control" is thus both a normative and a teleological concept.

Finally, it is sometimes claimed that the term "function," in the teleological sense, may be dispensed with—in the field of evolutionary biology, at least—in favor of the term "adaptation." However, as Amundson and Lauder (1994; p. 447) point out, "[a] trait is an adaptation when and only when it has a function." So, merely translating "function" talk into "adaptation" talk does not ipso facto demonstrate the dispensability of teleological function in biological science. To do that, it would be necessary to show that "adaptations" are themselves wholly analyzable in non-teleological terms. For some, that will sound like an analytical truth, but in fact it is a theoretical claim—one that I will contest in Chapter 3, below.

1.5 A Note on Realism

The preceding discussion of the phenomenology of biological function and function-discourse has done three things: (1) It has provided us with a paradigm of biological function (bacterial chemotaxis) that, thanks to its generality, provides us with a concrete basis for discussing the teleological character of all biological functions as such; (2) It has established the "biological-function realism of everyday life"; and (3) It has shown that there is substantial warrant for maintaining that teleological ascriptions in everyday life and those in biological science share precisely the same phenomenological basis. This last demonstration, in turn, already provides us with substantial warrant for taking the objective reality of teleology in biology as at least a serious possibility.

But, of course, while ordinary usage may provide prima facie warrant for our beliefs about the furniture of the world, common sense alone cannot be determinative of our ontology. There is little question but that the outstanding cognitive success of the

natural sciences has earned them the right to play an important role in any effort that we make to construct a systematic view of the world. But ought not scientific theorizing, then, be determinative of our ontology? And if so, then how can we reconcile our teleology-impregnated commonsense view of ourselves and other living things with biological science which officially bans teleology from its theoretical picture of the world? This question is the main focus of the remaining chapters of this dissertation, but it would not be amiss to make some observations of a general and preliminary nature about what exactly is meant by "realism" with respect to teleology in biology, and by the concept of "teleology" itself, to round out our introductory discussion.

First, then, let me say a few words about what is meant by "realism" in connection with teleology in biology. As mentioned briefly above, for the purposes of this dissertation I am simply assuming naturalism and a broad version of general metaphysical realism. Therefore, I am excluding more or less without argument several ways in which the so different commonsense and scientific views of teleology might be reconciled. One of these would be by means of positing a radical ontological separation between the human and the physical realms. I am simply going to pass over here all such frankly dualistic attempts at a solution to the problem of teleology. Another would be by means of a purely epistemological reconciliation of our commonsense and scientific views, which regards teleology as an illusory though perhaps unavoidable projection of the human mind. I reject this sort of approach without comment, as well. On the general metaphysical-realist view I am adopting here, no dualistic, idealistic, or subjectivist view of teleology could constitute a real reconciliation of our commonsense and scientific worldviews. Rather, such a view of the nature of teleology would inevitably constitute

recognition of the unreliability of our commonsense view. In other words, if one accepts a general metaphysical realism, as I do, then accepting anything less than a robust realism with respect to teleology would be tantamount to accepting the determinative nature of the present scientific worldview with respect to ontology. And that would simply mean accepting the unreality of teleology, in which case, the supposed "reconciliation" would consist of nothing more than a total capitulation of one side. Of course, in the end it may not be possible to justify regarding teleology realistically. Capitulation to the scientific reductionist view may turn out to be unavoidable, after all. In other words, **TRB** may of course fail to convince, and so fail. But from a metaphysical-realist perspective, that is a conclusion ought to be arrived at only after a long struggle, whereas to accept a dualistic or idealistic viewpoint would be simply to give up without a fight.

In short, the sort of naturalism and metaphysical realism I am espousing here assume the reality of the conflict between the commonsense and the scientific views of teleology. Reconciliation of these views, then, assumes a real adjustment in our view of what exists, either on the side of common sense, or on the side of biological science, or both. The heart of this dissertation lies in negotiating the necessary adjustments to arrive at the view that seems, on balance, best to do justice to everything we know about life. This is a many-faceted undertaking. Many factors of common sense and of science must be weighed against each other. It is also a significant undertaking that must result in a substantial readjustment of our view of the world, whether on the side of common sense, or on the scientific side, or both.

The principal substantive claim that has been advanced so far—that teleology has a firm phenomenological foundation in both common sense and biological science—

helps to ease the burden of proof from off the shoulders of common sense alone and to redistribute it more fairly between common sense and science. This means that it is not known a priori how the necessary conceptual adjustments must be made. It may yet turn out that our commonsense view must be rejected, or it may be that the prohibition on teleology in natural science must be lifted. Or we may find that conceptual adjustments are required on both sides. All that I have endeavored to show so far is that we are entitled to embark upon this inquiry with an open mind, without the presumption that teleology is guilty until proven innocent. There are powerful reasons for us to take teleology seriously—not the least of which is that we cannot help but do so in our practical engagement with the living world, whatever we may say in our theoretical pronouncements about that world.

> To sum up, by "realism" about teleology in biology, I mean the following:
>
> Definition of Teleological Realism in Biology (**TRB**) To be a realist about teleology in biology is to hold that there is some objectively real principle in living matter itself that corresponds to our concept of teleology.

All I am aiming at in this definition is the basic idea common to all forms of realism—namely, that a given concept corresponds to something real, something objectively existing outside any human mind, something that is the case and would have been the case even if no human being had ever existed.

Realism, of course, raises a host of deep and difficult questions, from worries about the external world in general, to misgivings about the postulates of scientific theories in particular. So, it might be thought that I ought to begin by offering some general justification of realism, before addressing the particular postulate of **TRB**. By way of addressing this worry, I could mention that realism is the default position of both

common sense and natural science—which fact alone might be supposed sufficient warrant for its supposition here as a general metaphysical framework. But of course such a response is somewhat cavalier, and might be indefensible were it not for the following consideration.

TRB postulates realism within a very restricted domain—teleology as applied to living systems. For the purpose of examining the meaning and merits of this postulate, debates about realism in its various wider senses are irrelevant. That is because to succeed in my purpose here all that is necessary is for me to establish that the referent of the concept "teleology" as employed in biology has the same ontological status as that of any other empirical concept employed in everyday life or in the conduct of natural science. To do this, I am required only to show that the notion of teleology is (or can be made) coherent and that belief in its objective existence in living systems enjoys the same degree of empirical warrant as belief in the objective existence of the referents of other commonsense and theoretical concepts in biology. That is, I must show that teleology enjoys epistemological parity with such concepts as "matter," "energy," "atom," "molecule," "cell," "nucleus," "membrane," "flagellum," and so forth. From this, ontological parity may be inferred, on the principle that where no difference in epistemological status exists, no difference in ontological status ought to be posited.

In other words, **TRB** is essentially the claim that belief in teleology is as well warranted as belief in such scientific postulates as cells and flagella. Therefore, as long as teleology in biology can be shown to be no less mind-independent (or no more mind-dependent) than cells and flagella, **TRB** will be vindicated. Since ontological parity between teleology and other empirically ascertainable biological phenomena is all that is

being claimed here, it makes no difference to my project whether scientific realism is accepted or rejected. Even if one wished to understand cells and flagella in an instrumentalist or even an idealist sense, **TRB** would still succeed so long as teleology was granted a similar ontological status.

1.6 Some Preliminary Remarks on the Concept of Teleology

What the postulate of **TRB** amounts to in detail, and whether it is ultimately rationally warranted, will of course depend in large measure upon the interpretation given to the term "teleology." Indeed, teleology—both the everyday and the biological concepts, or cluster of concepts, as well as the biological phenomena to which those concepts refer—is the central focus of this dissertation. Let us now turn to a preliminary discussion of what the term will be taken herein to mean.

I begin by noting that the primary difficulty with teleological concepts in biology derives from the following triad of considerations:

(1) The paradigm case of teleological ascription is to human intentions.[22]

Q: Why are you running?
A: I am running to catch the bus (= for the purpose of catching the bus = I have the intention to catch the bus).

(2) Teleological ascriptions to all living systems whatsoever seem clearly coherent.

Q: Why is the bacterium "running" (swimming rectilinearly)?
A: To find and ingest food (= for the purpose of feeding).

[22] For present purposes, an "intention" will be understood as a purpose or goal that is held consciously, i.e., "before the mind." "Intentionality," then, will denote either the capacity of a particular organism to have such intentions or else the ensemble of such capacities considered as a natural kind. The context should always make clear which sense is intended.

(3) Intentional ascriptions to some living systems seem clearly incoherent

Q: Why is the bacterium "running"?
A: Apparently not because it has the intention of feeding (bacteria do not seem to be the sort of thing that can have intentions).

There seem to be only two ways to go here. If teleological ascriptions are to be identified with intentional ascriptions, then teleological ascriptions to some kinds of living systems (such as bacteria, protozoa, plants, and the like) must be disallowed. This move leaves us with the problem of explaining the apparent coherence of the disallowed teleological ascriptions. On the other hand, if teleological ascriptions to all living systems whatsoever are to be accounted coherent (as they appear to be), then we must not identify teleological ascriptions with intentional ascriptions. This move leaves us with the task of providing an alternative account of teleology, one that makes no mention of intentionality. The latter strategy is the one that will be pursued in this dissertation.

Theoretically, of course, another way out is to allow intentional ascriptions to any living systems whatsoever. Along with practically everyone else, I reject this approach. However, it is worth pausing for a moment to reflect upon the reason why our minds seem to revolt against the suggestion that bacteria swim up chemical gradients because they intend to feed. I will do this by way of making some distinctions that I hope will clarify the precise nature of the present project.

Crudely put, we balk at ascribing intentions to bacteria because an intention seems to be something which requires a mind, and it seems fantastic to suppose that bacteria have minds. It seems fantastic to suppose that bacteria have minds because a mind seems to be something that requires a brain (or at least nervous tissue), and we know for certain that bacteria do not have any nervous tissue. So, whatever one may

think of ascribing intentions to nonhuman animals with large brains, like dogs and cats, it seems entirely justifiable to withhold ascriptions of intentionality from small animals without brains, like bacteria.

This reasoning is sound so far as it goes, but it does not go far enough. It ignores the fact that the concept of "mind" is ambiguous, which means that the above argument is equivocal. Certainly, bacteria lack "minds" in the sense of being able to hold purposes consciously "before the mind," for that capacity requires not only a brain but also arguably human-like language capacity. But perhaps there is another, more restricted sense in which bacteria may be said to have "minds," and if so, then this sense of "mind" might suffice to justify ascribing teleological (and more generally, normative) predicates to bacteria and similar creatures, even while refusing to ascribe to them intentional predicates.

There are two points here, one semi-empirical, the other purely conceptual. First, we must ask: How can it make sense to say of a bacterium, or any other lower life-form, that it has a "mind," in however attenuated a sense, if it does not have any brain whatsoever or any capacity for intentionality in the sense of holding purposes "before the mind"? The answer is: It makes sense to do so because as a matter of fact all organisms, including bacteria, behave in a manner that we cannot help but interpret as purposeful. That is the empirical, or partly empirical, point: Bacteria and other single-celled creatures simply behave as if they were acting purposefully, whether we are willing to attribute intentionality to them or not.[23] The conceptual point is this: Our linguistic willingness to

[23] For many stunning examples, see Jennings (1962); for a more recent discussion, see Shapiro (2007).

ascribe purposeful behavior to bacteria (whatever our metaphysical queasiness) cannot be separated conceptually from our willingness to ascribe functions to them, for the two concepts are internally connected. Part of the meaning of "purposeful behavior" is having an end (however that "having" and that "end" are ultimately to be interpreted), and part of the meaning of "function" is supplying the means to an end. Thus, purposeful behavior and function are complementary concepts. That is so because ends and means are equally complementary concepts. That is, the concept of an end implies the existence of an appropriate function, or set of means. In other words, the concept of an end implies the possibility of its own attainment under some possible conditions, though of course not necessarily under actually existing conditions.[24] An end could never exist in the first place absent the recognition and successful employment of appropriate means necessary to bring it about. Simply put, nothing would count as an end, unless it were the sort of thing that required means to bring it about. It is even clearer that the concept of means implies the concept of end. So the concept of end and the concept of means are complementary. Thus, it seems quite clear that the concept of purposeful behavior and the concept of function are closely connected conceptually, for they are both bound up with the notion of a means-end relationship.

Of course, this fact might well be interpreted as further grounds for restricting teleological ascriptions to cases in which intentionality is clearly present. At least, that would be a plausible inference if purposive behavior could be plausibly ascribed only to

[24] This is true even for intentional ends, where the successful thinking of the object in question counts as attainment of the end. It is not necessary that an intentional end be possibly attainable though action for it to be an end in the sense of being an object of thought. But for something to be a possible object of thought it is necessary that some intentional agent be capable of thinking of it successfully in some circumstances. Thus, the principle that possible attainability is criterial of being an end applies to intentional ends, as well.

ourselves. If the one clear case we have of intentionality, namely, ourselves—that is, the one clear case in which we know beyond cavil that purposes are being held consciously before the mind—if that case were also plausibly the paradigm case of purposeful behavior, then the preceding argument, far from undermining the case for identifying teleology with intentionality, would actually strengthen it.

If teleology is not to be identified with intentionality, then what is the correct way to conceive of the relationship between these two concepts, and the natural phenomena to which they refer? To answer this question, we must consider two facts. First, consider the fact that our concept of teleology is broader than our concept of intentionality: The notion of teleology encompasses ideas that are lacking in that of intentionality. Von Wright (1971) puts this point as follows:

> One could divide the domain traditionally claimed for teleology into two subprovinces. One is the domain of the notions of function, purpose (fullness [sic]) and "organic wholes" ("systems"). The other is that of aiming and intentionality. Function and purpose figure prominently in the biological sciences, intentionality in the behavioral sciences, social study, and historiography. But the domains of biology and of behavioral science largely overlap and so, of course, do the domains of function, purpose, and wholes on the one hand and that of aiming and intentionality on the other. To distinguish them may nevertheless be useful. (ibid.; p. 16)

So, the first fact to keep in mind in reflecting upon the relationship between teleology and intentionality is that they are simply two different concepts; while their semantic ranges do overlap, they are by no means congruent. There is more to our concept of teleology than is contained in our concept of intentionality. If, as a matter of linguistic fact, the concept of teleology finds application to areas of reality beyond the bounds of applicability of the concept of intentionality, does this not give us at least prima facie

grounds for believing that something in the world corresponds to the concept of teleology that does not also correspond to the concept of intentionality?

Second, given that the concepts of intentionality and teleology form partially overlapping but not congruent sets, the question naturally arises whether they are merely intersecting sets, or whether intentionality is not a proper subset of teleology. It would seem that the latter must be the case. The reason is simple. The essential feature of intentionality is "aboutness," that is, the directedness of mental states like beliefs, desires, fears, etc. toward particular objects in the world. This directedness is clearly teleological in character (which is of course the reason why teleology has been traditionally identified with intentionality in the first place). No mental state that lacked this property of directedness would qualify as "intentional." We have already established that not every instance of teleology counts as an instance of intentionality. Since we now see that every instance of intentionality does, on the other hand, count as an instance of teleology, it is obvious that intentionality must be a proper subset of teleology, or, in other words, intentionality is a species of a broader genus, teleology. As Okrent (2007; p. xi) has put the point: "Intentionality is rooted in teleology . . ."

This way of viewing the relationship between teleology and intentionality is not only forced upon us by reflection upon the nature of the respective concepts, it is also a commonsense conclusion that many will wish to draw from our scientific understanding of the place of the human mind in nature. True, it is a delicate question whether any but human minds possess intentionality, properly speaking. But that is fortunately not a question upon which I need take a stand here. All that is required here is to point out that human minds, with all their properties and powers, including the power of intending and

referring, must be considered as having evolved from lower minds lacking this power. Since we spontaneously attribute teleology, but not intentionality, to all living things, it is obvious that a theoretical perspective that views intentionality as a proper subset of teleology, and thus as a relatively late acquisition evolutionarily speaking, is in better accord with our overall view of ourselves as natural beings than a perspective which insists on equating the two properties. As I am adopting a rigorously naturalistic perspective in this dissertation, this consideration must count rather powerfully in favor of the pan-biological attitude toward teleology adopted herein, as well.

Even so, I must admit that nothing I have said so far is conclusive, or even has very much tendency to block what was no doubt the initial response of many critical readers: All this ascription of "purposeful behavior" to organisms like bacteria is just metaphorical. No one takes it seriously (least of all in the laboratory). So, why should the mere fact that we all (biologists included) speak in this way warrant our ascribing real capacities to the lowly living systems in question? No brain, no intentionality; no intentionality, no purpose; end of story.

To put the point in somewhat more theoretical terms, the skeptical reader might well wish to point out the following: It is one thing to give reasons to doubt the identification of teleology with intentionality, and thus the presumption against the objective existence of teleology in biology, and it is something quite different to provide positive reasons for believing in the objective existence of teleology in biology as a scientific phenomenon. It will no doubt seem to many readers that no amount of reflection upon the meaning of our theoretical concepts will carry much conviction in the

absence of a positive account of teleology in biology that comports rationally with everything else that we believe about life and living things.

Ultimately, no purely philosophical response to this objection is possible; only a well-substantiated scientific theory of teleology will carry true conviction, at the end of the day. Accordingly, I will canvass some potential candidates for such a theory in Chapter 4, below. For the present, however, I would like to respond to this objection in the following way.

Any outright rejection of our natural way of talking about living things—including microorganisms—requires an alternative story about how things are with them that makes no use of teleological or normative concepts. Of course, those who wish to reject the natural way of speaking about lowly life-forms do believe that they are in possession of such an alternative account of how things are with them. I will attempt to show, in Chapter 3, that they are mistaken in their belief. It would be pointless to rehearse in unconvincing summary here the arguments to be expounded at length there.

Given, then, that there is no coherent reductive account of teleology, all that is required to shift the burden of proof is for me to provide prima facie warrant for looking upon our patently teleological vocabulary with metaphysical seriousness. For remember that the argument against objective teleology ran like this:

1. Teleology is necessarily bound up with intentionality.
2. Lower life-forms do not have a capacity for intentionality.
3. Therefore, lower life-forms do not possess objective teleology.

Without the presumption of a necessary connection between teleology and intentionality, the argument fails, and there are no longer grounds for automatically assuming that ascribing objective teleology to lower organisms is preposterous.

For these reasons, it seems to me that it is no longer appropriate simply to dismiss the notion of objectively existing teleology out of hand. Lyon (2006) calls the traditional method of starting from human intentionality and seeing how far down the ladder of nature cognition, teleology, and normativity extend the "anthropogenic" approach, while she terms the contrary effort to understand teleology (and agency) as a universal biological attribute, out of which human intentionality has ultimately grown, the "biogenic" approach. In a magisterial survey of the literature, she lays out both approaches in a fair way, while ultimately opting for the biogenic approach in these terms:

> what the biogenic approach seems to suggest is that agency, the capacity to adaptively change (act, interact) relative to an environment, and normative assessment, the capacity to opt to do this rather than that in the present circumstances, given present needs, are the dual bases of cognition . . . The capacity to infer relations between external circumstances and internal need to facilitate agency may be, to a first approximation, what cognition is. (Lyon, 2006; p. 27)

So far, I have described the behavior of the lowest life-forms only in terms of function and behavior, which enjoy an immediate, phenomenological warrant. In the cited passage, Lyon is less circumspect, introducing the notions of agency, normative assessment, and cognition. I have not yet said anything to justify these more ambitious claims, though I will do so very shortly, in Chapter 2. For now, I would simply like to point out that, if a case can indeed be made out for the proper ascription of such concepts to bacteria, and thus to organisms as such, then Lyon's "biogenic" approach—that is, the

project of explaining human intentionality and rational agency as a development of teleological and normative capacities already present in the lowest life-forms—is perfectly coherent. To charge the "biogenic" approach with "anthropomorphism" in such a way as to dismiss it out of hand is tantamount to assuming that only the anthropogenic approach is warranted. In short, the charge of "anthropomorphism" in this context is simply question-begging.

Even with the general character of my basic claim clarified in this way, there must still remain quite a bit of uncertainty in the reader regarding the specific nature of the sort of objective teleology so envisioned. Let me now turn to the task of spelling out in somewhat greater detail precisely what I mean by the term "teleology" (though the bulk of the discussion of this concept will take place in Chapter 2). To fix ideas, let me begin with the following stipulation.

> Preliminary Definition of Teleological Function in Biology A biological process or behavior is a teleological function if and only if it meets both of the following conditions: [25]
>
> (1) Determinative Condition: The final state ("end") of the process partly determines (produces, brings about) its corresponding initial states ("means");
>
> (2) Normative Condition: The end is "normative," in the sense that it establishes criteria of appropriateness with respect to the means (namely, the initial states of a teleological process or behavior are "appropriate" insofar as they tend to produce or bring about their corresponding final state, and are "inappropriate" insofar as they fail to do so).

[25] If it is objected that this definition wrongly excludes structures from the domain of the teleological, two observations may be made in reply: (1) structures are only teleological insofar as they perform or participate in functions (e.g., the teleological character of the heart derives from the function it performs, namely, circulating the blood); and (2) in any case, most if not all "structures" in biology are in fact steady-state processes (e.g., bone turnover, or remodeling, is continuous throughout life).

On this understanding of terminology, then, "teleology in biology" will refer in a general and collective way to all functions that meet the determinative and normative conditions. It follows that no biological process or behavior that fails to meet these two conditions will count as "teleological."

"Teleology" in this sense is to be sharply distinguished from "teleonomy," a term that was introduced into the literature by Pittendrigh (1958) as a working substitute for "teleology," and which was supposed to be free from the latter word's ambiguity and undesirable historical baggage and metaphysical connotations. The concept of teleonomy is not itself of crystalline clarity, and has been subject to various interpretations, but the following definition of Mayr's has been one of the most influential (1988; p. 45): "A teleonomic process or behavior is one which owes its goal-directedness to the operation of a program."

Whether teleonomy in this sense is a coherent concept will be investigated in detail in Chapter 3, below. However, it is immediately clear that on Mayr's view a "teleonomic" process, though "goal-directed" in a certain sense, does not meet the conditions given above for a biological process to count as "teleological." First, regarding the determinative condition, Mayr is quite clear (ibid.; p. 48): ". . . a program is (1) something material, and (2) it exists prior to the initiation of the teleonomic process. Hence, it is consistent with a causal explanation" (emphasis added).

What Mayr means by being "consistent with a causal explanation" is precisely that the final state (the developmental result) is entirely determined by the initial state (the "program"). As there is no room on this view for the initial state to receive any

determinative influence whatsoever from the final state, the determinative condition is not met, and "teleonomy," at least on Mayr's conception, is something quite distinct from "teleology" as defined above. Mayr does not discuss the normative condition, but it would seem to follow from his conception of teleonomy as a material process consistent with purely causal explanation that the normative condition would not be met, either.

Teleology and teleonomy, then, are antithetical and competing accounts of the nature of biological phenomena. The main point of introducing the concept of teleonomy was to "save the phenomena" of biology, which do give the appearance of teleology, without resorting to such metaphysically problematic ideas as the determinative condition and the normative condition. Teleonomy, then, may be viewed as essentially a theoretical framework for "reducing" phenomena that were traditionally understood as teleological to a congeries of purely material and "efficient" (push/pull) causal interactions.

As mentioned several times already, the main theoretical frameworks for carrying out such a reductive project are the analyses of functions as causal contributions-to-a-system and as selected-effects. We will see in Chapter 3 why it can be legitimately doubted whether either of these reductive schemes works, and therefore why it may also be doubted that the concept of teleonomy adequately describes the phenomenology of function in biology.

With these clarifications of what is meant by "teleological realism in biology" out of the way then, let us turn to the next task, which is that of showing that the concept of teleology implies agency, and that ascription of agency to organisms as such is in fact warranted.

CHAPTER 2:

WHAT IS THE PROPER SCOPE OF OUR CONCEPT OF NORMATIVE AGENCY?

2.1 Introduction

In the first chapter, I have already made some preliminary distinctions regarding the notion of teleology, which is perhaps the clearest case for my purposes among the elementary normative concepts,[26] in that hardly anyone would dispute that teleological notions are properly ascribable to organisms as such. In this chapter, I will be enlarging on some of the ideas already broached there, in order to show that teleology is best understood in relation to the notion of agency, in the normative sense, and that, to the extent that we are justified in conceiving of teleology as an objectively existing property of organisms as such, we are also justified in conceiving of agency as an objective property or capacity of all living systems. Needless to say, this is a counterintuitive claim that will require arguments of considerable force to motivate its acceptance. Indeed, some may find the claim so difficult to accept as to regard it as a sufficient refutation of any arguments I bring to bear. Anticipating this sort of reaction, I will proceed in a twofold manner. For the most part, I will be advancing a series of mainly conceptual arguments,

[26] It will be recalled that the "elementary normative concepts" were defined in Chapter 1 as normative concepts connected to prudential instrumental action generally, exclusive of the normative concepts that imply human rational deliberation. Examples of the elementary normative concepts would be, e.g., purpose, value, well-being, need, and the "instrumental ought" (normative requirement in the narrow sense—see below). Examples of normative concepts excluded as non-elementary would be, e.g., moral right, the moral virtues (courage, justice, honesty, beneficence, etc.), beauty, and other concepts whose scope of application is clearly restricted to humans qua reasoning, language-using, social, and cultural beings.

which will constitute the main substance of this chapter. However, in addition to these, I will also be bringing a number of preliminary empirical considerations to bear in direct support of my main thesis, to help diffuse its air of paradox. This latter approach will be greatly expanded upon afterwards, in Chapter 4.

As we have already seen in the first chapter, teleological language is universally used to describe the behavior of even the simplest organisms, both in everyday life and in biological science. This much is beyond dispute. However, it could, of course, be denied that such language ought to be taken at face value. It might, for instance, be argued that we know that teleology has no objective existence, because we already have in place theoretical frameworks—such as the discipline of molecular biology and the theory of natural selection—by means of which the apparent teleology in biological systems can be fully reduced to mechanism. In Chapter 3, I will be showing why this move will not work (in a nutshell, because both theoretical frameworks tacitly presuppose either teleology or normativity at some point in their explanatory structures). For this reason, I will not attempt to defend against objections of this sort in this chapter.

This chapter will be structured as follows. First, in the following section, I will discuss some distinctions relating to the concept of mind, in order to help clarify the precise nature of the claim I will be advancing. I will propose a working division of mind-like phenomena in biology into the three categories of sapience (rationality, possessed by human beings alone), sentience (feeling or subjective experience, which many animals, but probably not all organisms, possess), and appetence (the basic purposive behaviors and metabolic functions that constitute a physical system as a living thing). This will help to clarify the claim that organisms as such are normative agents, by

rephrasing it as the claim that appetence is a form of normative agency. Then, in Section 2.3, I will make some further important distinctions relating to the concepts of normativity and agency. Notably, I will distinguish two separate problems related to the claim that organisms are normative agents—namely, the Scope Problem and the Ground Problem. The Scope Problem will be the principal subject of the remainder of the present chapter, while the Ground Problem will be investigated in Chapter 4. Next, in Sections 2.4 through 2.6, I will present what I shall call the Scope Argument, leading to the Principal Claim of this chapter that the proper scope of our concept of normative agency is living systems as such. Finally, in Section 2.7, I will begin to make the transition from conceptual analysis of the elementary normative concepts and the notion of agency, to a consideration of empirical considerations bearing on my Principal Claim (this section will serve as an introductory bridge to the material in Chapter 4).

2.2 A Typology of Mind-like Phenomena

The basic issue that I am addressing here is the ontological status of teleology. I will be arguing that teleology is conceptually connected a broader range of normative concepts, which may be viewed as so many different aspects of the concept of action. But action, in our ordinary way of thinking and speaking, is closely connected with the concept of mind. Thus, it might seem that to claim that teleology is an objective property of organisms as such is to attribute mind to organisms as such. And given that many if not most philosophers will understand by "mind" the capacity for having conscious thoughts, and that it seems absurd to attribute conscious thoughts to very simple

organisms such as bacteria, it may appear that my argument must either founder in equivocation or else lead directly and by a very short path to paradox.

However, matters are not that simple. We must tread cautiously here in order to avoid begging the question, in either direction. What is needed is a typology that is apt for naming the phenomena under study, but which is also free from connotations that appear to presuppose either the truth or the falsity of the claim that teleology possesses objective existence. That requirement is not so easy to fulfill as it might appear. The reason is that our everyday vocabulary seems to assume a dichotomy between "mind," where conscious thought is fully present, on the one hand, and "mechanism," where conscious thought is wholly absent, on the other. But the goal of our inquiry is precisely to inquire whether there may not be some intermediate ground between mind and mechanism, so understood, and if so, what its nature might be. Whatever vocabulary is adopted threatens to undermine that goal, either by assuming mind to be fully present in the simplest organisms, which would be absurd, or by assuming it to be wholly absent in any sense whatever, which would beg the question against the Principal Claim of this chapter.

It seems to me that the way forward is to acknowledge that there exists something like a consensus—if not within the philosophical community, at least within the scientific community—that "mind" is not a univocal term, but something more closely approaching an umbrella term, and that there are in fact many different forms of mind in nature. Or, at least, it is certain that there now exists an extensive empirical literature that employs just such a generic concept of mind, and that endeavors to inquire into the details of the

nature of the various types of animal minds, including the ways in which they differ systematically from the human mind, as well as from each other.[27]

One traditional objection to the idea of attributing "mind" or "mind-like" characteristics to other organisms can be disposed of briefly—namely, the charge of "anthropomorphism." This is, of course, the idea that it is illegitimate as a general rule to attribute human characteristics, such as mind, to non-human animals. However, a moment's reflection is sufficient to show that any blanket prohibition of this sort is untenable, so long as one assumes the evolutionary continuity of human beings with other animals (as we are doing here). In that case, there can be no principled reason for refusing to countenance the possibility of shared characteristics of whatever sort between humans and other animals. On this view, which of our human features are shared with other living creatures and which are distinctive and ours alone becomes a purely a posteriori question, subject to empirical investigation. There are, of course, methodological worries about how we can determine whether other animals possess phenomenal consciousness and other mind-like characteristics, arising primarily from the fact that other animals lack language and we cannot simply ask them. But these methodological difficulties do not justify a blanket dismissal of the possibility that some or all animals do indeed share these features with us. As Sober (2005; 96) has put it, "There is no evolutionary presumption in favor of assuming that nonhuman organisms differ from human beings, either in terms of their mental or their non-mental characteristics." To adopt such an assumption would be to commit the inverse mistake to the one committed by one who

[27] For example, see Allen & Bekoff, 1997; Bekoff et al., 2002; Dawkins, 1993; Griffin, 2001; Hauser, 2000; Heyes & Huber, 2000; Macphail, 1998; Prete, 2004; Reznikova, 2007; Ristau, 1991; and Walker, 1983; as well as studies too numerous to mention devoted to the higher primates, in particular.

mistakenly attributes some human characteristic to some animal. If the term "anthropomorphism" names the latter type of mistake, then the inverse mistake may be referred to as "anthropodenial," i.e., the failure to recognize characteristics that in fact are shared by human beings and other animals.[28] To adopt either anthropomorphism or anthropodenial as a general principle for research would be a grave mistake. The correct methodological principle here is that of careful empirical investigation on a case by case basis.

However, there is another, weightier objection to the notion that "mind" is not a univocal term. Someone might object that the scientists mentioned in footnote 1 are committing a fallacy of equivocation by referring to the various phenomena they study as different specific sub-categories falling under a single generic category, "mind." And so long as the majority of philosophers understand by the word "mind" conscious thought, there is a great deal of merit in this charge, for what the scientists have in mind is not merely a typology of different ways in which conscious thought manifests itself. Rather, it is widely recognized that it is a serious question whether animals are capable of having conscious thoughts at all.[29] Fortunately, that very difficult issue need not detain us here. That is because our subject is the teleological and normative character of living things as such, and no one imagines that the simplest living systems such as bacteria are capable of entertaining conscious thoughts. What is important for us, on the other hand, is to find a

[28] The distinguished primatologist Frans de Waal recently coined this term; see, e.g., de Waal (1999).

[29] For example, most if not all of the authors cited in the previous footnote explicitly recognize this. Classic discussions of this issue include Davidson, 2001b; Heyes & Dickinson, 1990; Malcolm, 1977; Searle, 1994; M.D. Wilson, 1995. See, also, more recently, Bermúdez, 2003; Biesecker, 1999; Hurley & Nudds, 2006; Lurz, 2009; Wynne, 2004. For the historical background in Descartes and after, see Gunderson, 1964; Massey & Boyle, 1999; Radner & Radner, 1989.

non-tendentious way of articulating what it is that scientists like those mentioned in footnote 1 think they are studying, which seems to them to be a form of mind, and which in their view would still merit some such designation as "proto-mental" or "mind-like," even if it were known with certainty in a particular case that no conscious thought processes were involved at all.

Given, then, that there is an appearance in some or all animal behavior of a feature ordinarily associated with mind—namely, that it is purposive—one obvious way to try to avoid begging any questions would be to call such apparently purposive behaviors in animals "proto-mental" or "mind-like," rather than manifestations of mind as such, thus acknowledging the potential for great differences, as well as some fundamental commonalities. Without prejudging the precise nature of the differences or the commonalities, I think it is intuitively acceptable to put things this way. At least, it is surely true to say that most people would recognize that a dog or a cat possesses something that it is proper to say is analogous to a human mind, in certain obvious respects, but that a dog's or a cat's mind is nevertheless very different from a human's, in other equally obvious and extremely important respects.

If that is so, then it would seem to make sense to speak of a variety or a spectrum of "mind-like" phenomena in nature. And if such a variety or spectrum objectively exists in nature, then it must be a goal of science to study and characterize the differences and commonalities of such phenomena—that is, to create an empirically and theoretically adequate typology of them. While this is primarily a task for science, not philosophy, I believe it will be helpful to us in getting a grip on the philosophical issues in this vicinity to reflect briefly on what the most important types of mind-like phenomena might be.

I begin this reflection by turning to a commonplace distinction, which is not exactly current in everyday speech—though I do believe it is intuitively obvious to almost everyone, once pointed out—but which is widely invoked both in philosophy and the empirical sciences dealing with the mind-like phenomena (comparative psychology, ethology, cognitive science, etc.). I have in mind the distinction between sapience and sentience. Sapience is, of course, a name for the type of mind-like phenomenon specific to human beings, i.e., mind properly speaking. A sapient organism is one that is capable of such higher-order cognitive achievements as reflection, rational deliberation, conscious weighing of reasons, requesting and offering justifications, etc. Many philosophers believe that only sapient organisms are capable of acting in the strict sense of the term, as action requires the conscious weighing of reasons. We will be examining this claim in detail below. More than likely, sapience is closely connected with the capacities for language and a certain level of complex social interaction. Homo sapiens is the only sapient species we are aware of, but if we were to encounter rational beings someday elsewhere in the universe, we would not hesitate to classify them as sapient, as well.

Sentience, on the other hand, is a much reduced degree of mind-likeness, in comparison with sapience. A sentient organism is one that is capable of feeling—i.e., of having subjective experiences, or of possessing what is sometimes referred to as "phenomenal consciousness"—but not of carrying out reflection or rational deliberation. Some philosophers nevertheless believe that it is proper to speak of sentient organisms as acting, because it does appear to us that at least some of the higher sentient creatures like cats and dogs are the authors of their own actions. If that is so, then it must be possible to

act for a reason without being capable of reflecting upon or justifying the action, that is, without being capable of articulating the reason for the action as a reason. This idea, too, we will investigate in depth below. Nevertheless, nearly all philosophers who are prepared to ascribe action and (inarticulate) reasons to sentient organisms do so mainly because of the perceived similarity between sentient creatures and ourselves, a similarity that is primarily grounded in the capability of such creatures for feeling, subjective experience, or phenomenal consciousness—that is, a similarity that is grounded precisely in their sentience as such. Those philosophers would surely balk at ascribing acting for a reason to any organism that was not sentient. But this raises the question: Which organisms are sentient, and how do we know?

I know of no way to answer this question with anything like assurance. It is sometimes claimed that our willingness to ascribe sentience to an animal is a function of its similarity to ourselves, but a more plausible theory is that our intuitions with respect to the sentience of animals track some objective property, such as the capacity for various and flexible behavior. For example, I think that few people who are willing to ascribe sentience to a dog would hesitate to ascribe it an octopus, as well, while an oyster might be viewed as a much harder case to decide. It seems that it is above all poverty and stereotypy of behavior that causes our intuition of sentience to weaken. Scientists are beginning to take an interest in the question of the boundary between sentience and non-sentience (e.g., Balcombe, 2010; Braithwaite, 2010; Cabanac et al., 2009; Ford, 2000;

Panksepp, 2005; Seth et al., 2005), but we are not yet in a position to bring empirical observation to bear on this issue in a decisive way (if we ever will be).[30]

Note that the way we have been talking implies not only a sheer difference between sapience and sentience, but a difference that creates a sort of hierarchy or ordering. That is to say, we have been speaking (very intuitively, I think) as though sapient creatures had to a greater degree something that sentient creatures have to a lesser degree. Whatever it is that all the different kinds of mind-like phenomena have in common, apparently some creatures possess more of it than others. And it is intuitive for us that the ones with more of whatever it is stand "higher" on a scale, are "more advanced" than, the ones that have less of it. The idea is, I think, in crude terms, just that some organisms have mind-like faculties that provide them with a relatively broader scope of activities—that is, with greater powers—while other organisms have faculties that provide them with a relatively more restricted scope of activities, or more limited powers. Such a rank ordering of mind-like phenomena into "higher" and "lower" types is traditionally associated with the notion of the "ladder of nature," or scala naturae.[31] At any rate, there is little doubt that our intuition of sentience does weaken as we descend the scala naturae, with the result that while most people would probably be willing to ascribe sentience to dogs and cats, most would probably be unwilling to ascribe sentience

[30] One might think that anatomical research could settle this question. For example, fish exhibit pain-like avoidant behavior, and also have nociceptors in their skin, opioid receptors in their brains, and other structures similar to ours. But while this argument from analogy is suggestive, it is not conclusive. The fact that certain anatomical structures that we share with other animals are correlated with the subjective experience of pain in us does not prove the existence of the same experience in them. For discussion of this issue, see Allen, 2004; Allen et al., 2005; Carruthers, 1989; Mogil, 2009; Povinelli & Giambrone, 1999.

[31] While the idea of a "ladder of nature" is out of favor with professional biologists, I do not think that its intuitive appeal has been in any way diminished; rather, the "ladder" is just a metaphor naming the intuition described in the text, which I believe is very robust.

to spiders, ants, flies, snails, starfish, corals, sponges, and other small invertebrates with limited behavioral repertoires, not to speak of plants or microorganisms. It is not clear, however, what we ought to infer from this fact, given that we do not have a lot to go on here other than sheer intuition. Perhaps the most plausible inference is that sentience slowly drains away as we descend the "ladder." But whether a spider, say, has only a little sentience compared to a dog, or none at all, is not a question that present-day science has any way of answering.

Fortunately, the fact that we do not know where to draw the line between sentience and non-sentience is of no great importance to my project. The reason is that, as I will show in this chapter, the capacity of acting for a reason has nothing essentially to do with the ability to have subjective experiences. I intend to argue that the capacity of acting for a reason is an inherent property of organisms as such, whether sentient or not. Therefore, my claim would be in no way undermined, even if it turned out that sentience was restricted to the higher animals. However, I recognize that there is a problem in this vicinity relating to terminology. The problem is how I am to refer to non-sentient organisms in a way that does not beg the question of the propriety of the ascription of normative action to them. Conventionally, one makes a distinction between organisms that "act" and ones that merely "behave." Another conventional distinction is between animals with "minds" and animals that are nothing but "machines." But the central claim of this dissertation, regarding the objective existence of teleology, calls into question both of these conventional distinctions. To accept either the "act/behave" distinction or the "mind/machine" distinction would be tantamount to deciding in advance that the concept of agency cannot be legitimately applied to non-sentient organisms. On the other hand, I

cannot simply refer to non-sentient organisms as "agents" without begging the question in the other direction. Therefore, I need a more or less neutral way to refer to non-sentient organisms that acknowledges something of their mind-like properties—which are, after all, the reason why they are candidates for agency-ascription in the first place—without implying that the question of the propriety of ascribing agency to them is a settled matter. That is, I need a word that is parallel to sapience and sentience, but which describes only the most fundamental mind-like properties of living things, quite apart from their sentience (if any).

By "fundamental mind-like properties," I mean such capacities as moving toward a nutrient source and ingesting it, or moving away from a chemical irritant or other threat. Let us call these capacities "feeding" and "fleeing," respectively.[32] Other similar fundamental capacities would include growth, metabolism, self-repair, and reproduction. To my knowledge, there is no word in common use that marks these sorts of capacities, considered as giving rise to a set of philosophical problems that are conceptually distinct from the question of sentience.[33] Seeing that the nature of these capacities is the main topic of this dissertation, I would like to propose a new term here.[34] The purpose of

[32] These examples, it will be recalled, are taken from the discussion of bacterial chemotaxis in Chapter 1. There, I stressed the applicability of various elementary normative concepts to these examples. Here, I am bracketing that question. I will return to it shortly, but for now, I wish to describe the elementary mind-like properties common to all living systems in as neutral and intuitive a way as possible.

[33] Brunswik (1957) introduced the term "ratiomorphic," which was taken up by Lorenz (1977) and others, to describe non-sentient, mind-like phenomena, such as some instinctive behaviors. However, it is an inconvenient term for my purposes. One reason is that, by its etymology, it ought more properly to be an umbrella term referring to what I have been calling "mind-like phenomena" generally. Another reason is that, the way that Brunswik and company use it, "ratiomorphic" contrasts with "rational"—i.e., sapience—thus implying a dichotomous typology that leaves no conceptual space for sentience.

[34] While neologisms ought to be avoided, to minimize confusion some term is clearly needed to refer to the mind-like properties of living things, considered apart from the question of sentience. I believe that this need justifies introducing a neologism in this case.

introducing this new term is to facilitate the philosophical discussion of these capacities, while avoiding confusing them with the capacities for sapience and sentience. And since the mind-like properties in question correspond roughly to what Aristotle called the "nutritive soul," I thought that some term in keeping with the spirit of Aristotle's notion would be most appropriate.[35] For these reasons, I propose the word "appetence."[36] Here is how I stipulate that the word should be defined:

> Definition: Appetence is the capacity of a living system to perform those elementary functions (such as feeding, fleeing, metabolizing, etc.) that constitute it as a living system, whether the system is sentient or not.

By this definition, then, all sentient organisms are appetent, but not all appetent organisms are necessarily sentient (though they might conceivably be). All sapient organisms are, of course, appetent, as well. In short, all organisms are appetent and all appetent systems are organisms. Appetence is the mark of the living; "appetent" and "vital" are more or less synonyms.[37] As in Aristotle, only a relatively small proper subset of appetent organisms are (most likely) sentient, and only a much smaller proper subset still of sentient organisms are sapient. From an evolutionary point of view, appetence is the primordial mind-like capacity of living systems to perform those functions necessary to maintain themselves in existence; sentience (probably) arose at some, much later time; and sapience arose only quite recently (geologically speaking). Each subsequent type of

[35] Aristotle's term is "threp*tikē psykhē*" (e.g., De anima, II.iv.415^{a}24). I merely note the general resemblance between my viewpoint and Aristotle's, without implying that I take his views to be in any way normative for my project, which I do not.

[36] The words "appetence/appetent" do already exist, but they are archaic and seldom encountered today, and so I think available for technical adaptation. In any case, the original meaning of the terms has to do with longing or yearning or striving, which is not altogether inapt.

[37] Here, this claim is more or less stipulative, but I will provide ample empirical and theoretical justification for it in Chapter 4.

mind-likeness was superadded to the previous kinds, so that human beings possess all three: appetence, sentience, and sapience.

With this definition in hand, then, it becomes a much simpler matter to express the Principal Claim of this chapter. It is this: Agency, in the full normative sense of the word, is already implied by appetence, and has nothing essentially to do with either sapience or sentience. That is to say, appetence already consists in the capacity to act for reasons, even without the capacity to reflect upon, justify, or otherwise articulate those reasons, and also without the capacity necessarily to feel or experience anything at all. If this claim can be justified, as I hope to show in a moment that it can, then it is no exaggeration or looseness of language to speak of appetence as a type of "mind-likeness," for it is appetence that is primarily associated with agency, not sapience or sentience, and agency is surely a property of mind, if anything is.

The claim that agency is an essential feature of appetence is a strong one that will require considerable support if it is to be accepted. I will be turning to the task of providing that support in a moment. First, though, I need to make some further distinctions regarding the notions of normativity, agency, and natural ground, as well as two different aspects of the claim that organisms are genuinely normative agents: the Scope Problem and the Ground Problem.

2.3 Some Further Distinctions, and Two Aspects of the Principal Claim

I will be using the term "normativity" in two senses, one narrow and one broad. In the narrow sense, normativity is requirement—the fact that there is something that a given agent is required to do in a certain situation in order to attain a particular end. In

human terms, requirement is often expressed through the imperative mood of the verb, as well as in the form of laws, rules, maxims, and other types of commands, prohibitions, and recommendations, often by means of such auxiliaries as "must," "ought," or "should." In this sense, requirement is usually referred to as "prescriptivity." An important further distinction is between moral and prudential forms of prescriptivity. I will assume that moral prescriptivity is restricted to the human case, and will give it no further consideration here.

But what of prudential prescriptivity—that is, the type of practical requirement that human beings incur just by virtue of being agents with ends? Can it be generalized? Given that there cannot literally be commands in the absence of a commander, it would seem not. And yet the notion of requirement does appear to be more widely applicable than just to the human case. For instance, it is natural to say things like: "Plants must have water"; "Dogs ought to get plenty of exercise"; and "Hearts should circulate blood efficiently." This makes it seem as though there is a kind of requirement that is more general than prescriptivity, or, in other words, that prescriptivity stands in relation to requirement as species to genus. If that is so, then it is natural to ask: What is the nature of this broader concept of requirement, and of the phenomenon to which the concept refers? This is one of the questions I will be investigating in this dissertation.

Though the notion of requirement will be used in a sense that is already broad with respect to that of prescriptivity, it is nevertheless comparatively narrow in relation to another way that the term "normativity" is sometimes used—namely, as an umbrella term to designate a family of closely related concepts for which we seem to have no collective name in colloquial English. Normativity in this broad sense includes such notions as

value, purpose, intention, well-being, and need, in addition to requirement. Much of the discussion below will be directed toward justifying the claim that there is in fact a natural kind corresponding to this umbrella concept.

Normativity, in both senses, is intimately connected to agency—which I take to be the capacity of an agent to act. For instance, take the narrow sense of requirement. If requirement is the fact that an agent must, or should, or ought to do something in a given situation in order to attain a particular end, then normativity in the narrow sense clearly implies agency. Whether, conversely, agency implies normativity is not as clear. If actions are held to be somehow controlled or guided by reasons, and if reasons are held to be metaphysically distinct from causes, then reasons may be said to indicate what should, or ought to, be done in a given situation. This makes it seem as though agency implies normativity. Unfortunately, there are two difficulties with this claim.

The first difficulty lies in determining to what kinds of things the concept of agency may be properly applied. Call this the Scope Problem. The problem arises from the fact that many commentators feel that reasons may properly be said to exist only where the capacity for their conscious weighing, or rational deliberation, exists. Accepting this claim would of course mean that only human beings could qualify as agents in the normative sense. According to this way of thinking, one ought to take care to say that human beings "act," while other animals merely "behave," where actions are held to be guided by reasons, in contradistinction to behaviors, which are merely caused.[38]

[38] I shall define "caused" here as "produced solely by natural law as currently understood by mainstream, contemporary physical science." This somewhat tortured locution is intended to leave open the

Nevertheless, we find it natural to speak of the "reasons" that (at least some) non-human animals have for doing the things that they do. For example, if I observe my cat jumping down from the windowsill and going into the kitchen, and I know that the kitchen is where her milk bowl is located, then I may infer the reason why she went into the kitchen: namely, to get a drink of milk. All of this seems closely analogous to my own behavior when I go into the kitchen from time to time to get a drink of water. If I say that getting a drink of water is the reason why I go into the kitchen, why should I not say that getting a drink of milk is the reason why my cat goes into the kitchen? It is true that my behavior may sometimes be complicated by the existence of countervailing reasons ("Shall I have a glass of beer instead?") and the need to weigh them in a way that my cat's is not. But I see little reason to doubt that our motivations in this case are basically similar—that when my cat is thirsty she experiences something similar to what I experience when I am thirsty; that the pleasure she takes in her milk is not so different from the pleasure I take in my glass of water; and so on. And, indeed, it may often happen that my behavior may be nearly as simple and unreflecting as hers (say, if I go into the kitchen for a glass of water with my mind on something else). If my unreflecting behavior nevertheless qualifies as acting for a reason—that is, qualifies as an action in the normative sense—then why should not her behavior so qualify? It may still be objected that I am trading on an ambiguity in the notion of "a reason." There is also a causal use of the concept, as in asking for the "reason" for an airplane crash or a mining accident. Therefore, one might wonder why my cat's reason for going into the kitchen should not be construed as a purely causal reason of that sort. Of course, one would then have to

---

possibility that, while present-day physical science may lack a theoretical perspective apt for the proper understanding of acting for a reason, such a perspective may for all we know be developed in the future.

explain why that construal of the concept should apply to my cat's behavior, but not to my own behavior. However, that would be a superficial reply. And, besides, there may be some readers who would be prepared to see my own reasons given this same sort of causal construal. Therefore, to address this worry adequately will mean digging deeper, and attempting to elucidate the fundamental difference between causes and normative reasons. Indeed, in a sense, that may be viewed as the central aim of this dissertation. But, in that case, I cannot accept the charge of equivocation yet, as it amounts to the claim that there is no important difference between causes and normative reasons, which begs the main question at issue here.

If my cat's behavior really is so similar to mine as to justify counting it as a case of normative action, still it cannot be denied that it differs importantly from mine in that in my case the potential for rational deliberation is always there, while in her case it is not. This is certainly a significant difference, and it needs to be marked by a terminological distinction. Let us call the cat's form of acting "sub-rational." But then, the question arises: Is subrational action truly normative? To the extent that we are comfortable explaining the cat's behavior by reference to reasons, it would seem that it is. But if we accept this, then obviously we cannot associate the concept of acting for a reason with rational deliberation alone, nor can we sustain a distinction between action and behavior in the traditional way. There are several ways to go here. One would be to deny that subrational behavior is truly action. Another would be say that not all action is truly normative, but a sort of "sub-normative" action also exists. Yet another would be to bite the bullet and admit that our original distinction was misguided, and that the higher animals (at least) are fully capable of action in the normative sense. But since this last

way involves rejecting the association of acting for a reason with rational deliberation, the question would then arise: How are we to understand the capacity of acting for a reason, or (as I shall say from now on) normative agency?

This brings us to the second difficulty involving the claim that agency implies normativity. This difficulty lies in understanding how something like normative agency can exist in nature at all, given the rest of the world picture painted for us by contemporary natural science. Call this the Ground Problem. I note in passing that the Ground Problem is just as much a problem for those who hold that the concept of normative agency is essentially connected with rational deliberation as it is for those who would widen the concept's scope of applicability to include (at least) the higher animals. However we resolve the Scope Problem, the Ground Problem still remains—which is not to say, however, that some solutions to the Scope Problem may not lend themselves more readily than others to a solution to the Ground Problem.

In the remainder of this chapter, I will argue in favor of a radical solution to the Scope Problem that views normative agency as a property of living things as such—that is, normative agency is a feature of appetence. In other words, I claim that all organisms are normative agents, and that only organisms are normative agents in a literal, original, and underived sense. This claim is supported by the master argument that is being pursued throughout this chapter. I will further motivate this claim by showing, in Section 2.7, how it connects to a certain way of posing the Ground Problem—namely, in terms of the claim that what distinguishes organisms as a natural kind is that they must act in order

to preserve themselves in existence.[39] While I have no definitive solution to the Ground Problem to offer, in Chapter 4 I will review some empirical and theoretical considerations from contemporary natural science that I believe hold the promise of opening up conceptual space for something like normative agency.

Having explained what I mean by "normativity," it would be as well to spend a few pages justifying my use of the terms "agent" and "agency," as well.

In everyday parlance, the notion of an "agent" is a very broad one, covering such disparate things as Agent Orange and Agent 007. In this everyday usage, the concept seems to contain two key ideas. The first is that of activity. An agent is simply something considered insofar as it is capable of acting, which we may think of roughly as exerting a causal influence. We may call this capacity for acting "agency." Agency is best seen as a relative concept, as surely no physical system is either purely passive (incapable of acting) or, for that matter, purely active (incapable of being acted upon). Nevertheless, in common speech, we are not usually mindful of such nuances, and simply mean to say, by calling something an "agent," that it is capable of acting (or of performing an "action")—that is, of doing something.

The other key idea in the everyday use of the term "agent" is that of instrumentality. Both Agent Orange and Agent 007 are instruments, or means, to the fulfillment of ends determined by something or someone else. Note that both the notion of a means and that of an end imply action. A means is something that must be done or

[39] I acknowledge many difficulties in specifying what is to count as an "organism" (what do we say about viruses, colonial organisms, cancers, hives, and other doubtful cases?), but cannot consider the problem in detail here. For present purposes, we may consider an individual prokaryotic or eukaryotic cell as the paradigm of an organism.

made to happen in order to bring about an end, and an end by its nature requires something to be done or made to happen in order to for it to be realized. So, the idea of instrumentality introduces an element of teleology, and hence of normativity in a broad sense (see below), into the notion of agency.

However, the ordinary-language notion of agents as instruments is not the concept of agency that is of primary philosophical interest. Rather, the concept of agency that seems to be revealed by considered reflection upon the nature of human action, in particular, is one that is normatively far richer than the mere idea of instrumentality. There is more to James Bond than being a tool of MI6. For one thing, the human agent seems to possess the power of having or setting ends for itself, as well as that of pursuing the means to their fulfillment. And this power, in turn, seems to entail others, such as the capacity to evaluate states of affairs as good or bad, and, above all, a sensitivity to action-guiding reasons of the so-called "justificatory" sort. In this more refined sense, a human agent's actions must be starkly contrasted with the motions of other kinds of objects under the sole influence of the laws of nature. Exactly what all of these claims amount to is something that must emerge from the discussion itself, little by little. What can be said already, however, is that agency of the human type implies quite robust normativity.

By an "organism," I mean a living system (a system that is "alive"). Exactly what being a living system amounts to is likewise something that must emerge from the discussion.

By "normativity," in the strict sense, I mean prescriptivity—that is, the idea of an imperative that something should or ought to or must be done. Normativity, then, implies agency, and, indeed, an action-guiding reason is the paradigm of a normative concept.

Finally, this is perhaps the best place to reflect on the Ground Problem, though our treatment of that problem will be postponed to Chapter 4. What, exactly, does it mean to speak of a "natural ground" of anything, in general, and of normativity, in particular? By a "natural ground," I mean the feature of the natural world that is original or primary or fundamental with respect to some range of phenomena. So, for example, one might say that human reason is the "natural ground" of language (in the proper sense of the term). This would mean that language originated in association with human reason, in a temporal sense (before there were human beings, there was no language, properly speaking); that secondary linguistic phenomena (such as systems of writing) are derived from the primary phenomenon of human speech; and that reason is the necessary condition for language (no entity lacking reason could have language in the proper sense of the term).

But what might it mean to inquire into the natural ground of normativity, in particular? For many readers, phrases such as "natural ground of normativity" and "naturalized normativity" will have the air of an oxymoron. In the philosophical literature, one is accustomed to seeing "natural" used in contradistinction to "normative." And there is, of course, a good reason for this. It is indeed not obvious how normative phenomena, realistically considered, are to find their place within the world described at present by the natural sciences. Mackie (1990), for example, thought that moral phenomena were so unlike the phenomena described by the natural sciences as to be

metaphysically suspect—or "queer," as he famously put it (ibid.; 38–42)—that is, scientifically disreputable and therefore rationally out of bounds. Garner (1990; 143) usefully elaborates Mackie's point as follows:

> It is the peculiar combination of objectivity and prescriptivity . . . that makes moral facts and properties queer . . . It is hard to believe in objective prescriptivity because it is hard to make sense of a demand without a demander, and hard to find a place for demands or demanders apart from human interests and conventions.

And what goes for moral phenomena, in particular, goes for normative phenomena, generally, as D. Phillips (2010; 95) has recently argued.

For this reason, the most common way of understanding the project of "naturalizing normativity" is the anti-realist, or "debunking," approach. On this interpretation, "naturalization" consists in demonstrating that normative discourse fails to refer to any objectively real phenomena. This is a view according to which normative concepts are nothing more than a human "projection" upon the world—at best, a convenient or even indispensable "fiction," at worst, something to be "eliminated" altogether from our now superseded "folk-psychological" vocabulary. But while it may be a common way of understanding the project of naturalizing normativity, this sort of debunking approach is not the one that I will be following here. Rather, I will be inquiring into a natural ground of normativity analogous to the natural ground of language discussed above. Success would be indicated by a picture or account capable of showing how objectivity and prescriptivity could indeed be combined in a feature of the natural world. Such an account would have to do full justice to the distinctiveness of normative phenomena (not merely explain them away), while at the same time showing

how they can be understood as citizens in good standing of the natural world—i.e., not metaphysically queer.

One finds the concepts of normativity and agency discussed mainly in two places: the literatures of ethics, practical rationality, and action, on the one hand, and those of the philosophy of biology, theoretical biology, and cognitive science, on the other. One important question, of course, is what relation the concepts discussed in these two disparate (and mostly non-interacting) literatures bear to each other, if any. In this dissertation, I will be advancing a claim that lies at the interface between these two literatures, and so I will be drawing on both of them in an effort to arrive at a synthetic view of normativity and agency. My claim will likely meet with many objections from both the philosophical and the scientific sides, so my argumentative strategy will be to pay about equal attention to both sets of concerns. For this reason, the discussion of the Ground Problem has been postponed until Chapter 4.

The Principal Claim of this chapter is the following:

> Principal Claim:
> The proper scope of application of our concept of normative agency is to organisms as such.

This claim leads immediately to the First Corollary:

> First Corollary to the Principal Claim:
> The natural ground of normativity lies in the capacity for agency possessed by all living things.

Together, the Principal Claim and the First Corollary tell us that the phenomena of normativity and agency are inextricably intertwined with the nature of life itself. This thesis is robustly realistic; that is, it recommends viewing normativity and

agency as real or objectively existing phenomena, on a metaphysical par with other natural phenomena, and so amenable in principle to empirical investigation and theoretical articulation by the natural sciences.

I will be arguing only that the Principal Claim is plausible, given the way we use normative concepts and in light of everything we know about the nature of living systems, and that it should therefore be viewed as a "live option." I do not claim to be able to provide a conclusive demonstration of its superiority to the anti-realistic, or "debunking," alternative. In order to convincingly demonstrate the superiority of the Principal Claim to the anti-realist view, I would have to be able to provide a rigorous scientific account of what it is for an organism to be a normative agent—something that I am unfortunately not in a position to do. I do, however, hope to supply evidence (beginning in Section 2.7 of this chapter, but mainly in Chapter 4, below) that a conception of organisms as normative agents is at least not inconsistent with anything in our contemporary scientific world picture, properly understood, and that the Principal Claim is indeed a live option, so far as empirical science is concerned.

Before concluding this section focused on distinctions and definitional matters, I would like to explain a few usages I am adopting here for the sake of convenience. Sometimes, I will use "normativity," "the normative," and related locutions in a broader sense, as a convenient way of gathering under one umbrella term both normativity in the strict sense and some other closely related notions connected with action, notably the concepts of value and purpose.

Another term I will employ frequently is "normative agency." Since normativity and agency imply each other, this term is strictly pleonastic (there is no such thing as

"non-normative agency" or "normative motion"). Nevertheless, it will be convenient to use this term, both as a shorter way of referring to both ideas together without always having to say "normativity and agency," and also as a salutary reminder that agency is indeed a normative concept.

Last but not least, I will refer from time to time to the general claim that normative phenomena are an objectively real feature of the world on a metaphysical par with other natural phenomena as "normative realism."

2.4 The Scope Argument

Having distinguished the two important aspects of the Principal Claim of this chapter that organisms as such are normative agents—namely, the Scope Problem and the Ground Problem—and having said a few words by way of anticipation about what is meant by the Ground Problem, which will be taken up in Chapter 4, let us now turn back to the Scope Problem, which will be the focus of the remainder of this chapter.

In the philosophical literature on action, it is frequently assumed that the concepts of agent, agency, and action have application only to sapient organisms, that is, to rational beings. On this view, a rational being is one that is capable of weighing reasons for or against a particular behavior, of justifying a behavior by citing reasons, of behaving in accordance with consciously held norms, and of reflecting upon and criticizing the norms and reasons that guide its behavior. According to this conception, only behavior that is "reason-guided" in the sense of being (potentially) subject to rational deliberation ought properly to be accorded the status of "action." Therefore, only rational beings

possess the capacity for action—that is to say, the faculty of "agency"—and only beings possessing that faculty are properly called "agents." On this view, human beings are the only organisms that count as agents, because (so far as we know) only human beings are capable of deliberating about and justifying their behavior in the required way.

The first point I wish to make is that, by claiming that that this view of agency is mistaken, I am in no way questioning that human beings are the only sapient or rational organisms. That is, I am not claiming that any other organism possesses the unique human faculty of rationality. This faculty certainly endows us with a great many capacities that are not only unique to Homo sapiens, but so far exceed in scope and power the ancestral capacities from which they may be derived as to constitute a real ontological rupture between us and all other animals. I wish to make it absolutely clear that I intend nothing I say herein to call the metaphysical distinctiveness of human beings into question. The claim I will be arguing for is not that some other animals possess human-like rational powers, unbeknowst to us. Rather, I am simply claiming that it is a mistake to view normative agency as essentially connected to rationality. The human form of normative agency is far more sophisticated—far more excellent, I would even venture to say—than any of its animal counterparts. Nevertheless, I believe that it is a fundamental error to identify agency with sapience. Agency is a property that can exist in degrees, and the less sophisticated, or more primitive, forms of agency with which nonhuman animals are endowed still deserve to be considered "agency" in the full normative sense of the term. Or so I will argue.

Once the essential connection between agency and sapience is disputed, the question naturally arises: What is the proper scope of our concept of normative agency?

If normative agency is not to be associated primarily with sapience, it seems that there are two other ways to proceed. Either we may associate agency primarily with sentience, or else we may associate it with appetence. As we shall see presently, a growing number of philosophers (though still a distinct minority) are prepared to question the essential connection between sapience and agency, but nearly all who do so choose the path of associating agency with sentience. Or, perhaps, no other possibility occurs to them, since for the most part their arguments are directed against the mainstream association of agency with sapience, and they almost entirely neglect the other possibility.[40] These arguments will form the subject of the next three sections. Then, in Section 2.7, I shall argue the more radical thesis that there is no good reason to believe that there is any essential connection between sentience and our concept of normative agency, and that the correct way to proceed is to associate agency primarily with appetence.

Here, then, is the informal argument that will occupy us for the rest of this chapter (the "Scope Argument").

First, although it is difficult to provide necessary and sufficient conditions for something's counting as "normative," nevertheless it is apparent that the elementary normative concepts are intimately related to one another conceptually. None of the concepts stands on its own two feet, as it were, but rather each leans heavily on its neighbors for support. Each of the elementary normative concepts is somehow incomplete on its own. For example, it is very hard to explain what we mean by

[40] Burge (2009, 2010) is the main exception to this rule.

"purpose" without appealing to some notion of "value" (Bedau, 1992a, 1992b, 1993). It seems, then, that the elementary normative concepts come as a package deal.

Second, though it is difficult to say precisely in what the "family resemblance" among the elementary normative concepts consists, one feature that surely unites them as a group is that each of them is partly constitutive of agency, in the normative sense. That is to say, each elementary normative concept constitutes an aspect of our complex concept of normative agency. For example, "having a purpose" is part of what we mean by "acting" in the normative sense. A motion that had no purpose (in the sense of "goal" or "end") would not count as an "action." (Snowing is not an "action.") Moreover, "having a purpose," or "end," implies a need to act—namely, to find and employ the "means" appropriate to realizing the end. A state of affairs that no agent ever brought about by taking the appropriate instrumental actions would not count as an "end." (My snow-covered yard is merely the result, not the purpose or goal or end, of its having snowed.[41]) So, the logical entailment between purpose and action runs in both directions.[42]

[41] To be sure, a snow-covered yard might be transformed into an end by human intentionality, as in a child's desire for a "white Christmas," and perhaps someday our improved control over the weather might even permit means to be taken to bring about such an end, but these examples only reinforce the tight conceptual link between purpose and action.

[42] It might be objected that I have simply stipulated that this be the case by excluding those concepts not constitutive of prudential instrumental action from my notion of an "elementary normative concept" in the definition above. However, even in the wider case, moral concepts would seem to be just as closely linked to action as prudential instrumental concepts. After all, to be morally good is to act rightly (justly, beneficently, etc.) towards one's fellow human beings. While it is true that there may be a few normative concepts specific to the human domain for which the link to action seems looser (beauty comes to mind), nevertheless, the link seems very tight in the elementary cases, not just by definition, but rather due to inherent features of our concepts of normativity and agency. And, in any case, the elementary normative concepts are the ones that concern us here.

Third, certain of the elementary normative concepts (e.g., purpose, need, well-being) are clearly properly ascribable to organisms as such.

From the foregoing considerations, we may conclude that all of the elementary normative concepts, as well as the concept of agency, are properly ascribable to organisms as such—i.e., organisms are properly regarded as agents in the full normative sense of the term. In other words, the proper scope of application of our concept of normative agency is living systems as such.

Let us now look more closely at each of these claims in turn.

2.5 The Elementary Normative Concepts and Agency

I would like to begin this section by addressing a twofold difficulty that no supporter of teleological and normative realism can avoid laboring under. First, many readers will find it difficult, when they hear such words as "good," "right," and "ought," to resist construing them in a moral sense. Now, it is perfectly standard in the philosophical literature to make a distinction between moral goodness, morally right action, and the moral ought, on the one hand, and prudential goodness, prudentially right action,[43] and the prudential instrumental ought, on the other. For example, Kolnai (2008) expresses the importance of the distinction in this way:

> When we speak of the good the agent is pursuing (perhaps efficiently, with the appropriate means, and successfully) or of "the good of man," and when we speak of the goodness of conduct or of a "good man," we mean by "good" sharply different things whatever relations we may on closer enquiry discover between

[43] That is, "right" in the sense of "correct" action, or action that is "adequate" in matching means to ends.

> them; and to force them into a conceptual frame of short-circuited unity cannot but result in a distorted vision both of morality and of human wanting, "happiness," and practice. (ibid.; p. 66)

One might have thought that this obvious, indeed commonsense, distinction between the norms relating to moral action properly so called, and the norms relating to merely prudential action, ought to go without saying. Nevertheless, I feel it is not amiss for me to emphasize that nothing I say in this dissertation should be construed in the moral sense, but rather always in the prudential sense.

Another distinction I would like to draw attention to at the outset is between two uses of the term "natural." In one sense of the term that is current in the literature on ethics, practical rationality, and action, "natural" is contrasted with "normative" in such a way that certain aspects of our biological nature get counted as "natural." For example, "justifying reasons" are held to be fully normative, while so-called "enticing reasons"[44] are held to consist of such biological phenomena as desires, inclinations, urges, and other psychologically motivating factors, and for this reason are held to be "natural," and hence non-normative. However, there is another sense of the term in which nothing counts as "natural" that is not expressible in terms of the basic physical sciences. This is the sense in which it is relatively uncontroversial that normativity cannot be "natural," since we can all agree that nowhere in the ontology of physics or chemistry as presently constituted is there to be found anything like prescriptivity, values, purposes, or reasons. Even this latter usage might be disputed on the grounds that our knowledge of physics is not complete, and that we ought not to prejudge what eventually is going to get counted as

[44] This terminology derives from Frankena (1976).

"natural."[45] But setting that point aside for the moment, it is surely the case that such things as desires ought not to count as "natural" on this latter reading, either. Desires no more appear on the list of the fundamental ontological posits of physical science than do values, purposes, and the rest. Therefore, at the very least, we ought to insist that the status of phenomena like desires and other motivating states of organisms should remain an open question, and should not be automatically assigned to the "natural" side of the supposed natural/normative divide. For, to assume that desires and such are "natural" is simply to beg the question against the normative realist. The status of biological phenomena is precisely what is at issue between the normative realist and the anti-realist, and that status ought to determined at the end of inquiry and argument, not at the outset.

> Kolnai (1980) has also given us a crisp statement of this crucial distinction:
>
> > . . . even though value criteria like pleasure, desire, preference, will or decision of the subject (or of a community of subjects to which he belongs) undoubtedly are naturalistic in that they express prevalent tendencies of nature or appetitive facts recognized as sovereign principles . . ., they still do not in any way refer to Nature in a comprehensive and overall sense. Nor do they refer to any concrete power or strength prevailing on the plane of brute factuality. Rather, hedonistic criteria connote an aspect of autonomous evaluation and thereby the hint of a departure from pure Naturalism. (ibid.; p. 15)

Rosati (2003; 502) discusses this crucial distinction, and labels the two senses of naturalism "hedonistic naturalism" and "brute naturalism," respectively. One way of looking at the aim of the present chapter is as an effort to justify and expatiate upon Kolnai's insight into the normative character of our hedonistic (i.e., appetent) nature as biological beings. In Chapter 4, I will attempt to flesh out what Rosati's notion of a hedonistic naturalism might look like in terms of contemporary scientific research.

[45] In Section 2.8 and also in Chapter 4, below, I will argue that normative agency may in fact end up being included in the ontology of physics.

With these preliminaries out of the way, let us turn to the task of inquiring into the proper scope of our concept of normative agency. What arguments are there, then, in support of taking normative concepts to be properly and literally applicable to organisms—that is, to living systems as such?

First of all, as a matter of fact, natural language does sanction such application. For example, it is as uncontroversial as anything can be that a human being, a dog, and a tomato plant all three "need" water in exactly the same literal sense: namely, without it, they will die. So, already we have a concept that is clearly applicable to a very wide range of living systems, and moreover is so applicable precisely because it is conceptually connected with what it is to be a living thing at all. Somehow, to be alive is to have needs. And this fact will loom large in the more empirically oriented discussion in the next section, below. But before turning to those empirical questions, there is much clarificatory conceptual work to do. First, we must ask: Is it certain that the concept of "need" is in fact a normative concept?

It seems hard to deny that it is, at least in my own case. Satisfying my own vital needs appears to me as among the most peremptory of all the commands I am subject to. This fact becomes especially clear when one of them runs an actual risk of not being satisfied. For example, if I am lost in the desert, there is little that will appear to me under the aspect of a higher duty than that of securing some water to drink, in order to save my life.[46] It is true that one of the things that distinguishes me from most if not all other life

[46] If anyone is tempted to say that the point of securing the water is primarily to satisfy my thirst, not to save my life, he is raising an interesting issue that opens out into a number of side-paths. For example, sometimes shipwrecked sailors may drink sea water, even in full knowledge that doing so spells death. There is no space to explore this complication adequately here, but let me make two quick points. First, the sailors will surely hold out against their thirst as long as possible, so long as their reason and will

forms, is that there is indeed one thing that may appear to me as a higher duty, and that is saving another human being's life. So that if I happen to have a last swallow of water in my canteen, I may well give it to my wife or my child or my friend, or even a perfect stranger I happen to be thrown in with. But, notice that the point of my sacrifice is still to preserve life. I am unlikely simply to pour my last mouthful of water into the sand, at least so long as my reason and will do not fail me. Therefore, it seems that satisfying vital needs constitutes the highest of all normative imperatives, whether conceived of prudentially, in relation to the preservation of my own life, or morally, in relation to the preservation of the life of other human beings. Moreover, not only is need (at least in the vital sense we are investigating here) a normative concept itself, it can be shown to be very near to the fons et origo of all the other normative concepts. Let us see how.

From the concept of need immediately flows the concept of value: For a system to have needs is already for it to partition its environment into valenced categories. There are things to be pursued, and things to be avoided, that the needs may be satisfied. "Good" and "bad" are concepts of an immense semantic richness; nevertheless, there are really no more appropriate terms with which to describe these things that are to be pursued or avoided, based on our vital needs.[47] From the idea of pursuing the good proceeds directly from that of end-directedness (or purposiveness), for what else does it

---

are intact. This proves that in their own minds the end of quenching their thirst is secondary and instrumental to the end of preserving their life. Second, at the end of the day we must explain the very existence of thirst in terms of the need of the organism for water, which again shows that the preservation of life is conceptually prior to the quenching of thirst.

[47] Stuart Kauffman offers the suggestion of "yum" and "yuck" (Kauffman, 2000; Kauffman & Clayton, 2006; Kauffman et al., 2008), which, in addition to wit, has the virtue of minimal ambiguity. His intended application of these terms to single cells may be controversial, but at least in human terms, who would deny that when I say "yum," I am saying of something that I find it "good," and likewise for "yuck" and "bad"?

mean to pursue the good than to have achieving a certain good (and thereby satisfying a certain need) as one's end or purpose? As Aquinas famously noted, the concept of value (good and bad) implies the concept of having a purpose or pursuing an end (Summa Theologiae, IaIIae.94.2): "bonum est faciendum et prosequendum, et malum vitandum" [the good is to be done and pursued, and the bad avoided]. From this, the "instrumental ought" (normative requirement) follows immediately; indeed, the "instrumental ought" is already tacitly relied upon in the grammatical form of Aquinas's formulation of this point: faciendum . . . prosequendum . . . vitandum [is to be done . . . to be pursued . . . to be avoided]. If one has the end or purpose of satisfying one's need for water (even short of saving one's life in the desert!), then one ought to seek water to drink. Which means, in turn, that the need for water provides an excellent reason for whatever steps must be taken to secure the water. And so we arrive at the concept of "action."[48]

It is also the case that good can be seen directly to imply should, as Burge (2003) has pointed out. As he puts it: "goods generate shoulds" (ibid.; 513), or, a little less apothegmatically, "goods imply standards for achieving them" (ibid.; 516). McLaughlin (2009) agrees, noting that:

> When we view a causal chain as a series of means and ends, we presuppose something that stops the regress, something that has a good. And this applies whether it is an intentional agent, an organism, or simply anything that can be said to have interests—whether or not it consciously takes interest in them. We presuppose an entity somewhere down the line which has some kind of interests that (ceteris paribus) ought to be served. (ibid.; p. 98)

In this way, one can see the direct conceptual connection between value and action. The connection between need and action is, of course, even more readily

[48] Actually, the concept of action was already implicit in those passive periphrastic gerundives: faciendum, prosequendum, vitandum.

apparent. Finally, the above considerations show us that all of the elementary normative concepts are connected in one way or another with the concept of acting for a reason. Conversely, a direct analysis of the concept of acting for a reason reveals its fundamentally teleological (means-end) structure (behavior lacking a teleological structure does not count as action),[49] from which flow the concepts of value and the "instrumental ought," from which in turn flow the concepts of need and well-being. Agency—the capacity of acting for a reason—then, is implied by the elementary normative concepts, and the elementary normative concepts imply agency. Agency is not something over and above the elementary normative concepts. Rather, agency is a complex concept consisting of a number of different aspects, and some of these various aspects are captured by the individual elementary normative concepts.

One way of summarizing much of the dense network of mutual implication formed by these concepts—a way that is pithy and highlights the central role of the notion of need—is the following:

> Just as a true belief is one which corresponds to fact, so a good action is one which corresponds to need. In another idiom, just as facts are the truth-makers of true beliefs, so needs are the goodness-makers of good actions. (Lowe, 2008; p. 209)

To sum up the discussion so far: While I have been proceeding very quickly, and while each of the concepts mentioned could of course be individually explored at much greater length, nevertheless, I believe that enough has been said to show that there is good prima facie reason for believing that need (at least in the vital sense) is indeed a normative concept. Furthermore, there is excellent reason to believe that the concept of

[49] See, e.g., Delancey (2006), Foot (2001), Okrent (2007), Schueler (2003), Sehon (2005), and G.M. Wilson (1989).

(vital) need lies near the heart of a cluster of tightly interrelated concepts for which there seems to be no colloquial English name, but which I shall refer to here by the umbrella term "normativity."

However, though the concept of vital need lies close to the center of normativity in the broad sense, it does not quite lie at the very center. Need is not quite basic. That is because most of the vital functions that we associate with needs are in fact instrumental, not intrinsic. For example, most living things need to consume water in some form or other. One might suppose that water is an intrinsic need of, say, human beings, if one judged solely from the pleasure that we derive from drinking water when we are thirsty. But of course we all know very well that it is not the quenching of thirst per se in which our vital need for water really consists. Rather, thirst is merely the sign by which our need for water is brought to our conscious awareness. A man lost in the desert might well be able to put up with mere thirst, no matter how terrible, if he did not know that the need represented by the thirst must be fulfilled if he is to go on living. The point is an obvious one that does not require belaboring. To put it in the most general way:

> Vital Need. A biological function is constituted as a vital need only in relation to a normative state of affairs such that the state of affairs can only be preserved by the proper exercise of the function.

In real terms, what is the normative state of affairs that is logically prior to the concept of need? There are two candidates. One is "life" (or, perhaps, "survival" and "reproduction"). The other is "well-being" (or "welfare" or "flourishing"). I will postpone the discussion of the definition of life until Chapter 4. For now, let us focus on the latter concept, of well-being or flourishing.

Kraut (2007; p. 5) states the basic idea of well-being or flourishing very simply: "For most living things, to flourish is simply to be healthy: to be an organism that is unimpeded in its growth and normal functioning." He goes on to show how the concept has nothing whatever to do with sapience or sentience, but is clearly properly ascribable even to plants:

> Such terms as "welfare," "well-being," and "utility" are seldom, if ever, applied to plants. But it is just as obvious a point about plants as it is about animals that some things are good for them and other are not. If something can flourish or fall short of flourishing, that by itself shows that we can speak of what is good for it. (ibid.; pp. 6–7)

In another passage, he is even more explicit on the main point at issue:

> Plants do not have minds. And yet some things are good for them: to grow, to thrive, to flourish, to live out the full term of their lives in good health. Whatever impedes this—diseases, droughts, excessive heat and cold—is bad for them. (ibid.; p. 9)

In other words, logically speaking well-being is not connected with sapience or sentience, but is connected with the fundamental vital functions as such, or, as we would say using the terminology developed above, well-being is essentially connected with appetence.

Foot (2001) makes a very similar point, though she uses the slightly different terminology of "natural goodness"; from the context, though, it is clear that she could just as well say "well-being" or "flourishing":

> . . . "natural" goodness, as I define it, which is attributable only to living things themselves and to their parts, characteristics, and operations, is intrinsic or "autonomous" goodness in that it depends directly on the relation of an individual to the "life form" of its species. (ibid.; pp. 26–27)

Finally, Shields (2012; pp. 122–123) has this to say:

> . . . it makes sense to ask the following question of every living being: Is it flourishing? It is difficult to grasp how this question should be permanently present in the absence of the kinds of norms, whatever their origin, against which appraisals of life may be tendered and debated.

Here, we have finally reached rock bottom in our search for original or underived normativity. The notion of well-being or flourishing is as basic as it gets. The only way to go deeper is to pass from our everyday vocabulary altogether and venture onto the terrain of the natural sciences, in order to investigate in what the well-being and flourishing of living things consists, from a scientific point of view. That is, to go deeper we must pass from the Scope Problem to the Ground Problem, and inquire into the physical nature of life itself—a task that is reserved for Chapter 4.

For now, we must be content with the progress we have made toward solving the Scope Problem, by justifying the claims that our elementary normative concepts are a package deal, and that they are constitutive of agency. In this section, we have seen that the principal elementary normative concepts all do imply one another, and that normativity in the broad sense is essentially connected to the concept of agency. Let us now turn to the question of the applicability of the elementary normative concepts to organisms as such.

2.6 Appetence and Agency

Everything that has been said so far tends to reinforce the intuition we began with—namely, that it is perfectly proper to ascribe normative concepts in a literal way to living systems as such. If only one or two of the concepts were clearly so ascribable—say, need or purpose—then one might perhaps dismiss that fact as a quirk of the language. But if all of the elementary normative concepts are so ascribable, and especially if all of them seem to stand in the same, densely interconnected, network-style

relationship to one another when considered in their application to living systems generally as when considered in their application to human beings, then it becomes much more difficult to argue that the identity of the conditions of application of the concepts in the two cases is merely accidental, and of no importance for our understanding of the real nature of things. On the contrary, there seems to be a genuine mystery here that cries out for an explanation. Why do the world and our way of thinking and talking about it seem to conspire to give every appearance that normativity and agency are objectively real features of organisms, if in fact they are not?

We have already shown that some of the elementary normative concepts, such as purpose, need, and well-being, are clearly ascribable to some of the lower life forms, such as plants. Indeed, this is abundantly clear from ordinary language and our everyday experience of the world. Plants need water (need). Water is good for plants (value). It is unhealthy for a plant to go too long without water (well-being). Some plants turn their leaves toward the sun in order to capture more light (purpose). To capture more light is the reason why some plants turn their leaves toward the sun (having a reason for action). So much is, or ought to be, tolerably obvious.

Nevertheless, for many readers, I suppose that the conclusion of the Scope Argument—the proper scope of application of our concept of normative agency is living systems as such—will seem so difficult to believe as to constitute grounds for rejecting the Scope Argument as a whole. If one looks for a claim to dispute as a result of taking the argument to be paradoxical, that claim will most likely be the one relating to the proper ascribability of any of the elementary normative concepts to organisms as such.

For this reason, I will spend a little extra time attempting to provide independent motivation for the acceptance of this consideration.

The crucial point is to see that the ascription of normativity to living systems (organisms) as such is not only a matter of how we ordinarily speak. If that were the case, then indeed we could not accept the truth of this claim with such certainty. After all, ordinary language might be mistaken on this point, since it developed before so much was known about the material constitution of organisms. But it is not just ordinary language that sanctions the ascription of normativity to organisms, it is biological science itself. Let us see how.

Take, for example, bacteria. Many bacteria, such as E. coli, swim about by means of a faculty known as "chemotaxis."[50] Such bacteria are capable of engaging in two forms of locomotion, or "motility." In the first form (called "running"), the bacteria swim in a straight line. In the second form (called "tumbling"), they move about at random. At the molecular level, the bacteria contain a locomotory assemblage, which is basically a protein motor that makes external appendages called "flagella" rotate, either counterclockwise (for running) or clockwise (for tumbling). This motor is connected to a sensory assemblage, consisting of a complex, transmembrane, protein-receptor array that is sometimes referred to as a "nanobrain" (e.g., Webre et al. 2003). The inner workings of this nanobrain, as well as its chemical linkages to the motor, are immensely complicated, but, in a nutshell, the organ enables the bacterium to sample its external environment for a large number of chemical compounds, to compare the concentrations of these

[50] For brief descriptions and interpretative discussion, see Shimizu & Bray (2003); Wadham & Armitage (2004); and Webre et al. (2003); for full technical details, see Stock & Surettte (1996).

compounds at different times, in this way to determine whether the concentration of a given compound is increasing or decreasing between samplings, and thus to determine whether it is traveling in a favorable or unfavorable direction (where "favorable" means traveling toward an attractant or away from a repellent, and "unfavorable" means the reverse). Finally, by means of its nanobrain the bacterium adjusts the setting of its motor so that if it finds itself swimming in a favorable direction it continues running (i.e., it continues traveling in the same direction) and if it finds itself swimming in an unfavorable direction it begins tumbling (i.e., it tries a different direction).

The elucidation of many of the molecular details of all of this, which are of staggering complexity, represents an outstanding achievement of contemporary science (even if many things remain to be worked out). The precise nature of the relationship between those molecular details and the apparent normative agency of the bacterium in exercising its locomotory faculty is an important theme that I will address in the next section, below. For now, I would like simply to point out that the concepts of normativity and agency do indeed seem to apply in the case of bacterial motility, as just described.

Thus, we may begin with the observation that bacteria need various nutrients, such as lactose, sucrose, and other sugars. Without such nutrients, a bacterium will die. This of course presupposes that self-preservation in life is normative, and death something to be avoided. Indeed, "health," "vigor," "vitality," "viability"—all of these are descriptors that scientists commonly use to refer to the well-being of living things, including individual cells. For example, I.D. Campbell (2008; 2386) claims that "[m]echanical forces, generated while cells migrate, are important for maintaining a healthy cell," while Lloyd and Hayes (1995) expressly ascribe the notions of "vigor,"

"vitality," and "viability" to microorganisms.[51] Given this norm of well-being and the needs generated by it, nutrients then may be said to be good for a bacterium—that is, they are "to be pursued." Thus, the bacterium's motility is end-directed, or purposive. Moreover, a bacterium "should" swim toward its nutrients (if it does not, there is something wrong with it). If it senses that it is swimming in the right direction (toward its nutrients), then it has reason to continue swimming in the same direction, that is, to run (by rotating its flagella counterclockwise). All of this makes it seem natural to say that swimming toward its food is something that the bacterium does, not something that happens to it. In short, bacteria act.

All of this may be said quite naturally, without in any way forcing the language. There is no sense that in describing a bacterium's swimming toward its food as the bacterium's acting, we have slipped somewhere from speaking the literal truth to speaking in poetic fancies or metaphors. That is not to say, of course, that how such descriptions sound to the untutored ear settles the matter. There are certainly objections that can be raised at this point, and I will address some of them presently. Nevertheless, in the ensuing discussion, I think it is important for us to keep in mind that this way of describing the even faculty of motility in the lowly bacterium is perfectly natural, and that this fact is a significant one.

One objection can be dispensed with fairly quickly. One might say that the biologists themselves do not use this sort of normative language to describe bacterial

[51] Of course, such usage of normative concepts by scientists does not in itself show that the concepts cannot be given a reductive analysis. While there is an extensive philosophical literature on the concept of "health" (Ereshefsky, 2009), most of it focuses solely on human beings, and simply presupposes the natural/normative dichotomy at issue here. Wachbroit (1994) importantly shows that the notion of biological "normality" is irreducible to a nonnormative, statistical concept.

motility. Or, to be more precise, they attempt to avoid using such language wherever possible, though they are seldom successful in suppressing normative vocabulary entirely for any length of time.[52] Still why not take our cue from the biologists' own practice? Rather than speak of the bacterium's "pursuing the good," or even "swimming toward its food," why not just speak of its "following a positive attractant gradient"? But notice that this locution is itself a metaphor. After all, it is not as though bacteria are "attracted" up a chemical gradient in the same way that iron filings are "attracted" to a magnet.[53] Bacterial motility is not a matter of a direct reaction to impressed forces or of a tight coupling to an external field. Chemical gradients do not "pull" bacteria along; rather, bacteria carry their own principle of motion within them. They move, as we might say, "of their own accord." That is, they control what they do in such a way that they swim up only those gradients that are good for them. Therefore, motility is not something that merely happens to bacteria, but rather something that bacteria achieve or accomplish. And that is just another way of saying that bacteria "act." Therefore, in point of fact, it is the commonsense normative, agential descriptors of bacterial motility that are literal, and the descriptors that employ physico-chemical terminology known not to be strictly applicable that are metaphorical. Such metaphors amount to a kind of euphemism—an effort to avoid the natural way of describing phenomena such as bacterial motility in terms of normativity and agency.

[52] Cf. almost any page of any molecular or cell biology textbook, to say nothing of works on physiology or animal behavior.

[53] Historically, I suppose, the metaphor must have run the other way—from personal or sexual attraction to magnetic "attraction." But if biologists today speak of a bacterium's food as an "attractant," it is surely in order to assimilate its behavior more closely to that of iron filings, and not that of young lovers.

However, there is a more penetrating form of the foregoing objection that cannot be dismissed so easily. Some might claim that, instead of quibbling about describing bacterial motility at the whole-system level as "pursuing the good" versus "following an attractant gradient," we ought to consider the fact that both sorts of descriptions have (supposedly) been rendered redundant by our knowledge of the molecular details of the chemotaxis subsystem. The idea would be that both sorts of whole-system–level descriptions are little more than convenient verbal summaries that stand in for the myriad physical and chemical details of what is transpiring at the molecular level. In principle, then, if not in practice, one should be able to explain bacterial motility by referring to events exclusively at the molecular level. And indeed if it were true that all the causal work was being done at that level, then, by the parsimony principle, we really should avoid ascribing any ontological significance to whatever purely verbal formulations we may use to summarize those events for our own convenience at the whole-system level.

This sort of objection might seem open to the same reply as before—namely, that living systems are not passively swept along by external causes, but rather are active in the pursuit of their own interests. However, this time, when the objection is expressed in its more radical form, a ready rejoinder becomes apparent. That is the following claim. Science has now (for all practical purposes) fully explained in molecular detail how systems like the bacterial chemotaxis subsystem work. That is, we are now in possession of a (for all practical purposes) complete understanding of the internal "mechanisms" that give rise to the behavior of bacterial motility. While it is true that that type of behavior is very different in detail from the movement of iron filing in a magnetic field, nevertheless, we are now in a position to see that there is no deep difference in principle. Everything is

still happening according to the laws of physics and chemistry; it is just that those laws work themselves out in a special way in certain kinds of systems, which we call "organisms." But that is no problem, because we can fully explain that special way the laws of physics and chemistry have of working themselves out in the case of organisms, by supplementing those laws with a few metaphysically unproblematic auxiliary concepts, such as "negative feedback control," "fitness," "natural selection," and a few others. The capstone of this line of thinking is the observation that we ascribe normative, agential descriptors to manmade machines, as well as to organisms. For example, I might well say that my car "needs" gasoline; that the "purpose" of the gasoline is to make the car go; that if the fuel tank is nearly empty, then gasoline "should" be added; that a nearly empty fuel tank is a "reason" for gasoline to be added; etc. And an automobile, too, is not ordinarily moved about willy-nilly by external forces, but rather contains its own principle of motion within it. In this sense, it too moves "of its own accord."

Since the "machinery" of bacteria is now known to be no different, in principle, from the machinery of automobiles—or so it is claimed—and since we ascribe the same sort of normative, agential descriptors to both kinds of systems, should we not then view organisms and machines as belonging to the same natural kind? Not to put too fine a point on it: Shouldn't we simply say that organisms are machines? And if that is so, then we need not worry about which vocabulary we use. Just as I feel free to say that my car "needs" gasoline, all the while realizing that this is just an elliptical way of describing how the car operates internally, so too (on this view) I should feel free to say that E. coli "need" sucrose, all the while realizing that this is just an elliptical way of describing how bacteria operate internally.

There are two kinds of responses that one might make to this suggestion. One would be to retreat to the position that there is no fundamental difference between organisms and machines, after all, and give up the aim of naturalizing normativity altogether, except by elimination. This is the way urged upon us by Lenman (2005). In a penetrating discussion of McDowell, Foot, Hursthouse, and other "liberal" naturalist authors, he refuses to accept their finding of normativity in the natural inclinations of living things. For example, he writes:

> A nurturing polar bear father . . . is certainly behaving in a way that may surprise ethologists and we may classify it accordingly as defective in a very deflated sense of that word. But surely that's just classification. How does something that deserves to be called authority get into this picture? That's the mystery. A greenhouse full of plants is a space full of healthy and less healthy specimens, specimens that promise to reproduce and live a long time, and specimens that do not. Sure it does. But, except when you are inside it, there are no reasons in your greenhouse. No normativity, certainly no authority, merely a space in which certain natural dispositional properties are distributed in certain ways. (Lenman, 2005; p. 46)

On the next page, Lenman goes on to invoke Williams's (1995; 110) dictum that the complete absence of teleology from nature is the "first and hardest lesson of Darwinism," one which we have yet to take sufficiently to heart.[54]

Lenman's paper is of the first importance because it poses in stark and vivid terms the precise challenge to which any realistic effort to naturalize normativity must respond. But it is not as though there were an actual argument in the quoted passage; rather, Lenman simply assumes that organisms are mechanistic systems to which normative concepts may not properly be ascribed. But of course that is the very point at issue. The reason he is able to get away with such flagrant question-begging is that he is working

[54] On this point, see also Enoch (2006), Sommers & Rosenberg (2003), and Street (2006).

against the background of near-universal agreement with his presupposition that organisms are machines.[55] Therefore, in the final analysis there is no way to respond to Lenman's challenge effectively other than by providing an alternative account of what organisms could be, such that normative agency might be properly ascribable to them.

The other type of response would be to take the bull by the horns and explain why organisms are not machines—that is, why organisms constitute a natural kind, but manmade machines do not. It is easy enough to say (what is obviously true) that organisms have "original" or "intrinsic" normativity, while machines have "derived" or "extrinsic" normativity. But what does that mean? What is original or intrinsic normativity? After all, organisms are physical systems, are they not? How, then, exactly, do they differ from machines?

This is the master question. To pose this question is to ask about the ultimate ground of normativity in nature. I am sorry to say that I have no definitive answer to this question to offer. However, I will begin the preliminary investigation of this question in the next section, focusing on how it may be most fruitfully posed. Afterwards, I will propose some tentative answers in Chapter 4, below.

2.7 Agency and Organisms

We have been moving very quickly, and covering a lot of ground. Perhaps it would be well to pause at the beginning of this last section of the chapter in order to

[55] Davidson's seminal contributions (e.g., 2001a, 2001b) played an important historical role in framing the action debate in this way. For argument that Davidson's position is indeed question-begging in essential respects, see Finkelstein (2007; especially, p. 267).

entertain another type of reply to our master question. For, there is an entirely different way to go in response to the denial of the objective existence of normativity and agency in organisms as such. The other way is to retreat to the traditional view, in which the ground of normativity is held to lie, not in some principle inherent in organisms, in general, but rather in some principle inherent in human beings, in particular. So, I will say a few words now about why grounding normativity in the general sense in some aspect of human nature is an unattractive prospect.

It is true that there are some considerations that provide strong support for ascribing normativity and agency to human beings (together, perhaps, with some higher animals). The most impressive of these is the simple fact that the paradigm case of action for us is undoubtedly human intentional action (where "intentional" means consciously wanted or willed). Consider, once again, the following familiar scenario: It is a hot summer afternoon. I am thirsty. For this reason, I walk into the kitchen to get a drink of water. My walking into the kitchen, opening the faucet at the kitchen sink, and bending over to drink are all intentional actions. But no one supposes that bacteria are capable of forming conscious intentions. Therefore, if such a scenario is to be our model of what it is to act for a reason, then lower organisms like bacteria clearly do not act for reasons and so cannot be agents. In that case, we would indeed be justified in saying that only human beings (together, perhaps, with some higher animals) are agents in a literal sense.

There is of course a great deal to be said about this objection, but I must limit myself here to a few points. First, it seems tolerably clear that many higher animals are indeed capable of intentional action. We have already considered above the case of my cat seeking her milk bowl in the kitchen in the same way that I seek a drink of water. It

seems difficult to deny that she, too, was thirsty, that she knew where her milk bowl was located, and that these were the reasons for her actions. To be sure, the number of the beliefs and desires that contributed to the generation of my actions may have been far greater than those contributing to the actions of my cat, and many of them may have been very different qualitatively from anything ever experienced by a cat, as well.[56] Perhaps I had to deliberate whether to drink water from the tap or to indulge myself in the beer from the refrigerator. Any number of factors may have entered into such deliberations (the day of the week, the time of day, whether I was alone or in company, how many beers were left in the fridge, whether I was watching my weight, whether I was watching my budget, whether I was watching my alcohol consumption, a promise I had made to my spouse, etc., etc.). Any or all of these factors may have been considerations that entered into the generation of my intentional actions. In that case, we would say that my actions were not merely intentional, but the result of "rational deliberation." And surely nothing remotely comparable can possibly have entered into the generation of my cat's actions. Nevertheless, there does seem to be prima facie parity in the fundamental structure of our respective actions—namely, in the relationship between our actions and their reasons, which appear to be available to conscious awareness in both cases—that would seem to justify the same attribution of intentions to both of us. If my cat is capable of being thirsty, of knowing where her milk bowl is located, and of acting for those reasons, that would seem to be enough for her to qualify as having acted intentionally.

[56] Then, again, they may not have been. For, I am capable of acting almost as "automatically" as my cat—that is, without rational deliberation—say, if my mind were otherwise occupied. But in that case, I would still be acting intentionally.

My cat's actions cannot have been generated by rational deliberation, as mine may have been, but they will have been intentional just the same.

However, considerations of this sort merely succeed in relating agency to sentience. How would we go about taking the next step, necessary for my thesis, and relate agency to appetence?

Let us start with a distinction of Railton's (2009). He notes that much of our action is the result, not of rational deliberation, but rather of more or less automatic practical skills or competences—what he calls "fluent agency." Then, he notes that rational deliberation presupposes fluent agency:

> I have no quarrel with treating deliberate choice as one paradigm in the theory of rational or autonomous action—it is certainly an important phenomenon for any such theory to explain. My argument instead is that it cannot be the fundamental phenomenon, for it is built up from, and at every step involves, the operation of countless non-deliberative processes that are—and must be—quite unlike choice. These processes are not self-aware or reflective, yet they are intelligent and responsive to reasons qua reasons. They make us the agents we are, and give our agency its capacity for rational, autonomous self-expression. (Railton, 2009; p. 103)

Railton does not discuss the other animals, but his notion of fluent agency would seem to apply to them as well. Certainly, such notions as automatic skills or competences and fluidity of motion would seem to apply to the pouncings of cats and the acrobatics of squirrels in a perfectly literal way. There remains the issue of whether such behaviors are responsive to reasons qua reasons. This is, of course, the crucial point. As it happens, a number of philosophers have recently begun to argue that the behaviors of at least the higher animals are responsive to reasons in the right way, and thus do qualify as "actions" in the normative sense.

First, Steward (2009a) believes that it is not even necessary to ascribe intentions to the higher animals in order to accept that they are in an important respect the authors of their own actions. Thus, she writes the following, appealing essentially to our commonsense way of speaking and thinking about animals:

> And I should like to insist that the idea that an animal might be able to produce a bodily movement, so far from being a strange piece of metaphysical lunacy seems to be part and parcel of an everyday picture of the world with which we are very comfortable. It is not at all obvious that there must be something deeply wrong with it. Animals have many powers—what is so strange about the idea that one of the types of powers of which they are possessed is the power to control in certain respects movements (and other changes) in their own bodies? (Steward 2009b; pp. 303–304)

Korsgaard's (2009) view of the matter is similar. Though she is more willing than Steward is to ascribe intentions to the higher animals, her reasoning here, like Steward's, remains anchored in our commonsense way of understanding animal behavior:

> Human beings are, after all, not the only creatures who act. The distinction between actions and events also applies to the other animals. A non-human action, no less than a human one, is in some way ascribed to the acting animal herself. The movements are her own. When a cat chases a mouse, that is not something that happens to the cat, but something that she does. To this extent, we regard the other animals as being the authors of their own actions, and as having something like volition. (Korsgaard, 2009; p. 90)

Glock (2009) is still more explicit about the propriety of ascribing intentional states to the higher animals:

> Both in everyday life and in science we explain the behaviour of higher animals by reference to their beliefs, desires, intentions, goals, purposes. These psychological explanations are not causal, at least not in the sense of efficient or mechanical causation. Instead they are intentional in the sense explained above, just as our explanations of human behaviour. In both cases we employ intentional verbs, and we explain the behaviour by reference to the fact that A believes that p, desires X, wants to Φ, etc. (ibid.; p. 242)

Boyle and Lavin (2010; p. 178) agree, observing that the general form of explanation of which intentional explanation is an instance "can apply to nonrational

animals and indeed to plants. Its application marks the feature of living things we are tracking when we say that what goes on with them is subject to teleological explanation."

Finally, Hurley (2003) has addressed the issue of rational deliberation in this way:

> . . . acting for reasons does not require conceptual abilities—not, at least, the full-fledged context-free conceptual abilities associated with theoretical rationality and inferential promiscuity. I appeal to practical reasons in particular to argue that the space of reasons is the space of actions, not the space of conceptualized inference or theorizing. (ibid.; p. 231)

Hurley goes on to raise the issue of whether we can properly speak of a non-human animal's reasons for action as being the animal's own reasons, as opposed to its behavior's being merely conformable to reasons supplied by a human observer, as suggested by Dennett's (1987) notion of the "intentional stance." Here is how she puts this point:

> It may still be objected that while there may be reasons to act that an agent has not conceptualized, these cannot be the agent's own reasons, reasons for the agent, at the personal or animal level. (Hurley, 2003; p. 233)

And here is what she says immediately in reply:

> I disagree. I understand reasons for action at the personal or animal level in terms of the requirements of holism and normativity. Perceptual information leads to no invariant response, but explains actions only in the context set by intentions and the constraints of at least primitive forms of practical rationality. (ibid.)

In these passages, Hurley corroborates my conclusion that sub-rational animals may properly be said to act intentionally, and to be agents. Furthermore, she raises two important issues that I have yet to deal with. One is what it means for reasons of action to be a system's own reasons for acting, as opposed to the system's behavior's being merely conformable to some external reason or normative judgment. (Let us call this the "Intrinsicality Problem.") The other is the precise nature of the relationship between

normativity and holism. (Call this the "Holism Problem.") Both the Intrinsicality Problem and the Holism Problem will become particularly acute once we attempt to extrapolate beyond the higher animals to organisms as such. Accordingly, they will loom large in my discussion of the problem of determining the natural ground of normativity in Chapter 4.

Before turning to that pressing problem, however, I must take one more step in order to motivate taking that problem seriously in the first place. So far, I have only discussed reasons for ascribing literal normative agency to the higher animals. Apart from Steward, the reason cited was basically that the higher animals appear to have intentional states like ours. This material was rehearsed in order to respond to the traditional concerns of many if not most philosophers of action who have usually assumed that literal normative agency ought to be ascribed only to rational beings like us. But even if the position of Steward and the others were to be accepted, that would still leave me only half-way to my stated goal. For, I wish to claim, not just that normativity and agency exist objectively in relation to the higher animals (that is to say, in relation to sentience), but that they exist objectively in relation to organisms as such (in relation to appetence). That is a bridge too far for Steward and the others, and is denied with a greater or lesser degree of explicitness by all of them.

What are some of their reasons for resisting the more radical move I am urging? Interestingly, it does not seem to be the issue of intentionality, or even sentience, that is of primary concern to them (that is to say, none of them argues that action is conceptually linked to conscious intentions). Rather, they make two basic points.

The first point is that they are reluctant to ascribe normative agency to living systems that do not meet some threshold of flexibility of behavior, or "intelligence." The idea seems to be that if the system's behavior is sufficiently stereotyped, then it is simply "automatic" or "mechanical," and no longer meets the criterion of normative agency. Thus, Hurley (2003) contrasts animals with intentions to those supposedly operating according to "invariant" stimulus-response relations (ibid.; 235–236).

There are two different kinds of responses that one might give to this worry. First, as the details of the chemotaxis system outlined above suggest, the behavior of lower organisms is not really as stereotyped as one might think. In fact, it has been observed that no two bacteria can be counted on to respond in precisely the same way to identical environmental circumstances, not even if they are genetically identical (Zimmer 2008; 44–49).[57] In general, one may say that the idea of a rigid stimulus-response relation in the lower organisms is something of a myth. Most of the behavior even of the lower organisms is in fact endogenously generated (Brembs, 2010; Heisenberg, 2009; Maye et al., 2007; Prete, 2004; Simons, 1992; Trewavas, 2009). Moreover, it is now beginning to be acknowledged that the capacity for flexible, purposive behavior is the key to the "robustness," or stability, of the cell, and ultimately of all living things. For example, Kirschner & Gerhart (2005) have put this point as follows:

> The organism is not robust because it has been built in such a rigid manner that it does not buckle under stress. Its robustness stems from a physiology that is adaptive. It stays the same, not because it cannot change but because it compensates for change around it. The secret of the phenotype is dynamic restoration. (ibid.; pp. 107–108)

[57] See, also, Trewavas (1999).

Indeed, Kirschner (2010; p. 3803) goes so far as to claim that "all of biology is built on the dynamic and adaptive capacity of the cell."[58] On this view, "adaptive capacity" is tantamount to an elementary form of "cognition" or "intelligence" that is an inherent property of living things as such.[59]

Nevertheless, it would of course be foolish to deny that the behavior of bacteria is relatively speaking far more stereotyped than that of higher organisms like cats and dogs. The behavior of bacteria may be more flexible—that is, bacteria may be more intelligent—than is commonly assumed, but still I think it is safe to say that no microbe is going to beat a mammal on an intelligence test anytime soon. It is important, therefore, to add—and this is the second response to the first worry—that intelligence is not really a relevant criterion for assessing whether agency is properly ascribed to a system. Rather, responsiveness to reasons is the relevant criterion. And as we have seen above, however limited a bacterium's behavioral repertoire may be compared to a higher animal's, it clearly passes that test with flying colors,

The second worry raised by several of our authors relates to the fact that we commonly ascribe agency only to whole animals, and not to their component parts. Thus, Hurley (2003; p. 234) clearly states that ". . . I understand the subpersonal level as the level of causal/functional description at which talk of normative constraints and reasons no longer applies," and the other authors make similar remarks.

---

[58] See, also, Harold (2001). Piersma & van Gils (2011) and Turner (2007) take a similar view of the adaptive capacity of higher animals. This topic will be discussed in detail in Chapter 4, below.

[59] There is no space here to analyze this controversial claim, but for the idea that "intelligence" may be properly ascribed to living things as such, see Albrecht-Buehler (2009), Ben-Jacob (2009a, 2009b), Ben-Jacob & Levine (2006), Ford (2009), Shapiro (2007), and Trewavas (2003, 2005, 2010). For the closely related view that living processes are inherently "cognitive," see Calvo & Keijzer (2009), Heschl (1990), Lyon (2006), Stewart (1996), and van Duijn et al. (2006).

This makes intuitive sense, and does reflect common sense, which has been one of our chief guides so far. However, we must be attentive here to a distinction that is too easily blurred. It is one thing to say that agency is properly ascribable to whole organisms, and not to their parts. It is something else to say that whole organisms are endowed with a power of agency only over the movements of their bodies as a whole, or over the movements of the external parts of their bodies, and not over the processes internal to their bodies. I am going to argue that there is no good reason in principle to withhold ascription of objectively normative agency to an organism's control of its own internal processes.

I agree, of course, that agency is conceptually linked to the capacities of a system as a whole (and I will examine in detail what this condition amounts to, in Chapter 4). But it does not follow that internal processes cannot be actions of a system, for there remains the possibility that the system as a whole may actively control its own component parts.[60]

> Burge (2009) gives us a clear account of what the holism requirement involves:
>
> I think that the relevant notion of action is grounded in functioning, coordinated behavior by the whole organism, issuing from the individual's central behavioral capacities, not purely from sub-systems. (ibid.; p. 260)

I maintain that this criterion can clearly be met with respect to the active control of a whole system's component parts, just so long as the parts are controlled by the whole

[60] Frankfurt (1997) raises an objection to this line of reasoning when he asserts that the concept of control or guidance is intuitively linked to the conscious actions of whole persons. As he remarks of pupil dilation (ibid.; p. 46): "The guidance in this case is attributable only to the operation of some mechanism with which [the person] cannot be identified." But this objection fails to take into account the fact that it is the whole organism, not the person qua rational agent, with which such subpersonal instances of control are to be identified, as well as the fact that such control (or "regulation") is routinely attributed by scientists to biological systems.

system, and not the other way around. For example, consider the difference between voluntary and involuntary actions within your own body.

We have voluntary control over several of the component parts of our body. Examples include the thoracic diaphragm (breathing), tongue, lips, eyelids, face, mouth, larynx, pharynx, the upper esophageal sphincter (swallowing), the sphincters controlling the bladder and the rectum, the abdominal muscles, and striated muscles, generally.[61] Let us consider breathing. No one, I take it, will deny that by holding my breath for a minute while I am under water, I am acting. And yet, the same internal part (namely, the thoracic diaphragm) is being controlled just as surely when that control is involuntary (i.e., unconscious) as when it is voluntary (conscious). In both cases, the control has exactly the same function—that is, it occurs for basically the same reasons. In both cases, the reason for the occurrence of the internal processes is the introduction of air (containing oxygen) into the respiratory and eventually the circulatory systems. The only difference is that voluntary breathing permits an additional layer of control, permitting greater responsiveness to environmental contingencies. In short, from the point of view of why the body does and what it does, voluntary control of breathing is just more of the same of what is already provided by involuntary control of breathing. Therefore, it is hard to see what principled reason one could give for saying that the voluntary control of breathing qualifies as a normative action while the involuntary control of breathing does not.

[61] The case of the skeletal muscles includes the complication that the voluntary control of the internal part (the muscle) is simultaneously manifested externally (by the movement of the corresponding limb), and some might wish to ascribe the agent's control in such cases solely to the external manifestation. For simplicity's sake, I set this case aside.

I conclude from this example that there is no good reason to deny that, in principle, the whole organism can be in control of its component parts. Of course, it is one thing to say that this departure from our commonsense way of speaking is justified by the phenomena; it is another thing altogether to show how it is possible, from both theoretical and empirical points of view, for the internal processes of living systems to be under the control of the system as a whole. This Holism Problem will be discussed in the context of the Ground Problem, in Chapter 4, below.

2.8 Conclusion

To summarize what has been shown in this chapter, I conclude, on the basis of commonsense linguistic usage and conceptual analysis, as well as some empirical considerations, that there is no principled reason for maintaining that normativity and agency are properties restricted to human beings (sapient creatures) or even to the higher animals (sentient creatures). If that is the case, then we are faced with a decision (assuming we do not wish to be outright dualists) between accepting eliminativism and seeing ourselves as mere machines devoid of any genuine normativity, on the one hand, and seeing all living systems (organisms) without exception as normative agents, on the other. Nothing I have said here excludes our taking the eliminativist path. However, assuming that we opt to follow common sense in viewing ourselves as genuine normative agents, then the arguments I have deployed in this chapter lead to the conclusion that we are entitled to attribute the objective existence of normative agency to organisms (appetent creatures) as such.

In summary, then, in this chapter I hope to have provided strong grounds for accepting the Principal Claim—namely, that the proper scope of application of our concept of normative agency is to organisms as such. Furthermore, we are now in a position to state a Second Corollary to the Principal Claim, to wit:

> Second Corollary to Principal Claim:
> A living system is a physical system that is under a normative requirement to act in order to preserve itself in existence as the sort of physical system that it is.

What does this mean? What is the nature of this "normative requirement," in physical terms? How can such an idea be reconciled with our present scientific understanding of organisms? Is it really possible for us to understand the most primitive living things, such as bacteria, as anything other than mere machines? What, in short, is the ground of normativity in nature, from a scientific point of view? I will take up these important questions, relating to the First and Second Corollaries to the Principal Claim, in Chapter 4, below.

First, however, we must inquire whether there are adequate grounds for supposing that a successful reductive account of the apparent teleology manifest in living things has already been given in terms of molecular biology and the theory of natural selection. It is to this question that I now turn.

CHAPTER 3:
HAS TELEOLOGY IN BIOLOGY ALREADY BEEN SUCCESSFULLY REDUCED?

3.1 Introduction

In spite of the mainly conceptual arguments in the previous chapter that seem to provide us with considerable warrant for applying the concept of normative agency to organisms as such in a literal and univocal way, some readers may well feel that the argument must contain a hidden flaw, or else be simply beside the point, due to the fact that they imagine that the apparent teleology in biology has already been successfully reduced by the natural sciences. If someone believed that the achievements of the natural sciences—notably, molecular biology and evolutionary biology, especially the theory of natural selection—had already successfully demonstrated that teleology is reducible to mechanism, or is otherwise eliminable from our picture of nature, then no argument on the basis of the analysis of our normative concepts would be likely to carry much weight against such a conviction. For this reason, I will try to show in this chapter that such a belief is not in fact well substantiated, and that there is warrant for doubting it.

The first difficulty here lies in making out precisely what is at issue between the realist and the anti-realist with regard to teleology. As already discussed in Chapter 1, by "realism" with respect to teleology in biology I have in mind the thesis that the apparently teleological phenomena that are manifest in all living systems are objectively real. We have agreed to call this position "teleological realism." No heavy-duty

metaphysics is required to support teleological realism. It merely requires being prepared to accord to the manifest teleological properties of living things the same ontological status that we ordinarily accord to their other, non-teleological properties. For example, the claim that the purpose of bone is to support the body would be on an ontological par with the claim that bone is hard. In this way, "body-supportingness" and "hardness" would be properties of the same metaphysical type. This is a modest claim. It simply asks that we be serious about our own linguistic and conceptual practices, that we pay the same ontological respect to all the properties that we ascribe to biological systems, and that we not consign some of those properties (i.e., the teleological ones) to a purgatory of fictive or "as-if" ontological status for no sufficient reason. In other words, I claim no special ontological status for the teleological properties of organisms. A view that accords ontological parity to the teleological and the non-teleological properties of living systems will be realism enough for my purposes here.[62]

With this understanding of teleological realism firmly in mind, then, it is easy to define what is meant by "teleological reduction" ("teleoreduction," for short), without our having to enter into the complexities of the notion of reduction in all of its various interpretations: theoretical reduction, entity reduction, eliminative reduction, and so forth.[63]

> Teleological Reduction (Teleoreduction): To reduce a putative teleological phenomenon is to give an account of the phenomenon that is both empirically and

[62] This means that if someone were an anti-realist about scientific entities in general, but considered teleological phenomena like biological functions to be no less real (or more unreal) than non-teleological phenomena like matter, force, or energy, then that person would qualify as a "teleological realist" for present purposes.

[63] See, e.g., Hohwy & Kallestrup (2008), Rosenberg (2006), and Sachse (2007).

> theoretically adequate and that neither employs any teleological concepts nor presupposes any other teleological phenomena.

This is a minimal definition of teleoreduction, but it is adequate for our purposes. On any of the various accounts of "reduction" in the literature, it seems clear that if an explanation of a given biological phenomenon itself employed teleological concepts, or presupposed some other teleological phenomena, then that explanation could scarcely be said to have "reduced" the teleology manifest in the given phenomenon in any recognizable sense of that term.

This chapter will be organized as follows. In the next section, I will discuss the idea, seldom explicitly stated but often I think implicitly held, that our increasingly sophisticated knowledge of molecular biology in and of itself constitutes a reduction of teleology to mechanism. It will be concluded that molecular biology alone could never show the eliminability of teleology in biology, but could only do so, if at all, in conjunction with some other explanatory framework that carried out the teleological reduction. The theory of natural selection is most often invoked as the explanatory framework best capable of playing that teleoreductive role.

The following three sections will deal with three arguments intended to show that it is far from certain that the theory of natural selection can actually play the role that would be required of it to form the basis for a successful teleoreduction. In Section 3.3, I will present an argument to this effect from the conceptual structure of selection theory. In Section 3.4, I will present a mixed conceptual and empirical argument from the nature of the explanation of causal powers in general in the natural sciences. And in Section 3.5,

I will present a mainly empirical argument from examples of the inherent adaptive capacity of living things that cannot plausibly be attributed to natural selection.

In Section 3.6, I will raise and respond to two objections that might be posed to the foregoing arguments.

Finally, I should perhaps stress that I make no claim to offer a definitive case in this chapter, or indeed in this dissertation, for rejecting the anti-realist view of teleology in general. To do that would require a detailed critique of the empirical adequacy of evolutionary biology as a whole, a very difficult task that lies far beyond the scope of this dissertation. If someone were comfortable with a purely physicalist worldview that had no place in it anywhere for teleology in any form, then nothing I will say here would do much to discomfort that individual. All I claim is that, if one is already convinced of the rationality of taking at face value at least some of the teleological concepts that we employ both in everyday life and in biological discourse, then one is not required to relinquish that conviction on the basis of the notion that molecular biology and the theory of natural selection, either severally or jointly, have already settled the matter by providing us with a successful means of eliminating such concepts from biology. I wish to claim, in short, not that the anti-realist about teleology is rationally required to relinquish his view, but rather that **TRB** is deserving of being taken seriously as a competing account of the ubiquitous appearance of teleology in biology—that is, that it is a "live option."

3.2 A Note on Molecular Biology and Cybernetic-Control Theory

The conceptual basis of the idea that advances in molecular biology constitute a successful teleoreduction is Cummins's (1975) "causal-role" analysis of function. It was explained in Chapter 1 that "function" per se is not an object of analysis in this dissertation. However, it is of course true that most of the properties of organisms to which we intuitively ascribe teleological characteristics are what are called "biological functions." Cummins's analysis was an attempt to analyze the notion of function in such a way that the teleological connotations could be expunged from that concept.

According to this theory (ibid.), the function of a thing is determined by the causal role it plays in—that is, the contribution it makes to—the operation of some larger system. The fundamental problem with this approach lies in the difficulty in explaining exactly what it is about a "causal role" that makes it different from any other causal effect. That is, the challenge is to say what makes certain physical effects of causal processes count as functions, thus differentiating them from all other effects.

Cummins's idea is to link functions to the part-whole relation obtaining within both living systems and complex human artifacts like machines. Obviously, not just any part-whole relation will do, however. An individual calcium carbonate crystal, for example, is a part of a larger system of marble, and might even be said to play a causal role within it. But we do not want to say for that reason that the role played by the crystal within the larger marble system is a function. So, for the causal-role idea to work, we must specify some special kind of part-whole relation that will act as a norm or criterion in relation to which some mere effect may be properly judged a function. Such a criterion would then permit us to differentiate between the blood-pumping effect of hearts and all

other effects, like thumping noises. But what might that criterion be? Cummins himself does not say.[64] Some might be tempted to think that the concept of cybernetic control could provide us with just the help we need here, so let us look at that idea briefly.

The term "cybernetics" is not as fashionable as it once was, but it is still a convenient way of referring to the body of theory that has been developed over the years by mathematicians and engineers to facilitate the design and construction of complex machines. A number of scientists and philosophers have seen in cybernetics the key to understanding the teleomorphic character of living things (Burks, 1988; Rosenblueth et al., 1943; Sommerhoff, 1969, 1990). Let us look, then, at the claim that the concept of cybernetic control provides us with the conceptual tools to effect a teleoreduction of biological functions to mechanism. The chief concept upon which this claim rests is the notion of negative feedback control.

By "negative feedback control," I mean a system so arranged that any deviation from a preferred value in some variable of the system is automatically compensated for, so that the real value of the variable oscillates around the preferred value. This result may be accomplished either through direct physical coupling (as in the governor on a steam engine) or through self-measurement of the transient states of the system and corrective action to minimize the difference between them and the preferred state (as in a thermostat). The homeostatic character of many biological functions does indeed depend upon a sort of negative feedback control that seems highly analogous to that employed by us in the construction of governors and thermostats. This fact, then, is taken as strong

[64] For further discussion, see McLaughlin (2001), Millikan (2002), Mossio et al. (2009), and Nissen (1997); for a defense of Cummins's position, see Cummins (2002).

evidence that the homeostatic (or, more broadly, the goal-seeking) character of biological functions and living things generally is sufficiently explained, at least in part, by the notion of feedback control. If that were the case, then cybernetic-control theory would indeed seem to provide us with a successful teleoreductive explanatory framework.

Now, there can be no doubt that the notion of negative feedback control is a highly useful, even crucial, one for understanding living things. Furthermore, it seems to explain precisely the functional character of biological processes, by showing how physico-chemical structures and events are organized (how they articulate with each other) in space and time in such a way as to preserve a preferred goal state of the system as a whole. Nevertheless, whatever its utility in elucidating how organisms work at the level of mechanisms, the notion of negative feedback control is of no use as a framework for teleoreduction. The reason is the same reason one which defeated Cummins's original causal-role analysis of function. There is nothing within the theory of cybernetic control as such, anymore than in Cummins's idea of the part-whole relation, that allows us to distinguish the goal-state (or "set-point") of the system from any other system state. Superficially, it might seem otherwise, because we can of course trace causal correlations which seem to converge on a particular state. But this fact is no more indicative of the existence of a true goal-state within an organism than the fact that rainwater may converge by many different pathways upon a lake or river basin (see Nissen, 1997; p. 11). The point can be made clearer by thinking about how a simple cybernetic-control device such as a home heating system works.

If I set the thermostat on my home heating system to a certain temperature (say, 65°F), then it may seem as though that set-point is an objective fact about the system, and

that the property of negative feedback control by means of which a steady temperature of 65°F is maintained in my home provides a sufficient explanation of the set-point in wholly mechanical, i.e., non-teleological, terms. But this impression is an illusion. It is only I who determine that 65°F represents the set-point. Granted, the way the system is objectively organized will produce the desired temperature, but the point is that there is nothing about the cybernetic-control system in itself which can distinguish the setting that produces 65°F as the desired setting. If a malfunction should occur within my home heating system that caused the set-point to change to 35°F, there is nothing within the system itself that would know or care that the set-point had changed. In other words, the state of affairs we call a "set-point" in a cybernetic-control system is just a consequence of a particular sequence of causal events. There is nothing about this causal sequence in itself which constitutes the set-point as a goal-state. The fact that one causal sequence rather than another one leads to a state of affairs that is rightly describable as a goal-state in the sense that it is preferred, is a fact superimposed upon the physical system by its designer and its user. Merely shaping a sequence of causal events by means of negative feedback control does nothing to change the fact that no sequence of causal events in itself can constitute any state of affairs as a goal-state, properly speaking. In our manmade systems, it is always the human observer who decides which system state will count as the goal-state.

For this reason, in a biological system, cybernetic notions can only play a teleoreductive role in conjunction with some other account of how such goal states can occur in a non-teleological way. In practice, the theory of natural selection is always invoked to play this teleoreductive role. Therefore, I shall ignore issues relating to

cybernetics henceforth, and focus solely on the putative teleoreductive role of natural selection.

It might seem that the spectacular advances in our knowledge of molecular biology—by which term I mean the macromolecular detail of the "mechanisms" by means of which all cellular functions are carried out—have provided us ipso facto with all the teleoreduction we need. But that is not necessarily the case. The reason why was pointed out long ago by Aristotle. As Cooper (1986) has remarked, in an important study of Aristotle's concept of "hypothetical necessity,"

> Summarily stated, an organ or feature of a living thing is and is formed by hypothetical necessity if, given the essence of the thing (specified in terms of capacities and functions) and given the natures of the materials available to constitute it, the organ or feature in question is a necessary means to the creature's constitution. . . . Explanation by appeal to hypothetical necessity is not an alternative to explanation by reference to goals. It is a special case of the latter kind of explanation . . . (ibid.; p. 134)

In other words, not only is citing the material conditions and motions by means of which a particular state of affairs is brought about consistent with the state of affairs' being teleologically determined, if the state of affairs is in fact teleologically determined, then citing its hypothetically necessary conditions is a necessary component of a complete explanation of the existence of the state of affairs. In any event, merely citing the material conditions and motions by means of which a given state of affairs has been brought about proves nothing one way or the other, as to whether the state of affairs was teleologically determined. More would have to be known in order to decide that point.

In summary, to invoke molecular biology, even with the addition of cybernetic-control theory, as evidence of the successful reduction of the teleological character of organisms to mechanism is flagrantly question-begging. Reference to molecular biology

in this context begs the question of teleology because it always presupposes some other explanatory factor in order to account for the constitution of a particular state of a system as the goal-state. The question we must turn to now, then, is whether the theory of natural selection can provide us with a satisfactory teleoreductive account of biological function in its own right.

3.3 Argument from the Conceptual Structure of Selection Theory

The other main analysis of biological function is the so-called "etiological" analysis, originally presented by Wright (1973), and expanded upon by Millikan (1984, 1989). According to this influential analysis, a function is anything that owes its existence to the fact that it does what it does. In Wright's original formulation:

> The function of X is Z means (a) X is there because it does Z, [and] (b) Z is a consequence (or result) of X's being there. (ibid.; p. 161)

Although according to Wright's original formulation, only particular token X's are mentioned, Wright explicitly drew the comparison between his analysis and cases of biological functions assumed to have been put into place by the process of natural selection. On this view, then, it is the function of hearts to pump blood, rather than make thumping noises, because it is the fact that hearts pump blood that explains their presence in a given organism. As we shall see presently, in Section 3.4, the slide in this discussion from tokens to types will prove to pose problems for the etiological approach to teleoreduction.

At first blush, however, this seems to be an appealing way to reduce teleology in biology to mechanism. It was eventually taken up by Millikan (1984, 1989) and others,

and made the basis for a burgeoning literature in the philosophy of mind aimed at naturalizing mental properties (see, e.g., Macdonald & Papineau, 2006). The basic idea was to note the similarity between Wright's analysis of function and the Darwinian account of an "adaptation." For example, here is a standard version of the definition of an "adaptation," in the "product" sense:[65]

> Characteristic c is an adaptation for doing task t in a population if and only if members of the population now have c because, ancestrally, there was selection for having c and c conferred fitness advantage because it performed task t. (Sober, 2000; p. 85)[66]

Millikan's idea, then, was that the similarity between Wright's etiological analysis of functions and the Darwinian account of adaptations could be exploited to give an unproblematically naturalistic account of the apparently teleological and normative character of biological functions.

Whether this idea in fact works in the way that Millikan and others have claimed has been widely disputed in the literature, on a variety of grounds.[67] Clearly, it will not be possible to cover all of this territory here and, in any case, my focus is not the concept of function per se. What I propose to do in the remainder of this chapter, rather, is to look at three considerations which appear to cast doubt upon the claim that teleology in biology has already been successfully reduced. These three considerations are: (1) an

[65] As opposed to the "process" sense, meaning the process of natural selection itself, by means of which "adaptations" in the "product" sense come into being.

[66] A complete elucidation of this definition obviously requires an understanding of the theory of natural selection, as well as some familiarity with the tricky technical term, "fitness." Unfortunately, both natural selection and fitness are highly contested concepts; there is no consensus in the literature of the philosophy of biology on precisely how either concept should be understood. While these issues are not the focus of this dissertation, they will have to be taken up to a limited extent in a moment in order to properly address the question of teleoreduction, which is.

[67] See, e.g., Boorse (2002), Cummins (1975, 2002), Manning (1997), Maund (2000), McLaughlin (2001), Mossio et al. (2009), Nanay (2010), and Nissen (1997).

argument from the conceptual structure of the theory of natural selection (remainder of Section 3.3); an argument from the manner of explanation of causal powers in general (Section 3.4); and an argument from the general adaptive capacity of living systems (Section 3.5).

First, then, let us look in some detail at how the account of adaptation given above articulates with the broader theory of natural selection, with a view to assessing in a general way the prospects of the latter theory's forming the basis for a successful teleoreduction.

A number of somewhat different formulations of the theory of natural selection have been advanced over the years, but Godfrey-Smith's (2007) recent magisterial review will doubtless provide a benchmark for some time to come. In it, he makes the point that in crafting formulations of the theory there is bound to be a trade-off between generality and applicability to real-life situations. Full generality requires abstracting away from particular features of real organisms, which limits practical usefulness, while adding in particulars improves usefulness but results in loss of generality. For this reason, he gives two different versions of the theory, one stripped down (which he calls a "summary") and the other containing more detail (which he calls a "recipe"). Since our concern here is with the most general features of natural selection, I give his "summary":

> Evolution by natural selection is change in a population due to: (i) variation in the characteristics of members of the population, (ii) which causes different rates of reproduction, and (iii) which is inherited. (ibid.; p. 515)[68]

[68] If a primer on the subject is desired, see Sterelny & Griffiths (1999); for a more in-depth treatment, see Pigliucci & Kaplan (2006) and Sober (1984).

Thus, we may refer to "natural selection," understood in the process sense ("process-selection"), as broadly constituted by three factors, namely, "variation," "fitness,"[69] and "inheritance." Process-selection in this broad sense must be carefully distinguished from two other concepts with which it is easily confused.

One is "natural selection" the product ("product-selection"), that is, the result of the process, or, in other words, a particular distribution of traits in a population. The other is what one might call process-selection in the narrow sense. In using this terminology, one focuses on the actual differential survival and reproduction of some organisms in relation to others that occurs within a population over a period of time (the "biting the dust," as it were). On this understanding, process-selection-narrow-sense is what takes place in between the first two factors (variation and fitness) and the third factor (inheritance). Variation and fitness at time t_n are the result of process-selection-narrow-sense and inheritance prior to t_n, and give rise in their turn to a new round of process-selection-narrow-sense and inheritance after t_n, resulting in a new constellation of variation and fitness at time t_{n+1}. Multiple iterations of this process, then, would constitute process-selection-broad-sense. A portion of the considerable confusion reigning in the literature on selection and causation is due to a failure to attend to the distinction between the broad and narrow senses in which the term "natural selection" gets used.

With Sober's definition of adaptation and Godfrey-Smith's definition of natural selection under our belts, then, let us turn to a consideration of the literature.

[69] Fitness is itself a highly disputed concept; however, as it is not our focus here, I will simply follow Godfrey-Smith's lead and define fitness for our purposes as the factor responsible for different rates of reproduction within a population. For review, see Ariew & Lewontin (2004).

A number of authors have noted over the years that by virtue of its conceptual structure the theory of natural selection inevitably presupposes the existence of already-functionally-integrated organisms at each round of variation-and-selection. In short, a viable organism must already exist before it can be selected. For instance, over a century ago early critics of Darwin had already drawn attention to this problem by means of such catch-phrases as "the origin of the fittest" (Cope, 1887) and "the arrival of the fittest" (De Vries (1912).[70] More recently, Cummins (1975) claimed that:

> We could, therefore, think of natural selection as reacting on the set of plans generated by mutation by weeding out the bad plans: natural selection cannot alter a plan, but it can trim the set. (ibid.; pp. 750–751).[71]

Similarly, G.P. Wagner and coworkers (2000; pp. 822–823) recently put the point even more bluntly: ". . . to state that a genetic mutation led to a favored character, which, in turn, was selected is utterly uninformative in explaining innovation."[72] Finally, Sober (1984) put the matter thus in a passage that has attained classical status:

> Natural selection does not explain why I have an opposable thumb (rather than lack one). This fact falls under the purview of the mechanisms of inheritance (Cummins, 1975). There are only two sorts of individual-level facts that natural selection may explain. It may account for why particular organisms survive and why they enjoy a particular degree of reproductive success. But phenotypic and genotypic properties of individuals—properties of morphology, physiology, and behavior—fall outside of natural selection's proprietary domain. (ibid.; p. 152)[73]

Neander (1995) has dubbed this the "Negative View" of natural selection: namely, the view that natural selection (in the broad process sense) can explain changes

[70] Cope's (1887) carried this title (cited in Badyaev, 2011). De Vries's (1912) exact words were: "Natural selection may explain the survival of the fittest, but it cannot explain the arrival of the fittest" (ibid.; pp. 825–826; cited in McLaughlin, 2011; p. 203—see also Fontana & Buss, 1994; p. 1). De Vries attributes the phrase "arrival of the fittest" to the American botanist, J. Arthur Harris.

[71] Cited in Neander (1995a; p. 65).

[72] Cited in Love (2008; p. 875).

[73] Cited in Neander (1995a; p. 63).

in the distribution of traits in populations over time, but cannot explain the origin of those traits. The contrary view—the view that natural selection in the broad process sense plays a creative role in the origin of novel traits—may be called the "Positive View."

Neander's (1995a) article led to an explosion of literature on the general question of the explanatory power of the theory of natural selection. Unfortunately, a number of issues have tended to get run together that need to be kept separate. The main line of literature on the topic of the Negative vs. the Positive Views (the one of primary interest to us here) is represented, in chronological order, by Neander (1995a), Sober (1995), Neander (1995b), Walsh (1998), Forber (2005), Nanay (2005), Stegmann (2010), and McLaughlin (2011). Another discussion, with which the foregoing has tended to become partly conflated, has to do with the question of whether natural selection should be thought of as a causal process—perhaps as the resultant of distinct forces on the Newtonian model, such as selection (in the narrow process sense), mutation rate, random drift, etc.—or whether it should rather be understood as only the statistical sum of other causal processes, all of which are located in the individual organisms constituting a population over a period of time. Some of the most important papers and books in this line of discussion would include Walsh et al. (2002), Stephens (2004), Millstein (2006), Brandon (2006), Brandon & Ramsey (2007), Shapiro & Sober (2007), Walsh (2007), Glennan (2009), Stephens (2010), McShea & Brandon (2010), and Lewens (2010). Other, related debates involve the question of how the concept of fitness interacts with the causal versus the statistical interpretations of natural selection (Matthen & Ariew, 2002; Bouchard & Rosenberg, 2004; Matthen & Ariew, 2009), and the question of explaining,

not how a particular kind of trait comes into existence, but how a particular individual comes to have a particular trait (Matthen, 1999; Pust, 2001).

The last question—dubbed "origin essentialism" by Pust (2001)—has been distinguished sharply from the question of the origin of traits by Forber (2005), who shows that it requires for its solution extra-scientific assumptions.[74] However interesting the deeper metaphysical questions raised by origin essentialism may be, they are clearly far removed from our present concerns. For this reason, we can safely ignore this strand of the complex debate on the explanatory power of natural selection.

What about the strand of the debate dealing with fitness? While providing a coherent account of the concept of fitness and its role in the overall theory of natural selection has been an important problem of long standing in the philosophy of biology (see Ariew & Lewontin, 2004), it is far too complex an issue for us to delve into in a serious way here. In any case, our main focus here is not on understanding the logical structure of the theory of natural selection in all its complexity, but on inquiring into the claim that that theory has already provided us with a successful teleoreduction. With respect to this issue, it seems that the strand of the debate of greatest interest to us is whether natural selection (in the broad process sense) can be construed in a causal fashion at all. The precise way in which the concept of fitness articulates with the causal versus the statistical interpretations of the theory of natural selection would appear to be of secondary interest to the present discussion. For this reason, I also table discussion of the fitness strand of the debate.

[74] On the "origin essentialism" question, see also Nanay (2005) and Stegmann (2010).

This brings us back to the primary remaining distinction between strands of the debate: namely, (1) the discussion of the Negative View versus the Positive View as such and (2) the discussion of the causal versus the statistical interpretations of natural selection (broad process sense). Now, the first thing to notice in this connection is that the statistical interpretation of natural selection implies the Negative View, because if there is no separate causally efficacious process of natural selection, then clearly it cannot explain the origin of traits. However, the causal interpretation of natural selection does not imply the Positive View, because it might be the case that on the causal interpretation, it turns out that natural selection causes only population-level phenomena such as the distribution of traits within a population. On such a view, the causal interpretation of natural selection would be compatible with the Negative View of the origin of traits, because the origin of traits is an individual-level phenomenon. Inasmuch as an adequate discussion of the causal versus the statistical interpretations of natural selection would require an extensive detour into the literature of the metaphysics of causation, the prudent course seems to be to grant the causal view, for the sake of argument. Then, the question remaining will be this: Given that the process of natural selection (broad sense) is causally efficacious at the population level, is there reason to believe that it is causally efficacious at the individual level, as well?

It seems clear that the answer to this question must be a qualified Yes. The reason is that under a certain set of plausible assumptions,[75] cumulative selection is empirically demonstrable (Forber, 2005; McLaughlin, 2011; Nanay, 2005; Neander, 1995).

[75] These include sexual reproduction, and so meiosis and genetic recombination, as well as Malthusian constraints. The latter may be accepted as general for the sake of argument, but note the limitation on generality imposed by the first assumption.

"Cumulative selection" here refers to the fact that under natural selection the genetic structure of a population may change in such a way as to introduce a statistical bias with respect certain combinations of genes. This bias in the genetic structure of a population may, in turn, lead to a much greater than expected frequency in certain mutations, leading to a series of relatively fitter novel traits in successive generations. This is the general model for the evolution of complex traits such as the eye and so forth. If it is right that such genetic biasing can occur (and there seems no good reason to deny it), then it follows that natural selection in the broad causal sense, though basically a population-level process, can nevertheless have causal effects at the individual level. And this means that natural selection can (partly) explain the origin of traits.

But is this admission not fatal to the Negative View? Yes and no. It certainly establishes that the effects of natural selection cannot be ignored, when explaining the origin of traits, at least in sexually reproducing species. So, if the Negative View is the view that natural selection has no causal role whatsoever in the origin of traits (call this the "Pure Negative View"), then that view is clearly mistaken.

But the Pure Negative View is too strong for our purposes here. For our purposes here, we need to show, not that natural selection has no role at all to play in the origin of traits, but only that its role is not decisive, so far as the teleological features of organisms are concerned. Call this the "Mixed Negative View." As Nanay (2005; p. 1101) has remarked, speaking of the literature we have been discussing, "it needs to be emphasized that the question is whether cumulative selection can play a role in explaining adaptation, and not whether it can fully explain adaptation." But the aims of the contributors to this literature are not the same as our aims here. The question we must ask here, rather, is

whether, in spite of the refutation of the Pure Negative View, it nevertheless remains reasonable to deny that the causal role of natural selection in the origin of traits is clearly of sufficient scope as to adequately account for the apparently teleological features of living things. And I believe the answer to this question is also Yes—that claim may still be reasonably denied. Let us see how.

From a logical point of view, the Mixed Negative View and the Positive View are symmetrical. From the perspective of the Mixed Negative View, a viable organism already presents itself at each new round of variation and selection. Clearly, the functionally coordinated organism must already exist before it can be selected. On this view, we assume that the functional coordination of the organism is prima facie evidence of teleological determination, and since that functional coordination is presupposed by the theory of natural selection, the theory is in no position to reduce the apparent teleology in biology to mechanism. However, viewed from the perspective of the Positive View, one could as well say that at each new round of variation and selection, there already exists a mechanistic structure put in place by previous rounds of the process. On this view, we assume that the organism is mechanical. Then, each small random variation is merely added to what has already been constructed little by little during previous rounds of the natural selection process. In this way, it appears that the overall process might indeed be able to account for the origin, not just of novel traits, but of the entire, complex structure of organisms in a purely mechanistic fashion. In short, each view interprets the nature of the organism in accordance with its own presuppositions. How, then, can we adjudicate between them?

First, it should be noted that if there is epistemic parity between the two views, that is already enough to provide grounds for doubting that a successful reduction of teleology to mechanism has in fact been accomplished. But there is no question that it would be nice to be able to do better than that. To do so, we would need independent warrant for the Mixed Negative View's assumption that the functional coordination of the organism is prima facie evidence of teleological determination. I will argue in Section 3.5, below, that such warrant is not lacking. Here, though, I would like to make a weaker claim, which I believe will still be sufficient to cast doubt on the Positive View.

Both sides will acknowledge, I think, that in between random genetic variation and the process of selection in the narrow sense (the biting of the dust), there must intervene a process of ontogenetic development, during which the novel phenotype is constructed. But this is not, in general, a "random" process with respect to the viable adult phenotype. If one concentrates on a certain class of examples, such as industrial melanism, in which the novel trait is purely passive, it may seem that it is. But such cases must be the exception, not the rule. If, according to the Positive View, we are to view the entire organism, in all of its immense complexity, as having been constructed step by step through the process of natural selection (in the broad sense), then the vast majority of novel traits that will need to be accounted for will be of a far more active nature: digestive systems, circulatory systems, respiratory systems, sensory and nervous systems, locomotory systems, and so forth, all intricately coordinated internally and with one another. To be sure, on the assumption of universal mechanism, we can always view such systems as nothing more than fantastically intricate congeries of mechanisms, each put into place gradually by incalculably many past rounds of random variation and selection.

However—and this is the crucial point—what accords particularly ill with the Positive View's picture is the fact that each novel trait, even if generated entirely randomly at the level of the genotype, must still be functionally integrated into a novel viable phenotype. And this is a process that is very difficult to represent in anything other than teleological terms, for it is the whole living system that adjusts itself to accommodate the novel trait. This capacity for compensatory adjustment under the constraint of viability happens at each round of variation of selection; it is arguably the precondition for the success of the entire venture.

Walsh (2003) has articulated this point especially well:

> The bearers of biological form are organisms and each organism faces the tribunal of the environment as a corporate entity, not as a loose aggregate of independent traits. One consequence of this is that at each stage of its development from egg to adult an organism must be an integrated, functioning whole. Another is that for any form (trait) to arise in an organism at a time, it must develop from the materials and processes at the organism's disposal at that time. The requirement of integration and the processes of development that produce it leave their distinctive traces on biological form. It seems reasonable, then, to suppose that one might appeal to the processes of development in explaining the nature and distribution of biological form. (ibid.; p. 281)

Walsh then makes another important distinction, between what he calls the "transparent" and the "opaque" views of ontogenetic development:

> By "transparent" I mean simply that the magnitude and direction of changes in genotype space correspond closely to the magnitude and direction of changes in phenotype space. If the genotype-phenotype relation were transparent, changes to the kinds and frequencies of genotypes wrought by the processes operating over genotype space—replication, segregation, recombination, mutation, etc.—could be mapped straight on to changes in phenotype space. Changes in biological form could then be exhaustively explained by a combination of processes at the genotype level (e.g., Mendelian inheritance) plus selection operating exclusively at the level of phenotypes. The details of individual development would not matter much to the explanation of adaptive evolution. (ibid.; p. 282)

Finally, Walsh goes on to describe what the "opaque" view consists in, and what it implies:

> On the other hand, if the genotype-phenotype map were opaque, then changes in genotype space would not translate in any simple way into changes in phenotype space. Large transitions in genotype space may correspond to small or no changes in phenotype space, while small (or no) changes in genotype space may correspond to major phenotypic differences. If the genotype-phenotype map introduces changes of its own to phenotype space, then we need to invoke (at least) two sets of causal processes, or forces, in order to explain phenotypic evolution: the force of selection and the various processes of development. (ibid.; pp. 282–823)

In summary, the transparent view of the ontogenetic process is in fact a crucially important tacit assumption of the Positive View, without which that view loses much of its persuasive force. However, there is now reason to believe that the transparent view is no longer tenable. More particularly, it is now clear that DNA is only one causal factor within an unexpectedly complex, hierarchically layered, and multiply recursive network of causes regulating phenotype construction.[76] In other words, the old transparent view of the genotype-phenotype relation upon which the synthetic view of evolution tacitly rests turns out to have been a vast oversimplification of the actual situation. Since there is little doubt today that the opaque view is by far better supported empirically, it follows that the Mixed Negative View would seem to enjoy at least some degree of positive warrant vis-à-vis the Positive View.

Is it possible to specify with greater precision the implications of the opaque view of the genotype-phenotype relation for our understanding of the evolutionary process? Probably, it is too soon to say much on this head with confidence, but at a minimum one

[76] See, e.g., Bentolila (2005), Beurton et al. (2000), Griffiths & Stotz (2006), Jablonka & Raz (2009), Mameli (2004), Mattick (2009), Moss (2003), Piro (2011), Portin (2009), Shapiro (2009, 2011), and Stotz (2006).

may point to the burgeoning research program that has come into being in recent years, which is attempting to do just that. I have in mind, of course, the rather diverse group of scientists who have produced a distinctive body of empirical findings and theoretical concepts that goes under the collective label of "evolutionary developmental biology" ("evo-devo," for short).[77]

The essence of the evo-devo view, insofar as it is relevant to our present concerns, is that even if it is true that most or all genetic changes are random with respect to fitness, nevertheless, the responses of the epigenetic systems within which the genome is embedded are not at all random.[78] In the case of ontogeny (individual development), this is easy to see. Any genetic change in an offspring (in relation to the parent)—if it has any phenotypic effect in the offspring at all—will constitute a perturbation that must be compensated for by the developmental process by which the offspring's phenotype is constructed. Let us call developmentall constructed adaptive changes "phenotypic accommodation," following West-Eberhard (2003). Since all novel phenotypes are the result of phenotypic accommodation, the raw material presented to selection (in the

[77] For scientific work in the evo-devo tradition that is sensitive to theoretical issues, see, e.g., Amzallag & Lerner (1995), Badyaev (2005, 2009, 2010, 2011), Carroll (2005), Gerhart & Kirschner (2007), Kirschner & Gerhart (2005), Koonin & Wolf (2009), Laubichler & Maienschein (2009), Monteiro & Podlaha (2009), Müller & Newman (2003, 2005), Palmer (2004), Pigliucci (2009), Pigliucci & Müller (2010), Pigliucci et al. (2006), Shubin et al. (2009), Sultan (1992), Uller (2008), G.P. Wagner et al. (2000), and West-Eberhard (2003, 2005). For more philosophically oriented discussion of the significance of evo-devo, see Amundson (2005), Callebaut et al. (2007), Depew (2011), Kaplan (2008), Moss (2003), Müller (2007), Neumann-Held & Rehmann-Sutter (2006), Robert (2004), and West-Eberhard (2008).

[78] Bird (2007) points out that the word "epigenetic" is often used ambiguously, referring now to factors affecting non-heritable changes in the developmental process (ontogeny), now to potentially heritable non-genetic factors (e.g., DNA methylation patterns) and non-nuclear factors collectively referred to as "maternal effects" (oocyte structures, proteins, mitochondrial DNA, etc.). However, insofar as the former factors are subject to "genetic assimilation" (see below), the distinction will not matter for purposes of the present discussion. In both cases, the genotypes transmitted to posterior generations will have been partly shaped by compensatory processes acting under the constraint of viability, and hence will not be random with respect to viability.

narrow sense) will never be truly random with respect to the viability of the organism. This also supports the Mixed Negative View.

That is the first point. In addition, phenotypic accommodation may occur due to external perturbation of the developmental process, in the absence of any genetic-level change at all. Now, at first it might be supposed that a novel adaptive phenotype of this second sort would be of no evolutionary consequence. Even though it could reliably recur over a number of generations so long as the inducing environmental perturbation was present, nevertheless, (the thought would be), since it has no specific genetic basis, it could never become properly heritable. Or so it would seem. However, it has been demonstrated that such phenotypic changes can become stabilized at the genetic level after a number of generations, in such a way that the novel phenotype may come to be reliably constructed even in the absence of the original environmental inducer, by means of a process known as "genetic assimilation."[79] In such a case, phenotypic change clearly precedes and causally influences genotypic change. This phenomenon provides support even more clearly to the Mixed Negative View.

Again, I stress that this is a new area of research. Not only are there many admitted lacunae in our understanding of these phenomena (such as the precise way in which genetic assimilation works), but it is virtually certain that many of our present viewpoints and hypotheses will end up having to be substantially modified. Even so, it is not likely that the evo-devo viewpoint will be wholly overturned. If recent history is any basis for judgment, it is much more likely that it will continue to expand and encroach

[79] For details, see West-Eberhard (2003); for discussion, see Badyaev (2011) and Pigliucci et al. (2006). Genetic assimilation in this sense is a developmental generalization of the Baldwin effect, which as originally conceived of was linked to learned behavior (B.H. Weber & Depew, 2003).

upon the mainstream view. Of course, as Pigliucci (2009) and others have argued, there may be no need to look upon evo-devo as an alternative to the synthetic view. Rather, the two viewpoints probably can and perhaps should best be viewed as complementary. Nevertheless, if the following sampling of cutting-edge biological opinion is to be trusted, the modification to the mainstream view has been and will continue to be substantial.

Thus, Uller (2008) observes that:

> There is a growing awareness that evolutionary biologists need to redirect their focus away from a narrow gene-centered view and toward developmental aspects of phenotypic evolution, to fully understand the evolution of organismal form and function. (ibid.; p. 432)

Similarly, Badyaev (2011) asserts that:

> . . . in similar need of re-evaluation is the original view on the links between functionality (produced primarily by plastic and emergent features of phenotype) and inheritance that preoccupied evolutionary thinkers since the birth of the theory of evolution. (ibid.)

Monteiro & Podlaha (2009) have this to say:

> There is still much to do in order to fully understand how novel complex traits evolve. . . . This work is difficult and time-consuming, but the question at its core—the genetic origin of new and complex traits—is probably still one of the most pertinent and fundamental unanswered questions in evolution today. (ibid.; pp. 0214–0215)

If these opinions are to be taken seriously—as I believe they ought to be—then the evo-devo viewpoint is of vital importance for the present inquiry. The reason is that evo-devo elevates "phenotypic accommodation"—that is, the inherent compensatory or adaptive capacity of organisms—into a distinct cause within the overall evolutionary process. But if that is correct, then there is certainly quite a bit of room for doubt that the theory of natural selection as usually understood provides a sufficient basis for teleoreduction.

This notion of an adaptive capacity inherent in living things as such is worth pursuing a bit further, here, as we will also be relying on in it throughout the rest of this dissertation. Other names for this phenomenon that one encounters in the literature include "robustness," "adaptability," and "resilience." It is also closely related to such concepts as "homeostasis," "canalization," and "stability," on the one hand—indicating the system's ability to restore a previous dynamical regime following perturbation—and "plasticity," "distributed robustness," "degeneracy," and "evolvability," on the other—indicating the ability to discover a novel dynamical regime consistent with viability.[80] Given that the theoretical basis for these various distinctions remains to be clarified (and that the terms themselves are not always employed consistently in the literature), I will stipulate the following usages here. To refer to a life-token system's capacity for compensatory adaptive action to restore a previous dynamical regime, I shall use the term "robustness." To refer to its ability to discover a novel dynamical regime, I shall use the term "plasticity." I shall use the term "adaptivity" as an umbrella term to refer collectively to the properties of robustness and plasticity, where the distinction is not important.[81]

Since the process by which phenotypes are produced is clearly adaptive, the Mixed Negative View would appear to be correct and adaptivity is indeed presupposed

[80] See Edelman & Gally (2001), Kirschner & Gerhart (1998), Stelling et al. (2004), Tokuriki & Tawfik (2009), Wagner (2005), and Whitacre & Bender (2010). "Degeneracy," in its biological sense, refers to the ability of structurally diverse component parts to fulfill the same function (in contrast to "redundancy," which implies multiple copies of structurally similar parts). "Evolvability" refers to an organism's capacity to generate viable, heritable phenotypic variation.

[81] Following the usage of the Autonomous Agency Theorists (e.g., Barandiaran & Moreno, 2008; di Paolo, 2005). In the final chapter, I shall additionally speak of "dynamical stability" as a way of referring to a living system's mode of persistence. On this usage, then, adaptivity is the property of a life-token system that is (in part) responsible for the system's dynamical stability, or persistence as the sort of system that it is.

by the theory of natural selection at each step of the cycle of variation and selective retention. Since adaptivity operates under the global constraint of self-preservation of the life-token system, it is clearly normative. And since adaptivity is presupposed by the theory of natural selection, the normativity inherent in it cannot be "reduced" or otherwise explained by that theory.

For all of these reasons, a mechanistic view of the relationship between genetic changes and phenotypic variation is no longer tenable. As Shapiro (2009) has noted forthrightly:

> If we are to give up the outmoded atomistic vocabulary of 20th-century genetics, we need to develop a new lexicon of terms based on a view of the cell as an active and sentient entity, particularly as it deals with its genome. The emphasis has to be on what the cell does with and to its genome, not on what the genome directs the cell to execute. (ibid.; p. 23)

It is simply no longer plausible to view the developmental relationship between genetypic variation and phenotypic variation as one that is "random" in any meaningful sense. But without this hidden presupposition, the theory of natural selection fails to provide a reductive account of the teleological organization of living systems or the global normative constraint under which adaptive biological processes occur.

Still, when all is said and done, one must admit that none of these arguments is conclusive against the Positive View, and so against the claim that the theory of natural selection has already provided us with a successful teleoreduction scheme. One could, if one wished, always adhere to the Positive View simply by insisting that even adaptivity itself has been put into place by natural selection along with everything else, on the assumption of universal mechanism. Therefore, a stronger claim in support of the Mixed Negative View, adducing empirical evidence in support of the positive non-reducibility

of adaptivity, would be highly desirable, if such a thing were to be had. I will attempt to adduce such empirical support in Section 3.5, below.

First, though, I wish to discuss an argument from the nature of the explanation of causal powers in general in the natural sciences.

3.4 Argument from the Explanation of Causal Powers in General

Let us begin this section by noting that the Darwinian account of adaptations differs from Wright's original analysis of functions in one important respect—in the case of Wright's etiological account of functions, it is one and the same trait-token whose own causal effects explain its own presence in the system**,** whereas in the Darwinian case, current trait-tokens count as adaptations in case past trait-tokens of the same type had causal effects that accounted for the selection of a past trait-token-bearer (parent), which in turn accounts for the existence of the current trait-token-bearer (offspring). In short, in the Wright example, counting as a function is linked to a reflexive relation, according to which the causal powers of an entity have certain effects upon the same entity. In the Darwinian example, counting as an adaptation is linked to a relation between the causal powers of one entity and their effects on a different entity. It turns out that this seemingly minor difference poses a grave difficulty for Millikan's project.[82]

Why is that? There are a number of different facets to this problem, but I believe the fundamental issue underlying all of them is this. The Darwinian account of

[82] In a sense, it also poses a problem for Wright, since without the separation between tokens his analysis is subject to counterexamples (see Boorse, 2002; Cummins, 1975; Nissen, 1997); however, that is of no direct concern to us here.

adaptation depends upon a pattern of explanation that is peculiar from a causal point of view. The basic explanatory pattern is as follows. We begin with the counterfactual:

> If trait-token X_i possessed by parent-organism Z_i had not allowed Z_i to φ in the past, then offspring-organism Z_j would not now possess trait-token X_j (because it would not exist).

This is unexceptionable. So far, so good. However, from this we then infer that:

> Z_i's having φ-ed in the past by virtue of possessing X_i explains Z_j's ability to φ now by virtue of possessing X_j, where X_j counts as an adaptation.

Why is this peculiar from a causal point of view? Because, in the physical sciences, we do not usually say that some current entity Z has the power to φ because a previously existing entity had this power. Rather, we say that Z has the power to φ because it possesses a token-property X, which endows Z with the power to φ by virtue of Z's microstructure. To put matters in the way that the Darwinian explanation of adaptation does is to confuse the microstructural explanation of an entity's causal powers with the historical explanation of how an entity with that particular microstructure came to exist in the first place—which is a very different matter.

Whatever else they may be, biological systems are also physical systems. Therefore, it is not obvious why anyone should expect an appeal to an ancestor's properties to provide anything like a satisfactory explanation of a current organism's properties and causal powers, be they apparently teleological or otherwise. Rather, one should expect that the causal powers of a given biological system should arise from that system's particular microstructure, as with any other kind of physical system. History may, of course, have a role to play, but it will be the more limited one of explaining how a particular token-system has come to have the microstructure that it has. It is the

microstructure itself, once in place, that will be the basis for the explanation of the token-system's current causal powers.

Of course, it is not as though this distinction has been entirely overlooked until now. On the contrary, it was recognized and codified with respect to biology long ago in Mayr's (1988) well-known distinction between "proximate" and "ultimate" explanations. However, it seems to me that the philosophical implications of this distinction have not been fully recognized. In particular, in Sober's account of adaptation, it appears that the "ultimate" (i.e., selectionist) explanation is being called on to provide the sort of explanation that only a "proximate" (i.e., microphysical) explanation can legitimately give. At any rate, that is what I will be arguing in this section.

The problem of the confusion between microstructural and historical explanations in science, then, is a difficulty with selection-style explanations quite generally (McLaughlin, 2001). There is an awareness of a conceptual difficulty here on the part of mainstream philosophers of biology, who are careful to distinguish between "adaptations" with and without "current utility" (Sober 2000; 85). As Sober also puts it (ibid.): "a trait can be an adaptation now without currently being adaptive." The terminology here is confusing, but the point at issue is clear. Sober is tacitly acknowledging that history per se does not provide us with a full account of "utility," in the sense in which a particular trait may contribute to the functional coordination, and hence the survival and reproduction, of an organism. By distinguishing between

"adaptation" and "current utility," Sober and other selectionists are in no position to identify "being useful to the organism" with "having been selected."[83]

Of course, there are a couple of ways for the selectionist to respond to this point. One way would be to say that "utility" is to be understood in terms of "adaptedness," i.e., "fitness," and will ultimately be reduced in terms of future selection.[84] But as we have already seen, this is question-begging with respect to the origin of teleology, because it is arguable that it is the functional coordination of the organism that explains survival and reproduction, not the other way around. Another way to go would be to directly identify "viability" as the concept in terms of which both "utility" and "adaptedness" are to be cashed out. Indeed, this seems to be the most plausible analysis of our concept of biological function (Wouters, 1995), and arguably of the concepts of "adaptedness" and "fitness," as well, but of course once again it is hard to see how this idea can help with project of teleoreduction unless one can point to an independent means of reducing the concept of "viability." The selectionist could then gesture back towards molecular biology, but we have already seen that that would be of no help, either, as molecular biology too presupposes the teleological character of life, and thus cannot help to reduce it.

This large conceptual difficulty lying at the very heart of the teleoreduction enterprise, as conceived of by mainstream selection theory, is frequently highlighted by means of the "Swampman" thought experiment (Davidson, 1987). Swampman is an

[83] For further discussion, see Brandon (1990).

[84] This would involve understanding instances of "utility" in terms of such notions as "preadaptation," "exaptation," and the like.

exact replica of Donald Davidson, which was instantaneously created in a swamp through a fortuitous concurrence of elementary particles in a cosmic coincidence. The idea is that since on the Wright-Millikan analysis sheer dispositions become normative functions only by virtue of their selection history, an instantaneously created being cannot have functions. Thus, Swampman's behavioral dispositions, though identical to Davidson's, are in fact non-normative and non-functional, despite all appearances. To many (e.g., Boorse, 1976), the Swampman scenario has seemed like a sufficient basis for rejecting the Wright-Millikan analysis, but even so the latter's defenders have generally bitten the bullet and accepted the scenario's counterintuitive implications. Thus, Millikan (1996; p. 110) claims that "if [Swampman's] brain makes good cannibal soup, that is as much its function as thinking". Admittedly, some proponents of the so-called teleosemantic approach in philosophy of mind have been less complacent about this problem. For example, in the following passage, Dretske's (2000; pp. 256-257) intellectual discomfort is palpable: "All I can do is ask what else, besides historical antecedents, could ground a semantic connection. If nothing else could, then intentional concepts are history-laden. Swampman is a zombie. If something else could, what is it?"[85]

> Fodor (2000) has put his finger on the crux of the problem here in noting that:
>
> . . . my heart's function has less to do with its evolutionary origins than with the current truth of such counterfactuals as that if it were to stop pumping my blood, I'd be dead. (ibid.; pp. 86–87)

[85] In his earlier work, Dretske (1981, 1986) was far more sanguine about the possibility of reducing intentional concepts using a special "indication" relation explicated in terms of information theory. However, more recently he appears to acknowledge that any reduction of intentionality must ultimately depend upon a reductive account of biological function, and that "only etiological functions, functions a thing has in virtue of its history, are up to this task" (Dretske, 2006; p. 73). Nevertheless, the passage cited in the main text reveals his continuing misgivings with this approach, as well.

In other words, functions are essentially modal, not historical, concepts (McLaughlin, 2009; Nanay, 2010). And, as McLaughlin (2001) has aptly observed:

> An individual function bearer (token) is, even according to natural selection, not there because of what it itself does but because of what other things like it once did. This is the rational core of a dispositional argument against the etiological interpretation of function ascriptions as causal explanations of the presence of the function bearer. Juggling types and tokens won't solve this problem. (ibid.; p. 163)

Of course, if juggling types and tokens fails properly to account for the normativity of functions, the conventional dispositional analysis does not do any better, as we saw in section 3.2, above. Where does that leave us? In Chapter 4, below, we will look at some contemporary ideas about how a different sort of microstructural analysis might do the job. But for the present, what is certain is that if the teleological and normative character of biological function is capable of being naturalized at all, it will not be along the lines of the Wright-Millikan analysis. The heart of the mystery is adaptivity, which is a real causal power of living systems. We have every reason to believe, therefore, that the explanation of adaptivity must lie in that aspect of the microstructure of the living state of matter that gives rise to that power.

There is also another way of looking at the problem with selectionist approaches to grounding normativity: They confuse the ontic and the epistemic issues. What they provide is, at best, a convention for speaking of certain effects as functions. But they do not give us a way of understanding what it is that actually differentiates functions from other effects at the ontic level. Let me try to clarify this claim with a concrete example. Consider a batch of some material that happens to be in the liquid state (a "melt"). Any material will do, but for the sake of definiteness, let us say it is a melt of silicon dioxide (SiO_2). If we lower the temperature of the melt according to a particular regime of

pressure and rate of cooling (the details of which will depend on the nature of the melt), then we will get a solid substance with one set of properties—say, low thermal conductivity and optical isotropy (transparency) (Zallen, 1983; p. 25). If we lower the temperature of the selfsame melt via a different (slower) cooling rate, we will get a solid substance with another set of properties (higher thermal conductivity and optical anisotropy). We say that the first solid is a "glass" (in this case, common windowpane glass), while the second is a "crystal" (in this case, quartz). Now, the question is, Do the window glass and the quartz crystal owe their different properties to their different histories?

In one sense, they do, while in another sense, they don't. It is true that the two melts have the different microstructures that they do because of their different histories, if by "history" we just mean a shorthand way of referring to the very different sets of physical interactions they have undergone. In this sense, anything in the world has the physical makeup that it does because of its history. But in another sense, it is obviously not the difference in the histories per se that explains the difference in the properties of the two melts. Rather, it is the difference in their material constitutions—one is amorphous, while the other is crystalline.

The mistake of the natural selection theory of normative functions is to imagine that there is some causal power inherent in the notion of "history", above and beyond the physical interactions themselves that this word stands for. I believe that this mistake is due, at least in part, to imagining that "selection history" could confer normative value on a biological function in the same way that pedigree confers value on a horse, or provenance on a painting. But it is human beings, within the context of the institutions of

horse racing and the art market, who confer value externally on these kinds of histories as such, quite apart from the intrinsic properties of the entities whose histories they are. Pedigrees and provenances are artifactual in this respect. However, this external type of ascription of normativity makes no sense in the context of biological function. To ascribe value to selection histories as such in the same way that we ascribe value to pedigrees and provenances as such commits one of two different sins, depending upon one's interpretation. If we interpret this move ontically, then it is tantamount to saying that the selection process is capable of imposing extrinsic normative criteria in the same way that human beings do. In this case, it reifies history and anthropomorphizes natural selection. To think in this way flies in the face of everything we know about the way the natural world operates, and is fundamentally anti-naturalistic. If, on the other hand, we interpret the move epistemically, then the entire Wright-Millikan analysis collapses into a linguistic stipulation. This is the way that we will agree to talk about functions (which presumably then do not really exist at all). Such a move may interest analytical philosophers, but it is of no interest to anyone who wishes to understand what functions really are.

Now, it may be objected that solidification and evolution are two very different sorts of processes, and in many respects, no doubt they are. But in the relevant respect, I believe the analogy holds. The point about history is quite general, and may be illustrated in any number of other ways—for example, with respect to ontogenetic learning. Say I have a choice between studying French and studying German, and I opt for French. Then, at the end of many years' instruction I may emerge with a knowledge of one language instead of the other. I will have traversed a different life history than I

would have done had I opted for German instead—I will have sat in a different set of classrooms and listened to a different set of instructors making a different set of sounds. So, there is a sense in which one might say that my particular life history explains why I know French, but am ignorant of German. However, if we focus on my current linguistic competence, we can see that that is not really explained by my life history per se. Rather, my current linguistic competence is explained by the current powers and dispositions of my brain. How those dispositions and powers came into being is one thing, how they are now constituted is something else. When examined more closely, my personal history really refers to the dynamical evolution of certain neural networks in my brain. This evolution resulted in the formation of one set of attractors rather than another, corresponding to this set of sounds rather than that. But my current competence is explained by the dynamical state resulting from this process, and not the process per se. At the end of the day, my knowledge of French still reposes upon the current dispositions of my brain.

I believe the case is precisely the same with phylogenetic learning. While "selection history" may be a perfectly acceptable shorthand way of adverting to the fact that particular lineages of organisms traverse particular paths through morphospace, at the end of the day it is still the case that the properties and powers of organisms derive from their physical constitutions, not from their selection histories per se.

In sum, "selection history" may be a perfectly acceptable shorthand way of adverting to the fact that particular lineages of organisms traverse particular paths through morphospace; as such, it may be a convenient way of speaking about the immensely complex nexus of dynamical, microstructural processes that led to the origin

of a particular type of organism. But it is by no means an explanation of the current causal powers of that organism. To say otherwise would be like pointing to the Hertzsprung-Russell diagram to explain why gold is yellow and shiny. We summarize the causal history responsible for the origin of gold in the H-R diagram's graphic summary of stellar evolution, but it is the laws of quantum mechanics that explain why this particular wedding band is yellow and shiny, not that history as such. Similarly, the "selection history" of a particular organism can do no more than point to the complex causal process that gave rise to the current microstructure of the organism. It is that current microstructure itself that explains the organism's causal powers, including the property, whatever it is, that makes some of the causal effects within the organism, but not others, perform the role of functions. At the end of the day it is still the case that the properties and powers of organisms derive from their physical constitutions—their microstructures—not from their selection histories as such.

In sum, the fundamental problem with the Wright-Milllikan account of biological function is that it confuses the reason a thing has the properties that it does (its microstructure) with the reason a thing has the microstructure that it has (its history). Functions, realistically interpreted, are properties of organisms. Therefore, microstructure is relevant to their explanation, but history per se is not. The moral of this lesson has been expressed by Callebaut and coworkers (2006; p. 42) in a particularly clear and concise way: ". . . biological research should substitute for past causes the 'traces'—state variables—left in the present by the operation of those causes" [references omitted].

Of course, this counsel is considerably easier to formulate than it is to put into practice, as we shall see in the next and last chapter. First, though, let us return to a

question that was raised above in a preliminary fashion: namely, the empirical evidence that calls into question the mechanistic assumption underlying the Positive View of the creative power of natural selection, and in so doing provides further warrant for the Mixed Negative View, which in turn opens up conceptual space for **TRB**.

3.5 Argument from the Adaptive Capacity of Living Systems

The empirical evidence I have in mind comprises various cases exhibiting radical adaptive capacities that cannot be plausibly accounted for using the conventional selectionist pattern of explanation, because the capacities in question cannot plausibly be held to have ever contributed to the fitness of any ancestor of the organisms in question. I will briefly mention three empirical studies that support this point directly, and then supplement them with a fourth set of studies that seem to warrant our generalizing the existence of a general, intrinsic adaptive capacity in organisms beyond the class of animals with brains to that of living things as such.

During World War II, a Dutch zoologist published an anatomical study (Slijper, 1942) of a goat born without forelimbs, which learned to hop bipedally—i.e., on its hind legs—as its regular mode of locomotion. Upon dissection, it was found that much of the animal's skeleton and musculature had been radically remodeled, rather along the lines of those of a kangaroo.[86]

It has been known for more than a century that human beings subjected to an inverted visual field will gradually adapt quite successfully to this anomaly (Stratton,

[86] For discussion, see West-Eberhard (2003, 2005). Similar cases of bipedal dogs and pigs may be viewed on YouTube. For instance, the animal called "Faith the Dog" has a strikingly human-like walking mode of locomotion.

1897). Heisenberg & Wolf (1984; pp. 194–204) have demonstrated an analogous capacity in fruit flies.[87]

Likewise, it has been known for several decades that blind human subjects are able to form visual images on the basis of sensory input via the skin or tongue (Bach-y-Rita, 2004; Ptito et al., 2005), while more recently Sur and colleagues (Sharma et al., 2000; von Melchner et al., 2000) have demonstrated that newborn ferrets whose optic nerves have been surgically redirected to the auditory cortex eventually learn to see well enough to support much of their normal visually guided purposive behavior.[88]

In each of these cases, we are presented with striking evidence of an inherent power of adaptivity (brain and general physiological plasticity) in living things, which cannot be plausibly explained by appeal to the theory of natural selection, since none of the behaviors in question can have been manifested at all, much less have contributed to fitness, in any ancestor of the experimental animals. Rather, the animals seem to possess a latent power of adaptivity far beyond anything that could be predicted on the basis on selection theory. These experiments strongly suggest the existence of an adaptive power that is an intrinsic property of all living systems.

But what of single cells? Might it not at least be argued that the cases mentioned above are crucially dependent upon the fact that behaviors in question are all mediated by brains, however primitive? In other words, do such cases demonstrate the existence of an adaptive power inherent in life as such, or only in brains? A number of experiments

[87] See Heisenberg et al. (2001) for further examples of behavioral plasticity in Drosophila.

[88] For other similar experiments, and discussion, see Newton & Sur (2005). For discussion of the molecular "mechanisms" underlying neural plasticity, including evidence of homeostasis at the level of individual neurons maintained by endogenous activity, see Tropea et al. (2009).

strongly suggest that similar powers of adaptivity are present in even the simplest living systems.

For example, Nakagaki and colleagues have shown that the plasmodium of the slime mold, Physarum polycephalum (i.e., an amoeboid protozoon), has such capacities as that of finding the shortest route through a maze (Nakagaki et al., 2000), of organizing itself into a maximally efficient hub-and-spoke system (Tero et al., 2010), and of anticipating oscillatory events of arbitrary periodicity (Saigusa et al., 2008).[89] Now, these capacities are arguably less far-removed from Physarum's natural ecological needs, and so are more amenable to explanation via the standard Darwinian scheme. However, note that these capacities seem to be very similar to the adaptive capacities which in the other experimental organisms mentioned are mediated by brains.

Even more striking are so-called "knockout" experiments, in which genes are inactivated through recombinant DNA technology—typically in early embryos of mice—with the result that the mature animal lacks a particular type of enzyme. To the astonishment of experimenters, the adult mice very often showed no ill effects, even when the enzymes supposedly "knocked out" seemed to be crucial for normal functioning (see Strand and Oftedal, 2009). One might suppose that the reason is simple redundancy—that is, the existence of ordinarily superfluous back-up enzymes that step in to do the job when the normal ones are eliminated—but it has been shown that this is not the case. Rather, these organisms exhibit what A. Wagner (2005) calls "distributed robustness." As Wagner puts it, ". . . absence of phenotypic effect . . . camouflages profound systemic changes that have to take place to compensate . . . " (ibid.; p. 131).

[89] For further details, see also Nakagaki et al. (2009) and Tero et al. (2005).

That is, rather than a simple substitution effect, what happens is that "[w]hen one part fails or is changed through mutations, other parts can compensate for this failure, but not simply by standing in for the failed part" (ibid.; p. 239). Although knockout experiments have mainly been carried out in mice, it is clear that we are dealing here with a generic capacity of living systems that has nothing to do with brains.

Finally, there is also a well-known class of experiments in which metazoan organ tissues (heart, kidney, etc.) and even entire primitive animals (e.g., sponges, sea urchins) are dissociated into their individual constituent cells, which (given proper care) have the power of spontaneously reaggregating into their original functional form, or something close to it.[90] Perhaps the most striking such experiment—and one more directly pertinent to the present discussion—has to do with the protist Euglena. This single-celled creature has been subjected to centrifugation such that many if not most of its constituent parts (small molecules and macromolecular assemblies alike) stratify into separate bands or layers within the animal's outer membrane. Then, after some time, these internal parts find their way back to their original locations, and normal functioning is restored (Kempner & Miller, 2003). Such experiments clearly show that there is no need to restrict the domain of the inherent power of adaptivity to animals with brains. Rather, this capacity is demonstrably present in lower forms of life, as well.

Of course, the selection-theorist could always counter such examples by explaining such capacities and powers as "spandrels"—free-riding properties accompanying adaptive properties that are actually selected for. Or else, one might speak

[90] See Moscona (1959). For a more recent, albeit popular, discussion, see Rensberger (1996).

of "hidden norms of reaction" or "cryptic variability" put into place by previous rounds of selection. However, it seems strange to posit a view in which a universal adaptive capacity is brought into being more or less incidentally to the process of natural selection—that is, a view in which the adaptive capacities of the supposed spandrel far outstrip those of any putative associated selected-for properties. That seems to be purchasing a power as remarkable as adaptivity at far too cheap a price. Moreover, the very existence of such a universal adaptive capacity, however it was put into place, is really all that is required for present purposes, which I remind the reader once again is nothing more than to cast doubt upon the claim that teleology has already been successfully reduced to mechanism.

3.6 Two Objections

The above considerations seem to indicate, not that the theory of natural selection has no substantive role to play in the overall explanation of the evolution of living forms, but rather merely that it does not provide us with an adequate framework for reducing the apparently teleological properties manifest in living things to mechanical causes alone.

Nevertheless, proponents of the mainstream anti-realist viewpoint might of course advance various objections against the teleological-realist perspective. I will end this chapter by examining two of the most considerable of these.[91]

[91] I suspect that most objections against the teleological-realist view will ultimately be found to contain, in their logical core, either the intuitive idea that the mechanistic details ought to reduce the appearance of teleology in biology, or else the idea that the theory of natural selection ought to do so. If that is right, then these two objections should be representative of a great many more.

3.6.1 Objection from the Conceivability of Artificial Organisms[92]

The first type of objection challenges the teleological-realist view on the grounds that building an artificial organism out of mechanical parts seems to many to be clearly conceivable. Of course, we are very far from such a capability at present, but that is nothing to the point. Many people are of the opinion that such an achievement is perfectly conceivable—at least, it is not obvious that it implies any contradiction. And if, as many are also convinced, conceivability implies real possibility, then it seems as though it ought to be possible to build an organism someday from scratch using inorganic materials. If this reasoning is correct, then teleological realism must be wrong. If the reasoning is not correct, then it is incumbent upon the teleological realist to explain where it goes wrong. This would involve explaining why future improvements in our present engineering capabilities must be limited in some way.

There are two sorts of responses to this objection.

First, it must be acknowledged that teleological realism is at bottom an empirical conjecture, and so inevitably is hostage to the future course of science and technology. Nevertheless, there are good reasons for placing one's bets against the possibility of building an organism from scratch. In a case like this, conceivability means little, because it is not logical possibility but nomological possibility that is at issue. The question is whether building an organism from scratch out of inorganic materials violates some presently unknown law of nature. To insist on the intuition that this is really possible is to

[92] Thanks to Grant Ramsey for pressing me on this point.

beg that very question. What we must do, instead, is adduce evidence for one side or the other.

The main evidence in favor of the real possibility of constructing an artificial organism is the apparent progress already made toward engineering bits and pieces of living systems (artificial tissues, organs, sensory apparatus, etc.), together with the difficulty in making out why this progress must be inherently limited. The main evidence against the real possibility of constructing an artificial organism is the great difficulty encountered so far in engineering more than small bits and pieces of living things. In all cases so far, we remain reliant upon the inherent adaptive powers of the living system within which the artificial piece is embedded to "capture" the piece and incorporate it into the ongoing life of the organism in question. In this respect, it is one thing to engineer a cochlear implant, or even a neural implant enabling someone to drive an automobile by thinking alone (as was recently accomplished), and it is something else again to engineer an entire living organ. All that has been accomplished so far is the engineering of devices that piggyback upon the inherent capacities of a living system. Engineering a device with inherent capacities capable of replacing a living system and standing on its own, but with functionality similar to the living system's, is a very different proposition. Impressive as they are in many ways, accomplishments in this field up to the present have actually been rather modest. Even a practical artificial heart—an organ that is supposedly nothing but a mere pump—remains far beyond our present technical know-how. If present trends are any indication of the likely future course of science and technology, it would appear that we are much more likely to grow tissues and organs from stem cell cultures in the future than we are to engineer such systems from

scratch. But even if this correct, it only raises the question of the reason why such engineering feats are so difficult. Here, of course, I am in no position to be dogmatic. To insist that an artificial organism is nomologically impossible because of past difficulties would be as wrong-headed as to insist that it is nomologically possible because of its apparent conceivability. Rather, we need to try to probe deeper, to get at the more fundamental underlying issues.

This, then, is the other sort of response to the objection from the apparent conceivability of artificial organisms. Let us grant for the sake of argument that it will someday be possible to build practical tissues and organs—say, a workable artificial heart—out of wholly inorganic materials (titanium, Dacron, silicon chips, etc.). Let us go even further and assume that it will be possible to build an entire artificial cell out of such materials—say, an artificial "bacterium." This means that we are envisioning the engineering of an artificial system that is capable of doing all the things that living bacteria do, from "metabolism" (self-sustaining manufacture of needed materials) and "chemotaxis" (locomotion away from dangers and toward needed materials), "replication" (manufacture of other such systems), and so forth. Still, even in that case, it would still remain open to question whether the nomological possibility of the artificial cell was ipso facto proof that the natural living cell operated according to the same physical principles as the artificial cell, and thus that the thesis of teleological realism was false. What do I mean by this?

If conceivability is our criterion, then it is perfectly conceivable both that an artificial cell might be nomologically possible and that the thesis of teleological realism might be true, at least in relation to natural living things. How so? Because from the point

of view of teleological realism, the two cases would remain quite distinct. The distinction would have to do with the source of the functional coordination of the parts of the system. In the case of the artificial cell, there would be no reason to believe that there was any inherent adaptive capacity lurking within the material components of the system. All of the functional coordination of the titanium and Dacron parts would have to be supplied by the engineers themselves, from the outside as it were, as a set of boundary conditions, precisely as an automobile is assembled. There would be no more inherent tendency for the individual parts of the artificial cell to metabolize, locomote, and replicate than there would be for the individual parts of an automobile to fall together and speed off on their own. It is precisely because of the unfathomable complexity of such an engineering task that its ever being realized strains credulity.

The case of the living cell is very different, from the teleological realist perspective. Here, by hypothesis, the functional coordination of the living cell arises spontaneously from within, out of the active dynamics inherent to the living state of matter, rather than being imposed on inherently inert parts from without. If this idea seems difficult to accept, it is only because we have at present little idea of how such an inherently active dynamics might be physically constituted. It will be the burden of the next and final chapter of this dissertation to attempt to dispel some of the aura of mystery that may be felt to surround the very notion of a sui generis active dynamics inherent to living matter. But while such empirical inquiries are required in order for teleological realism to carry much conviction, they are not really necessary to rebut the particular objection from the conceivability of artificial organisms. For, that objection purports to find in their conceivability sufficient grounds for rejecting the possibility of teleological

realism. But that is a plain mistake. For, all that is necessary for us to entertain the real possibility of teleological realism is to show that that thesis is itself conceivable. And in the last chapter I have already done more than that, showing not only that the thesis of teleological realism implies no contradiction, but that in fact we have good reasons for believing it to be true.

3.6.2 Objection from Population Biology

Finally, one of the most devastating criticisms of the teleological-realist view, if could be made to stick, is the objection from population biology. The basic idea here is that the teleological-realist viewpoint, with its emphasis on individuals and their causal powers, is inherently incapable of making certain necessary discriminations with respect to adaptive traits which come into focus only at the level of populations. Let us see in detail how this objection is supposed to work.

Consider a population of organisms, most of whose members possess a pair of complementary traits, say "strong" and "bold."[93] For example, a carnivore might have well-developed fangs and claws, together with an aggressive personality making it prone to attack its prey or sexual rivals quite fearlessly. Now, imagine that an individual within this population is born with a mutation resulting in the suppression of only one of the two complementary traits—say, the organism is strong but timid. Such a mutation, then, will result in a creature that looks like a lion but behaves like a mouse. Call this the "cowardly lion" scenario. The cowardly lion is supposed to constitute an objection to the

[93] The example is due to Grant Ramsey, whom I thank for pressing the objection upon me.

teleological-realist view by virtue of the following consideration. On the Darwinian view, it is obvious why the creature has well-developed fangs and claws (because it inherited them), while on the teleological-realist view, it is not clear why such an ill-assorted creature should exist. The general objection is that a population perspective is necessary in order to correctly identify adaptive traits. From the teleological-realist perspective, adaptive traits are apt to be misidentified or missed altogether.

The teleological realist may make several replies to this type of objection.

First, he must concede that if a population-level phenomenon (such as the frequency of distribution of a trait) is under consideration, then of course one must look to the population level for the explanation of the phenomenon. As an individual-organism –oriented perspective, teleological realism does not claim to replace population thinking in toto. It merely claims that considerations of the causal powers of individuals are logically prior to population considerations. That is to say, teleological realism does not claim to replace population thinking; rather, it challenges population biology's claim to have replaced essentialist thinking with respect to the causal powers of individual organisms.

Another way of putting the teleological realist's response to the objection from population biology is this. The objection from population biology only tells against the teleological-realist view if we assume with the population thinker that adaptations are ultimately population-level phenomena, inexplicable in essentialist or individualistic terms. But the truth of that assumption is, of course, the very point at issue between the population thinker and the teleological realist. Moreover, not only would such an

assumption be question-begging in this context, it would also be positively ill-founded, as may be seen by a little reflection upon the very case under consideration.

If one focuses on the fact that the cowardly lion example represents, on balance, a maladaptive phenotype (at least in the ecological context of the African savanna), then it becomes clear that we are indeed using an individualistic viability criterion and not a population-level criterion to make such a judgment. Our judgment takes this form: "This particular organism with this particular ensemble of attributes is likely not to flourish in this particular ecological context." This judgment depends in no way on knowing anything about the frequency of distribution of traits of the population to which the cowardly lion belongs. On the contrary, it is a paradigm case highlighting the fact that our judgments of which phenotypes are adaptive and which are maladaptive have nothing whatever to do with population thinking, but rather are ultimately grounded in the concept of the viability of the individual organism (Wouters, 1995).

But wait, wasn't the point supposed to be the fact that the population thinker was in a position to explain why the cowardly lion has those big fangs and claws that it is not disposed to use, while the teleological realist supposedly had no conceptual resources capable of either explaining the presence of those fangs and claws, or of even recognizing them as adaptations? Here, the teleological realist must make a twofold reply.

First, to deny that history per se can explain the particular causal powers or capacities that systems possess is not to deny that history can explain the presence of those powers and those capacities in a particular system, as opposed to their absence. In terms of our analogy with stellar evolution, take the case of a massive star that is at the

stage of fusing oxygen to create silicon. With respect to the capacity of the star to produce silicon, there is a difference between saying that the presence of oxygen may be explained by the past history of the star, and saying that the past history of the star per se explains the causal process by means of which silicon is produced. The first statement is true, but in no way tells against the teleological-realist view of causal powers inherent in living things. The second statement is misleading at best. Similarly, the presence of big fangs and claws in the cowardly lion may be explained by an appeal to history.[94] However, the fact that an organism with a mismatched pair of complementary traits has reduced viability is a fact grounded, not in the organism's history, but in its present capacities and causal powers.

Second, the fact is that for the cowardly lion, the big fangs and the big claws are not necessarily adaptive at all, as the population thinker would wish to claim. They may in fact be quite maladaptive. Of course, the population thinker will wish to insist that it is the cowardly lion's timid temperament that is maladaptive, but on what logical basis? After all, for all we know, the cowardly lion might be on the way to losing its fangs and claws, as well, and evolving into an herbivore, with traits more similar overall to those of an antelope. Obviously, if one of a pair of complementary traits is lost, there are two ways to restore adaptive equilibrium with respect to the pair: restore the lost trait, or vary the remaining original trait to match the changed trait. It seems wholly arbitrary to say

[94] This formulation actually concedes too much to the mainstream view, because as always "history" is really nothing more than a short-hand way of referring to a temporal sequence of dynamical events. Thus, even the presence of oxygen in an oxygen-fusing star is not really explained by "history" per se, but rather by the fact that in a previous cycle of fusion the star in question was burning neon to form oxygen. Similarly, the presence of big fangs and claws in the cowardly lion is, of course, really explained by a sequence of complex molecular and physiological events within the reproductive and developmental processes that brought the lion into being, and not by its "history" as such. But this point can be set aside for now, as even the weaker claim that history can account for the "presence" of a capacity but not the causal powers inherent in the capacity is enough to rebut the objection from population biology.

that one of these outcomes would be "adaptive," while the other would not be. But if both of these outcomes must be counted as "adaptive," as surely they must be even on the population view, then what grounds do we have for saying of the cowardly lion that it is its disposition that is maladaptive, and not its overgrown fangs and claws? In any event, whatever we are pleased to say about the cowardly lion's mismatched traits, it is obvious that it is only by evaluating the overall coordination of the animal as a functionally integrated individual that we are able to assess its individual traits as adaptive or not adaptive in the first place. Neither the animal's place within its population nor its phylogenetic history has any bearing at all on this assessment.

As Bouchard (2011; p. 111) has recently remarked: "Population thinking as it emerged in contemporary evolutionary thinking was intended as a way of abstracting away from individual circumstances in order to track only the mathematical properties of populations." That is why the traditional mode of "population thinking" is so misleading when it comes to trying to understand the general capacities of living things, including adaptivity. In such cases, the mathematical abstraction of population biology drops from view the very thing that most requires explaining. For, like any other real property of a real entity, adaptivity is a capacity of individuals, in this case organisms. Reiss (2009; p. 22) has recently summarized this point nicely by noting that "[a]daptedness is not a product of evolution; it is a condition for evolution."

As it turns out, the mathematical formalisms of population biology are applicable mainly to sexually reproducing metazoans; they are much less useful for other types of

organisms (Bouchard, 2011).[95] Therefore, while they are perfectly appropriate in their place, they are of less relevance to the sorts of foundational issues that are the subject of the present investigation, which finds its "model organism" in the bacterium. Moreover, the level of analysis of population biology is "phenomenological" (in the scientific sense of that term), in that it treats the individual organism as a "black box." But the problem of the viability of individual organisms is just as important to a complete understanding of life and evolution as the problem of the stability of matter is to a complete understanding of the behavior of a gas in a container. Natural selection without a deeper understanding of the inherent, adaptive compensatory capacity of all living things is as conceptually incomplete as the kinetic theory of gases would be without a deeper understanding of the quantum mechanical basis of the existence of atoms and molecules. In short, there is reason to believe that the theory of natural selection may eventually come to be seen as a special limiting case of a deeper theory of the dynamics of the living state of matter. Some avenues by means of which such a deeper theory may perhaps be fruitfully pursued will be explored next, in Chapter 4.

All of these difficulties with natural selection conceived of as the foundation of the teleoreduction project are of course well-known. Relatively little that I have said in this chapter is new. The difficulties are, accordingly, beginning to be widely recognized, and a great deal of discussion is now going on about the need to "extend" the Modern Synthesis (e.g., Pigliucci & Müller, 2010). Perhaps the most prominent proponent of such an "extended synthesis" is Pigliucci (2009), who argues that the Modern Synthesis requires emendation in various ways. Among the new ideas that need to be taken into

[95] See, also, Dupré (2010) and O'Malley & Dupré (2007).

account, Pigliucci (ibid.; p. 226) enumerates the following: "evo-devo" (i.e., a synthesis of evolutionary and developmental biology), network theory, epigenetic inheritance, complexity theory, niche construction, and the concepts of plasticity and accommodation, to name a few. Pigliucci is even moved to remark that:

> . . . living organisms are complex developing systems, not at all analogous to human-made machines (despite the popularity of the latter metaphor). As such, living cells, tissues, and tissue systems are endowed with the ability to react systematically, and often adaptively, to changes in the environment—both in the classic sense of the external environment and in the sense of internal, genetic, and developmental environments. [emphasis added] (ibid.; p. 224)

In spite of these remarks, Pigliucci is convinced that the conceptual core of Darwinism remains intact, and that no "paradigm shift" is required to accommodate Darwin's original insights to the new information now pouring in from laboratories around the world.

Whether the conceptual changes required to "extend" the Modern Synthesis in such a way as to achieve a better understanding of the foundational conceptual issues in biology amounts to a "paradigm shift" is not something I am concerned with in this dissertation. I have no stake in what is basically a semantic issue, and take no side on the question of whether all that is required is an "Extended Synthesis" or whether a more radical and more fundamental change of viewpoint is needed. What I am concerned with, however, is trying to achieve a better understanding of the inherent capacity of organisms to compensate for perturbations in a way that is consistent with their continued persistence, or viability. Or, in a word, what I am concerned to understand is the inherent and universal capacity of organisms that we have been calling "adaptivity." Accordingly, I will turn to a direct investigation into the sources of adpativity in the next chapter.

What Pigliucci's own words surely do help to confirm, at any rate, is that teleology has not yet been successfully reduced to mechanism, and that is what I set out to show in this chapter.

CHAPTER 4:

WHAT MIGHT AN ORGANISM BE, IF NOT A MACHINE?[96]

4.1 Introduction

In Chapter 2, I have attempted to show that there are weighty considerations, mainly of a conceptual nature, in favor of regarding teleology in biology as objectively real and organisms as genuinely normative agents. In Chapter 3, I have argued that claims that teleology has already been successfully eliminated from biology via molecular biology and the theory of natural selection may be legitimately questioned. However, even if someone found my arguments up to this point to be persuasive, it would still be only natural for him to demand to know what positive account of biological teleology—and of normative agency—I have to offer. The underlying worry would be that no genuinely scientific account of teleology and normativity is even conceivable, and that the seeming incompatibility between my conclusions and the precepts of biological science as it is actually practiced is by itself reason enough to justify rejecting all of the foregoing out of hand. In this chapter, I will address this understandable and important concern, by showing that, although we currently lack any well-established theoretical framework that would clearly support the teleological-realist view of organisms as normative agents, nevertheless, there are positive reasons for believing that such a view is at least conceivable, and for this reason—and in light of the considerations discussed in

[96] Van Gelder (1995) inspired this way of putting my problem.

Chapters 2 and 3, above—that it deserves to be considered as a "live option" in the ongoing philosophical debate on the nature of life, teleology, normativity, and agency.

Another way of formulating the goal of this chapter is reflected in its title. I have already shown in Chapter 3 that even a sophisticated type of manmade machine whose operation is governed by a complex cybernetic-control mechanism cannot properly be said to possess inherent (or original or underived) normativity, and so cannot properly be accounted a normative agent. It is, of course, commonly supposed that organisms—or, at least, appetent organisms such as bacteria—are at bottom nothing but very sophisticated machines. Therefore, any positive argument in favor of **TRB** owes the reader some account of how organisms differ from machines with respect to their capacity to possess inherent normativity, and hence genuine agency. In other words, while I have already explained what organisms are not—namely, machines—it would be highly desirable if I were now able to provide at least the rudiments of a positive account of what they are. In this chapter, I will attempt to make good this omission to the extent possible.

Obviously, it is not the place of philosophers to engage in speculation about matters that clearly fall within the province of empirical inquiry. So, if teleology in biology is indeed an objectively real phenomenon, as I claim, then it cannot be my role here to provide a scientific explanation of that phenomenon. Nor, of course, can I merely direct the reader's attention to some article or book that provides such an explanation. If such an explanation were already well established by the scientific community, and were widely acknowledged as such, the reader would surely have heard about it before now.

Rather, what I will do in this chapter is draw the reader's attention to a recent scientific literature that claims to point us in the right direction. I will be highly critical of

this literature, in the sense that I will show that the scientific concepts it draws upon are incapable of fully explaining normative agency. Nevertheless, it is a fact that the scientific project of explaining normative agency is now underway. If its present conceptual repertoire is inadequate to the task at hand, that does not mean that it may not develop more adequate concepts in the future. At a minimum, the very existence of this project shows that a scientific theory of normative agency is by no means inconceivable.

I will begin by considering in the following section the general question of the bearing that empirical research can be expected to have upon the issues treated of in this dissertation, and also by situating the present enterprise on the contemporary conceptual landscape as a particular variety of "naturalism." Next, in Section 4.3, I will consider the general question of whether it makes sense to think of life as having an "essence"—and if so, what its nature might be—as a way of posing as sharply as possible precisely what I take the problem of normative biological agency to be, from an empirical point of view. Then, in Section 4.4, I will argue that the most fundamental feature of life is "adaptivity,"[97] and that, from a physical point of view, one should expect the coherence and coordination of physical processes constituting adaptivity to be governed by an underlying physical principle (the "Ground Argument"). After that, in Section 4.5, I will review and critique some representative samples of a recent but rapidly growing body of work at the interface between theoretical biology and cognitive science that explicitly deploys the concepts of "normativity" and "agency." It will be found that most contributors to this literature either equivocate by effectively using "agent" in a non-

[97] It will be recalled that I introduced "adaptivity" in Chapter 3 as an umbrella term for the general capacity of all living things to take compensatory action in order to maintain viability.

normative sense, or else rely upon a notion of "self-organization" based (whether implicitly or explicitly) upon concepts borrowed from the scientific disciplines of nonlinear dynamics and nonequilibrium thermodynamics that are conceptually inadequate to serve the physical principle underlying adaptivity. I will show that while such concepts may be necessary for the naturalization of normative agency, they are not sufficient. Finally, in Section 4.6, I will review several lines of contemporary research, drawing especially on ideas borrowed from condensed-matter physics, to paint a picture of the living cell as a physical system endowed with inherent capacities that make it much more recognizably the sort of entity that might genuinely deserve the name of "normative agent." I will conclude that while this approach holds more promise than the other one of someday revealing to us what an organism might be, if not a machine, nevertheless, it too has a long way to go before that task can be brought to fruition.

4.2 A Note on Naturalism

Before moving to an examination of the various aspects of contemporary empirical research relevant to our question, it would be well for us to reflect for a moment upon the general character of the inquiry about to be undertaken in this chapter. What is it, precisely, that such an inquiry as that pursued here hopes to achieve?

I have already stated one desideratum: namely, that of independently motivating the acceptance of the conclusion of the argument in Chapter 2, which might otherwise be taken to be so paradoxical as to justify the wholesale rejection of that argument, and especially of the claim that teleological and normative language is properly and literally ascribable to organisms as such. But in addition to articulating this strategic aim, I also

need to situate the current chapter within the contemporary debate on "naturalism." The reason that this issue cannot be skirted is that my task may otherwise appear naïve or quixotic, in either of two ways. To some, the enterprise undertaken in this chapter may appear to be merely another species of reductionism, thereby undermining the whole "realistic" tendency with respect to teleology and normativity of the dissertation up to this point. To others, the present enterprise may appear excessively speculative, relying as it does upon ideas some of which have not been accepted by the mainstream of opinion in biology at the present time. To these latter, it may also seem that I am trespassing on territory that is the proper preserve of the empirical scientist. For the former sort of objector, one might say, crudely speaking, that I am being "too scientific," in the sense of dragging in empirical considerations where they do not belong, while for the latter sort of objector, one might say that I am being "not scientific enough," in the sense of arrogating to myself as a mere philosopher the right to question the scientific consensus of the day. For these reasons, I will attempt in this section to clarify precisely what I take the bearing of empirical research upon my project to be, while situating my project more generally within the conceptual landscape of contemporary debates about naturalism.

First, I would like to make it perfectly clear that I agree entirely with the following judgment (De Caro & Voltolini, 2010; p. 71): "The metaphilosophical constraint that philosophical views should not be at odds with science is both attractive and well established." The cognitive success of natural science has indeed earned it the right to a great measure of respect from everyone, including philosophers, and especially from philosophers (such as I) who wish to call themselves "realists" and even "naturalists." Now, following the "not at odds" principle would appear to condemn the

present enterprise before it is even begun. Admittedly, when it comes to human consciousness and human agency, it is at least arguable that these are phenomena with which we are acquainted first-hand, and for which there is no properly accredited natural science entitled to sufficient respect to cause us to defer to its claims to the extent of discounting our own first-person experience. But whatever one's views on those sorts of questions, on the question at the center of our inquiry here, surely (one might think) there can be no doubt that biological science has long since become a fully accredited intellectual undertaking, and so one to which as philosophers we are obliged to defer, even if it means relinquishing such seemingly unimpeachable intuitions as that biological functions serve the purposes of organisms, that things can go well or poorly for organisms, be good or bad for them, and so on. If biological science tells us that these intuitions have been demonstrated to be simply untutored prejudices with no foundation in fact, are we not obliged to pay heed to them?

Two major considerations bear on the question of the propriety of critiquing an established scientific opinion from the outside (sociologically speaking), on mainly a priori grounds (as opposed to critiquing it from within, on mainly a posteriori grounds). One is an "internal" or interpretative issue, and the other is an "external" or sociological issue. The internal or interpretative issue is this. The metaphilosophical constraint that "philosophical views should not to be at odds with science" presupposes that we are in possession of a clear criterion of what counts as "science" for this purpose—that is, it presupposes that we know precisely where the boundary lies that we as philosophers must not cross. But this is not in fact the case. We do not know how to draw such a boundary, and, indeed, one might say that the difficulty or impossibility of drawing such

a boundary is the very reason why philosophical disputes arise in relation to science in the first place. By and large, philosophers are well content to stay on their side of the boundary wherever it is clearly visible. For example, few philosophers, I think, would be interested in entering into a critique of the details of the chemical composition of the "nanobrain" controlling bacterial chemotaxis discussed in Chapter 1. Or if a philosopher with extensive training in biochemistry did embark upon such a critique, it would be qua biologist and not qua philosopher. In short, questions that can be answered by straightforward empirical means (in this case, for which well-established physical or chemical assays exist) fall clearly on the far side of the line, within the province of the biologist.

By the same token, questions that are mainly conceptual in nature—such as, for example, the nature of a species—fall clearly on the near side of the line. In such cases, not only are the contributions of philosophers accepted into the scientific discussion, I believe they are even welcomed by most scientists, who understand that their problem is mainly one of a lack of conceptual clarity, as opposed to a lack of empirical knowledge or theoretical insight.

Then there are problems that lie in those nebulous regions where the frontier is not clearly marked. In my view, the problems of teleology and normative agency are frontier problems par excellence. That philosophers are not to be regarded simply as interlopers in this region seems evident by the very existence of the voluminous and long-lived debate on the nature of functions, though I believe that this is a philosophical discussion that has had relatively little impact on the actual practice of biology. But while many biologists would perhaps be prepared to acknowledge that there is a conceptual

issue of some interest in the region of teleology—if only because of the glaring discrepancy between the anti-teleological ideology and the teleophile practice of biological science—nevertheless, I suspect that many if not most would feel their territory beginning to be infringed upon by the present project. And, indeed, the present project is a mixed conceptual-empirical one, and the present chapter above all. I would like to try now to allay such suspicions by specifying where I conceive the present project to lie within the conceptual landscape of contemporary attempts to "naturalize" teleology and normativity.

Broadly speaking, this dissertation can be viewed as a contribution to the project of "naturalizing normativity"—a project that is proceeding along a broad front of contemporary philosophy. And yet, for many philosophers the concepts "normative" and "natural" remain antithetical, and the idea of "naturalized normativity" is an oxymoron. For this reason, a few words about what the project of naturalizing normativity does and does not entail are necessary.

The project of naturalizing normativity is a highly various and complex enterprise, but perhaps it would not be oversimplifying matters too much to distinguish three main approaches. The first approach is the effort to eliminate normativity from our ontology altogether. On this view, normativity is "naturalized" by showing that it does not really exist, and that in reality the "natural" (understood here as a contrast class to the "normative") is all there is. This may be achieved, it is supposed, either by showing that the putative normative phenomena (such as actions) to which our normative concepts seem to refer can be ontologically "reduced" to nonnormative phenomena, and so are redundant, or else by showing that the putative normative phenomena do not really exist

in an objective sense, and are merely a subjective "projection" of human concepts and behavioral response patterns onto the world—i.e., a sort of "illusion." The justification for the eliminative approach may be expressed by means of something like the following argument (the "Eliminative Argument"):

> (1) The picture of the world painted for us by the present-day physical sciences (including chemistry and biology) is complete in all fundamentals. Call this the "present physical picture."
>
> (2) Our ontology—that is, our list of the things that really exist in an objective sense—ought to correspond to the present physical picture.
>
> (3) The present physical picture makes no mention of normative phenomena.
>
> (4) Therefore, normative phenomena do not really exist in an objective sense, and ought to be eliminated from our ontology.

Now, this simple picture would have to be complicated in numerous ways if a faithful account of the state of play in the literature were our goal here. For one thing, it would have to be acknowledged that there are relatively few philosophers who explicitly embrace eliminativism (e.g., Churchland, 2007; P.S. Davies, 2009). This should not be surprising, since to deny flat-out that normativity exists is a very strong and highly counterintuitive claim. But it does mean that the many philosophers who subscribe to one form or another of "reductionism" owe us a clear explanation of exactly what they take the ontological status of the "reduced" higher-level entities to be. To see this, let us set aside the many complex epistemological and semantic issues, and look toward the metaphysical implications of the basic reductionist idea—that a higher-level "reduced" entity is "nothing but" or "nothing over and above" the lower-level entities and relations

of the reduction base.[98] It would seem that the reductionist is faced with a dilemma. After the "reduction" has been carried out, the reductionist must say either that the higher-level "reduced" entity still exists as a real entity with causal powers of its own, or that it does not. If the "reduced" entity is held still to exist, then the position of the reductionist will be difficult to distinguish from that of the nonreductive physicalist (to be discussed below). If not, then the position of the reductionist will be difficult to distinguish from that of the eliminativist. Either way, the reductionist position will be revealed to be unstable.

In any case, my goal here is not to stake out a position on reductionism for its own sake, but rather to limn the conceptual alternatives available for "naturalizing normativity." For this purpose, it is enough to define "reductionism" with respect to normativity as follows:

> Normative Reduction: To reduce a putative normative phenomenon is to give an account of the phenomenon that is both empirically and theoretically adequate and that neither employs nor presupposes any normative concepts.

If an empirically and theoretically adequate account of a putative normative phenomenon (such as action) could really be given in entirely nonnormative terms, then surely we would be entitled to deny the reality of the normativity of the putative normative phenomenon. Whether one takes an "epiphenomenalist" or a frankly "eliminativist" attitude toward the "reduced" putatively normative phenomenon, then, would seem to be of comparatively small interest. What is of signal interest is that under the scenario we are considering we would appear to have little reason to allow the

[98] In a more adequate discussion, several different forms of reductionism would have to be distinguished: epistemological vs. ontological, and with respect to the latter, causal vs. compositional forms, to name only a few (see Gillett, 2007).

putative normative phenomenon onto our list of the real features of the world. For all intents and purposes, then, reductionism with respect to normativity is virtually indistinguishable from eliminativism, and so there is little reason for us to consider it here as an independent position within the conceptual landscape of "naturalized normativity."[99]

At the opposite extreme from eliminativism is so-called "liberal" (McDowell, 1998) or "naïve" (Hornsby, 1997) naturalism.[100] This second main approach to naturalizing normativity is a view that takes common sense rather than natural science as the arbiter of what is to count as "natural," i.e., as belonging to "nature." Liberal naturalism assumes that human beings are members in good standing of the natural world. This means that all the properties of human beings—indeed, all phenomena associated with, or pertaining to, human beings—are likewise natural. On this view, "natural" contrasts with "supernatural" (what "transcends" nature), but not with "normative." The normative, as a feature of the human, is to be viewed as a subset of the natural. This of course raises the question of how the normative natural phenomena and the nonnormative natural phenomena (let us call them the "physical phenomena") are related. However, liberal naturalism considers itself under no obligation to explain this relation. Rather, liberal naturalism is content to point out the limitations of natural science. Science is cognitively authoritative as far as it goes, but it only goes as far as the

[99] For further discussion of these issues in terms of the realism/anti-realism debate, see Fine (2002).

[100] One might suppose the opposite of eliminativism to be not liberal naturalism, but dualism—by which I mean the positing of a fundamental ontological discontinuity between normative and physical phenomena. For dualists, the natural is to be identified with the physical, understood as the "nonnormative," such that the "normative" and the "natural" become contraries. That being the case, it seems more appropriate to classify dualism, not as a pole within the naturalization project, but rather as the repudiation of that project altogether.

physical phenomena. Its writ simply does not extend to the entirety of nature. That is, liberal naturalism denies premise (2) of the Eliminative Argument outright. But while it is assuredly true that at present the normative phenomena lie beyond the ken of natural science, it is not clear why this limitation should be one of principle, true for all time. The problem with liberal naturalism is that by elevating the present limits of natural science to a matter of principle, it can seem to come perilously close to dualism. For if it is true that the normative is a part of nature, then there must be some connection between the normative and the physical, and what reason can there be in principle why natural science should be forever forbidden from coming to understand the nature of this connection?

In between the two extremes of eliminativism and liberal naturalism is nonreductive physicalism. This third main approach to naturalizing normativity exists in a great variety of different forms, but they all have in common the idea that premise (2) of the eliminativist argument ought to be, not denied outright as in liberal naturalism, but relaxed in such a way as to make it possible for us to admit into our ontology the normative and other higher-level phenomena, which are conceived of as standing in a certain admissible relation to the present physical picture, even though they are not formally a part of that picture. The trick here is to specify the exact nature of the admissible relation between the normative phenomena and the present physical picture. The two main candidate relations are supervenience and emergence.[101] Unfortunately,

[101] Supervenience is the relation between a higher-level ("supervenient") entity or property and a lower-level, acceptably physical ("subvenient") base such that there can be no change in the former without a corresponding change in the latter. It is important that the supervenience relation be conceived of as asymmetrical, in the sense that all causal influence flows from the base "upwards" to the supervenient entity or property. (For discussion, see Savellos & Yalçin, 1995.) Emergence is conceived of in a variety of ways, but in its most important, synchronic sense, it is basically the denial of this last condition, such that at least some causal influence is conceived of as flowing "downwards" from one or more higher-level entities or properties to the base. A further important component of the emergence relation is the idea that the

there are good reasons to believe that the supervenience relation collapses back into epiphenomenalism—and hence, for all practical purposes, eliminativism—while the emergence relation has been criticized as being underspecified and mysterious (see Kim, 1998).

In this chapter, I will pursue a strategy that has affinities with both liberal naturalism and nonreductive physicalism, but which accepts premise (2) of the Eliminative Argument according to the principle that it is desirable that our picture of the world be unified. Instead, I will deny premise (1). That is, I will claim that we have good reason to believe that the present physical picture is radically incomplete. Completing our physical picture will mean enlarging it to make room for the normative phenomena, considered as objectively real. Call this position "normative realism." No heavy-duty metaphysics is required to support normative realism; it merely requires being prepared to accord to normative phenomena the same ontological status that we ordinarily accord to nonnormative phenomena. In other words, ontological parity between normative and nonnormative phenomena will be realism enough for our purposes here.[102] In this way, we will be able to vindicate the liberal naturalist's insistence on according full ontological status and dignity to the normative phenomena, without walling them off from the physical phenomena on principle. At the same time, the nonreductive physicalist's postulate of a relation between the normative phenomena and the physical

---

higher-level entities and properties are not exhaustively determined by the causal properties of the base, which notion is often expressed by the slogan "the whole is more than the sum of the parts." (For discussion, see Bedau & Humphreys, 2008; Clayton & Davies, 2006; Corradini & O'Connor, 2010.)

[102] Thus, if someone were an anti-realist about scientific entities in general, but considered normative phenomena like normative action to be no less real (or more unreal) than nonnormative phenomena like matter, force, or energy, then that person would qualify as a "normative realist" for present purposes.

phenomena will be vindicated, and the relation itself clarified and shown to be admissible, by means of the notion of the nonreductive "grounding"[103] of normative agency in physical phenomena of a certain sort that remains to be specified, but is capable in principle of being fully incorporated into our future scientific world-picture.

4.3 Does Life Have an Essence (and If So, What Is It)?

One more preliminary matter must be attended to before moving to an examination of the various aspects of contemporary empirical research relevant to our question: namely, we need to fix firmly in mind precisely what it is that such research is being adduced to explain. I will begin this section, therefore, by recapitulating briefly some signal points established in Chapter 2 and 3, above.

We must admit that various elementary normative concepts, such as purpose, need, value, and well-being are indeed commonly ascribed to manmade artifacts, notably machines. Thus, my car has a purpose: to get me where I want to go. Instrumentally to that purpose, it may be said to need certain things like fuel, coolant, lubricant, etc., which things are accordingly good for it. Other things, like sugar in the gas tank, are certainly bad for it. Finally, a car may be in a better or worse state of repair— arguably, a form of well-being. We have also seen, however, that there is nothing in the internal workings of

[103] In the sense of Fine (2002; p. 23), in which "Its being the case that S consists in nothing more than its being the case that T, U, . . ." This may sound like a formula for reduction, but as Fine points out:

> "A statement of reduction implies the unreality of what is reduced, but a statement of ground does not. Thus in saying that the fact that P $\wedge$ Q reduces to the fact that P and the fact that Q, we are implying that the conjunctive fact is unreal; but in saying that the fact that P $\wedge$ Q is grounded in, or consists in, the fact that P and the fact that Q, we are implying no such thing. We are adopting a metaphysically neutral stand on whether there really are conjunctive facts (or truths)" (ibid.; p. 24).

a machine—not even a sophisticated machine operated by a cybernetic-control mechanism—that can conceptually underwrite such ascriptions of normativity. A machine considered in and of itself is just a locus of efficient causes. There is nothing in a machine as such that can explain how any particular state of the machine becomes constituted as its normatively preferred state (i.e., end state). Rather, the preferred state of a machine is always determined by a human being. Therefore, all of the other normative ascriptions are logically dependent upon human intentionality, as well. We summarize this state of affairs by saying that the norrmativity we colloquially ascribe to our machines is metaphorical or secondary or derived. It is intuitively clear that even the simplest appetent organisms have literal or primary or original normativity (nobody thinks that the fact that sucrose, say, is good for a bacterium has anything to do with human intentionality). But it is far from clear what it is about organisms that conceptually underwrites this intuition. This is the challenge that I hope to begin to meet in this chapter: to rationalize the distinction between organisms and machines with respect to primary or original normativity.

Now, it is also true that scientists commonly refer to organisms as "machines," and that there are certain respects in which the label does seem apt. However, if organisms were phenomenologically (especially, behaviorally) exactly like manmade machines, then I do not believe that we would experience the deep intuition with respect to the derived/original normativity distinction in the way that we do. But as it is, organisms are in fact very different indeed from any manmade machines behaviorally, and I believe it is this fact that underwrites the intuitive distinction at the phenomenological level. It is this difference, above all, that must be given a scientific

rationale at the microphysical level, if living systems are ever to be accepted as a natural kind. Let me explain what I have in mind. From here on out I will simply assume that the proper scope of our concept of normative agency is life, i.e., organisms as such. But even if this claim is true, it does not yet tell us very much about the metaphysical ground of normativity—the Ground Problem. More specifically, it tells us little about the nature of the relationship between normative biological phenomena and nonnormative physical phenomena. It is to this question that this final section of the chapter will be devoted.

On the assumption that normative agency is coextensive with living things, it is clear that the question of the natural ground of normativity is closely related to the question of whether life has an essential nature, and if so, what it is. In other words, if living systems constitute a natural kind, what are the criteria for membership in it, and what do these criteria have to do with normativity?

For a long time, it was fashionable to deny that living systems constitute a natural kind at all. And yet, while it is true that there are some difficult cases, it is also true that both common sense and biological science operate on the assumption that living systems are essentially different from nonliving systems. After all, there is no mistaking a cat for its saucer. Nor, for that matter, is there any mistaking a free-living cell like a bacterium for a mote of dust, or even for a collection of Bénard cells or for the contents of a Belousov-Zhabotinsky reactor. This striking and highly intuitive difference between living and nonliving phenomena would seem to provide more than adequate warrant for regarding the question of the essential nature of life as a legitimate object of scientific inquiry.

Nevertheless, scientists have traditionally been reluctant to undertake the task of defining necessary and sufficient conditions for life, preferring instead to produce lists of "signs of life," such as nutrition, growth, self-movement, sensitivity to the environment, metabolism, reproduction, evolvability, and the like. Yet, as Lange (1996) has pointed out, such "signs of life" only count as signs of life, as opposed to some other property, because systems already recognized as living exhibit them. In other words, a conception of life as having some kind of essential nature is presupposed by any such list. As Lange (1996) puts the point:

> The "signs of life," while neither individually necessary nor jointly sufficient for something to be living, bear a special relation to vitality. Certain things display a given "sign of life" because they are alive (while certain other, nonliving things display a given "sign of life" for some other reason). In other words, that a given thing is living explains why in certain circumstances it can reproduce, metabolize, move, and so on. This distinguishes the "signs of life" from other properties, including their simulations. And, I will argue, it is in connection with such explanations that the concept "life" performs its work in biology. (ibid.; p. 231)

Lange's point might seem to be an obvious one: that a system is able to move itself because it is alive, and not alive because it can move itself. But, of course, this raises the question of what the essential nature of life is, precisely, such that it can give rise to self-motion, metabolism, reproduction, and the rest of the "signs." And that is clearly an empirical question, if anything is.

Luckily, this way of viewing matters is no longer so much disputed as it used to be, as a rising tide of recent publications devoted to interrogating the essential nature of life attests.[104] Accordingly, I will not spend any more time defending the notion that

[104] See, e.g., Agutter & Wheatley (2007), Battail (2009), Bedau (1996), Bedau & Cleland (2010), Bruggeman et al. (2002), Cleland & Chyba (2002), Cornish-Bowden et al. (2007), Dürr et al. (2002), Gayon et al. (2010), Kolb (2007), Moreno & Ruiz-Mirazo (1999), Penzlin (2009), Popa (2004), Rizzotti (1996), and Ruiz-Mirazo et al. (2004, 2010).

living systems constitute a natural kind, but will simply assume that they do. The question, then, is: What is the fundamental principle or principles that account for life?

Although this is clearly an empirical question, I think that philosophers still have a legitimate role in clarifying it. First, in order even to begin to investigate this question seriously, we must decide precisely which phenomena we are attempting to account for. To do this, we must ask the question: Are some of the "signs of life" more fundamental than others with respect to our understanding of the difference between living and nonliving systems? If so, which one(s)?

Many would agree, I think, that the concept of self-preservation lies close to the heart of our concept of life. It is this idea that accounts, too, for the normativity we associate with life, at the most fundamental level. However, life as we know it involves two rather different concepts of self-preservation, one in relation to particular living things (individual organisms, or life-tokens) and the other in relation to classes of living thing (biological species, or life-types).[105] The first question, then, is whether the principle we seek is more closely connected to the former or the latter concept of self-preservation. To give them convenient labels, I will follow tradition by referring to life-token self-preservation as "metabolism" and life-type self-preservation as "reproduction."

The default position in both scientific and philosophical circles is generally taken to be that reproduction is the more fundamental concept. I think the reason for this is the

[105] The individual organism is itself sometimes construed as a "type" or "form" in relation to the matter of which it is composed, which is in continuous turnover. However, this seems to me misleading. The individual water molecules composing a hurricane are also in continuous turnover. So, this sort of turnover of the material constituents of a dynamically stable, nonlinear, steady-state thermodynamic process, or "dissipative structure," is not restricted to organisms (see Kondepudi & Prigogine, 1998; pp. 409–452). This shows that, while being a dissipative structure may well be a necessary condition for being a living system, it is not sufficient. This issue will be addressed in more detail below.

fact that reproduction is closely associated with the notion of evolution, and biologists and philosophers alike generally view the phenomenon of life through the lens of the theory of evolution. However, as Boden (1999), Bouchard (2011), Cornish-Bowden (2007), Moreno & Ruiz-Mirazo (1999), and others have pointed out, this is a mistake. The reason is very simple. A world in which metabolizing creatures are immortal—or else one in which each generation begins again de novo—is clearly conceivable, and life could clearly be instantiated in such a world, whereas a world in which "reproduction" (as in a spreading wildfire, a growing crystal, etc.) exists in the absence of metabolism is a world that is intuitively devoid of life. Moreover, in real organisms reproduction is simply one aspect of metabolic control—is, in fact, just another of the innumerable metabolic processes that collectively constitute life. However, this ought not to be mistaken for a merely empirical observation. The point is a conceptual one. As Cornish-Bowden and coworkers (2007) have put it:

> . . . staying alive [is] the problem that needed to be solved first: the early living entities could not begin to reproduce or evolve until they had learned how to stay alive, maintaining organizational invariance in the face of changing conditions. (ibid.; p. 844)

For this reason, metabolism, not reproduction, is clearly the more fundamental concept, and the one which comes much closer to embodying the essence of life. But what, exactly, do we mean by "metabolism" in this context?

Maselko and Maselko (2009) have recently given eloquent expression to the fundamental idea underlying the concept of metabolism:

> At the heart of the challenge facing the research community is that biological cells exhibit complex spatiotemporal organization. We cannot reproduce the complexity of cellular organization by simply mixing all the cellular components at the right concentrations. In life, the concentrations of a multitude of chemical compounds are organized in both space and time. Furthermore, the concentrations

> of different chemicals are continuously changing, and this spatiotemporal organization is incredibly precise. The possibility of successfully assembling all of the chemicals in their proper positions and correct concentrations is almost nonexistent. The difference in the precision of spatiotemporal organization of chemical versus biological systems is an important distinction between the two. . . . [O]ne of the issues associated with the synthesis of a prebiological entity is understanding and achieving the precise temporal and spatial chemical organization necessary for biological systems to function. (ibid.; p. 538)

So, metabolism—in the sense of the self-preserving, self-maintaining, or self-perpetuating capacity of the organism or life-token system—has something to do with the way in which thousands of chemical reactions are coordinated. Maselko and Maselko speak of "organization," but of course what they mean is "dynamic organization" (they do speak explicitly of "organization in time" as well as space). So, it is really the coordination of events we are talking about.[106] And this coordination must be considered normative to the extent that what happens happens under the global constraint of the self-preservation of the system as a whole. In short, we can see from these considerations that the essential feature of living systems is their capacity (within limits) to meet the various contingencies of their existence in such a way as to preserve themselves in existence. In Chapter 3, we agreed to refer to this universal compensatory property of living things by the term "adaptivity."

I have already offered in Chapter 3 a variety of considerations, both conceptual and empirical, showing that there is good reason to doubt the claim that the theory of natural selection constitutes an adequate framework for explaining adaptivity. This conclusion, together with the prima facie case for **TRB** established in Chapter 2, gives rise to the question whether we may not therefore legitimately entertain the postulate of a

[106] On this point, see especially Kauffman et al. (2008).

physical principle coordinating life processes. In the next section, I will offer a general conceptual argument (the "Ground Argument") from the point of view of physics for why we should expect there to be such a global coordinating principle.

4.4 The Ground Argument

We saw in Chapter 3 that a variety of recent empirical discoveries make it no longer plausible to view the phenotype of an organism as mechanically derivable from the genome. And we discussed a number of examples in which no plausible evolutionary explanation could be given for evident cases of the adaptive compensatory capacities of living things.

However, in a sense, it really does not matter whether we view the adaptive capacities and powers of living systems evidenced by the above examples as classical, neo-Darwinian selected "adaptations," or as "spandrels," or as "cryptic norms of reaction," or as something else, because whatever the case may be, the capacity in question still requires physical (that is, microstructural) explanation. After all, biological systems are a species of physical system, and it is not clear why biological explanation should differ from physical explanation, at least with respect to the properties of biological individuals.[107] From a physical point of view, then, it is the underlying dynamics of a living system that determines which parts of the system are more or less

[107] Respecting the evolutionary process as a whole, the "lawlessness" of biology is sometimes elevated to a fundamental principle, as in Gould's (1989) famous assertion that you "cannot replay the tape" of evolutionary history on earth. This is in fact far from certain, and indeed much evidence points to the contrary view (that there may be laws that lead evolutionary dynamics along particular pathways; see Conway Morris, 2003), but however that may be, there is surely no good reason to doubt that the causal powers of individual living beings derive from their microstructural constitutions.

crucial to the proper functioning of the whole. Therefore, instead of relying upon such conceptually opaque and operationally weak distinctions as "adaptation/spandrel," or "selection-for/selection-of," philosophers of biology would be better off following engineers in making more straightforward distinctions among traits based upon such conceptually transparent and operationally robust concepts as "central/peripheral," or "more/less critical," with respect to viability.[108]

At any rate, the basic point is this. Whatever the history of any real system may be, we always have still to ask the further question: According to what general dynamical principles is the system operating now? And this is just as true for life-token systems as it is for any other kind of physical system, a point which is beginning to be explicitly recognized by biologists. For example, Whitesides (2008) remarks:

> Today, we understand many aspects of the cell and many fragments of the network, but not how it all fits together. We particularly do not understand the stability of life and of the networks that compose it. Our experience with other very complicated networks (e.g. the global climate, air-traffic-control systems, the stock market) is that they are puzzlingly unstable and idiosyncratic. But unlike these and other such networks, life is stable—it is able to withstand, or adapt to, remarkably severe external jolts and shocks; and its stability is even more puzzling than the instability of the climate. We have a hard enough time understanding even simple sets of coupled chemical reactions. And we have, at this time, no idea how to understand (and certainly not how to construct) the network of reactions that make up the simplest cell. (ibid.; p. xiv)

Moreover, not only do we not understand the principle(s) underlying the stability of the cell, we will never do so, so long as we solely employ reductive means of investigation, which are necessary but not sufficient to address the problem of the missing stability principle. As Auffray and colleagues (2003; p. 1135) have put the point:

[108] Cf. Skewes & Hooker's (2009; pp. 290–291) distinction between "dominant" and "derived" norms, and Mossio and coworkers' (2008; p. 831) distinction between "primary" and "secondary" functions, which I take to be more or less equivalent to each other and to my suggestion in the text.

"An exhaustive knowledge of the structure, function and relation of the components of biological systems is necessary but insufficient to understand phenotypes." R. Phillips and Quake (2006) concur:

> . . . even a perfect understanding of each and every individual molecular machine would be inadequate for explaining what goes on in a cell, just as an understanding of the hydrogen atom is merely a prelude to explaining the electronic behavior of crystalline solids and, more dramatically, collective effects like the quantum Hall effect. (ibid.; p. 40)

Why is this realization now finally coming to be more widely accepted among practicing biologists? In addition to the recent discovery of the astonishing complexity of gene regulation previously mentioned, one reason is probably the disappointment that followed in the wake of the Human Genome Project, which vividly demonstrated that knowledge of DNA sequences does not of itself yield knowledge of cell structure or function. Another reason may be that recent results from such new technologies as green fluorescent protein (GFP) studies and the fluorescence recovery after photobleaching (FRAP) process have begun to allow us to investigate the dynamics of individual cells in vivo in real time for the first time. What these techniques reveal is the fact that many, if not most, of the putative "machines" within living things are very far indeed from the rigid structures typical of manmade machines. Rather, they are highly dynamical in nature, self-assembling and -disassembling moment to moment in response to the cell's metabolic needs.[109] These startling and unexpected discoveries are likely to do more than anything else in coming years to force a change in perspective upon philosophers, as

[109] For details, see Janicki & Spector (2003), Karsenti (2008), Kirschner & Mitchison (1986), Kirschner et al. (2000), Lippincott-Schwartz et al. (2000), Mayer et al. (2009), Misteli (2001, 2007), Nédélec et al. (1997), and Whitesides & Grzybowski (2002). For discussion, see Gierasch & Gershenson (2009) and Kurakin (2005, 2009).

well. “Machines” which self-assemble and -disassemble as needed are not “machines” at all, in any literal sense of that word. As Kurakin (2005) has put it:

> To summarize, the newly revealed and unexpected properties—such as steady-state character, transient self-organization on demand, stochastic dynamics and interconnectedness—that characterize cellular structures and molecular machines believed to exist as pre-assembled complexes designed for certain functions according to programs and blueprints, clearly suggest the inadequacy of expectations and assumptions based on the mechanistic intuition. (ibid.; p. 250)

I conclude from the foregoing that, if it is true that living systems may be properly understood as natural kinds whose essence is defined (at least in part) by the power of adaptivity, then we may legitimately raise a question concerning the sort of physical principle that underlies this capacity of living systems. This brings us to the more broadly conceptual argument I promised several pages back in support of the existence of such a global coordinating principle in living things.

Before proceeding to this argument, however, I need to discuss an idea that, while controversial, has become widely accepted in biology: namely, the claim that the laws of physics as currently understood cannot explain the fact that living processes are functionally coordinated, because the latter are contingent with respect to the former. The locus classicus for this claim is E. Nagel (1979).[110] In the context of an analysis aimed at distinguishing goal-directed systems, whether living or artificial, from non-goal-directed systems, Nagel says the following (the “variables” he refers to are those involved in his example of the homeostatic regulation of the water content of the blood):

> It should be noted that these variable are independent of (or “orthogonal” to) each other, in the sense that within certain limits the value of either variable at a given moment is compatible with any value of the other variable at that same moment.

[110] Nagel tells us that he was influenced by Sommerhoff’s (1969, 1990) cybernetic analysis of living systems.

> As will be seen presently, such orthogonality of variables is an important requirement. (ibid.; p. 287)

Nagel goes on to emphasize that he does not of course mean that there is no causal dependency at all between the variables in this type of homeostatic system. Obviously, there must be, if the homeostatic mechanism in question is to have any causal efficacy. Rather, what he means is that the causal dependency between system variables is not one that can be described directly in terms of the working out of the laws of physics or chemistry. The variables are such that they could take any one of a number of different values, so far as the laws of physics are concerned. That is, the values ultimately assumed are contingent with respect to the laws of physics; the laws of physics do not determine them. Nagel also expresses this idea by saying that such differing values are "compatible" with each other, meaning that no law of physics forbids them from being different from what they are in any given instance. All of this being the case, we must look to other factors to explain why the variables have the values that they in fact have. In the case of manmade artifacts, this other factor is a particular organization imposed from the outside by human intentions and agency. (One way of looking at the Ground Problem is as the challenge of articulating what the corresponding factor in living systems might be.) Once imposed, the particular set of physical dependencies that embodies the intended goal-directed organization of the artifact is stabilized by reliance upon such laws of physics as determine the rigidity of metals, for example. But there is no law of nature requiring that any particular set of physical dependencies be instantiated in just the way that it is upon a given occasion (or, as Nagel puts it, that the system variables have just the values that they in fact have). Rather, human agents have taken advantage of this lack of full physical determination in order to choose a particular configuration of system variables to

their liking. In summary, as Nagel put it (ibid.; p. 289), "the known (or assumed) 'laws of nature' impose no restrictions on the simultaneous values of the variables" in a goal-directed system. Let us call this the "Orthogonality Condition" on living systems.[111]

One reason that the Orthogonality Condition is important is that it seems to track closely our intuitive understanding of the difference between living and nonliving systems. Indeed, it appears to be crucial for explicating this distinction, inasmuch as other ideas such as "plasticity" and "persistence" that are sometimes advanced as definitional of goal-directed systems could just as well be attributed to such non-goal-directed systems as, for example, a marble rolling around in a bowl or rainwater flowing down a hillside. Nagel believes (ibid.; p. 288) that it is precisely because the Orthogonality Principle fails in the case of such systems that we are not inclined to say that these systems are "goal-directed," even though the marble and the rainwater appear to exhibit behaviors that might plausibly be described as "persistent" and "plastic" as they approach their equilibrium states (the marble will arrive at the same final state from many different

[111] For further discussion, see Monod (1972) (note that Monod uses the term "gratuité [gratuity]" to express what Nagel means by "orthogonality"). Maynard Smith (2010; pp. 133–134) suggests that the Orthogonality Condition be explicated by means of the notion of "symbol," that is, a sign whose meaning is contingent with respect to its physical instantiation. Pattee (2001) concurs, terming the fundamental distinction one between controlling, rate-independent "symbols" and controlled, rate-dependent "dynamics." Now, there is little doubt that "information" and related concepts, such as "signal," "message," "code," "computation," "communication," "control," "regulation," "transcription," "translation," "editing," "proofreading," and the rest, are indispensable in biology, and lie close to the heart of the problem we are investigating—whether in relation to cell biology (Beckerman, 2005; Bray, 2009; Loewenstein, 1999), collective cell behavior (Baluška & Mancuso, 2009; Barlow, 2008; Brenner et al., 2006; Mehta et al., 2009; Waters & Bassler, 2005), or our general theoretical understanding of the living state (Barbieri, 2003, 2007; Battail, 2009; Bruni, 2007; El Hani et al., 2009; Terzis & Arp, 2011). Nevertheless, it must be remembered that the concept of "information" implies the existence of an agent for which the information is meaningful. Absent an account of biological "meaning," use of the notion of "information" in biology is at best question-begging, at worst incoherent (Boniolo, 2003; Deacon, 2010; Griffiths, 2001; Jablonka, 2002; Queiroz & El Hani, 2006; Roederer, 2005; Sarkar, 2005). In other words, insofar as they are to be construed realistically, and not as mere subjective projections, "information" and related concepts refer to phenomena that constitute an important part of the explanandum of normative agency. It is for this very reason that we must not imagine that invoking these concepts gains us any explanatory ground with respect to the project of naturalizing normativity. Accordingly, I prefer to speak here of the "Orthogonality Condition" and set aside all information-related concepts.

initial states; the rainwater will move around obstacles placed in its path; etc.). Moreover, the still stronger claim is sometimes encountered that it is the Orthogonality Condition that makes life possible, in the sense that it supplies a certain "freedom" necessary for living systems to exist, the idea being that if all of the events occurring inside a living system were determined directly by the laws of physics and chemistry—that is, if all of the internal variables were directly dependent upon one another—then life as we know it could not exist.[112]

In summary, the upshot of the Orthogonality Condition is that the functional coordination of events in living systems cannot be explained solely by reference to the minimization of free energy or, more generally, of action, in the physical sense of that term (energy times time). The functional coordination of living things must, of course, be compatible with the least-action principle, as well as with the laws of thermodynamics and all the other laws of physics, but no energy- or entropy-related principle by itself can account for the functional coordination of events within living things.[113]

I am now ready to present the following informal "Ground Argument":

Life-token systems, i.e., organisms, are dynamically stable, in the sense that they persist for times that are long in relation to their thermodynamic relaxation rates. Since

[112] I should perhaps add that the "freedom" in question need not be construed as that of the metaphysical "libertarian"; I take it that neither Nagel nor Monod wishes to deny determinism. Rather, the "freedom" in question is only relative to the laws of physics themselves. Any particular functional organization, though "free" in the sense that it might have been otherwise insofar as the laws of physics are concerned, could still be explained deterministically by taking into account additional causal factors, such as, on this view, the process of natural selection.

[113] This point is analogous to the more widely discussed point that life does not violate the second law of thermodynamics, which becomes clear once friction, export of heat to the environment, etc. are taken into account; but this does not mean that the second law by itself explains the existence of life. Rather, while life is consistent with the second law, it seems to transcend it in a way that we are struggling here to explain.

the stability of all stable systems is due to some physical principle (stability always has a physical explanation; stable systems are not stable for no reason), the dynamical stability of organisms is due to some physical principle, as well. The dynamical stability of organisms consists in the coordination in space and time of the physical and chemical events occurring within it. Since such coordinated events are just the "functions" of an organism, let us rename dynamical stability "functional stability" (following Strand & Oftedal, 2009) to emphasize this feature of living systems. Therefore, functional stability is due to some physical principle. We know, in light of the Orthogonality Condition, that the physical principle giving rise to functional stability is not that of free-energy minimization or the least-action principle.[114] So, functional stability is due to a physical principle that is distinct from any energy- or entropy-related extremal principle. We have been calling the property of adaptive compensatory action that gives rise to the functional (or dynamical) stability of living things, "adaptivity." Therefore, adaptivity—and the capacity for "action" in the philosophical sense—must correspond to a physical principle that is distinct from action in the physical sense.

Why do I call this the "Ground Argument"? Because it seems to me to provide additional reason for believing that a global coordinating principle in living things must exist. And if it does, then it would be natural to identify such a principle with the natural "ground" of normative agency. But what can we say about the nature of this physical principle, which until now we have been content merely to label "adaptivity"?

[114] To repeat, the theory of natural selection is not relevant here; however one understands the origin of functional stability, it is always legitimate to ask the further question of how functional stability is physically possible at present.

4.5 The Concept of "Autonomous Agent" in Contemporary Theoretical Biology

Recently, a number of cognitive scientists, theoretical biologists, and philosophers have begun to speak quite openly of "teleology," "agency," and even "normativity" in an effort to get at what it is that distinguishes living from nonliving systems.[115] I will refer to these authors as the "Autonomous Agency Theorists." However, while all of the Autonomous Agency Theorists acknowledge the nature of the problem—indeed, some of them are quite explicit and even eloquent in their insistence that normativity and agency lie at the heart of life—none of them, in my judgment, comes fully to grips with the true depth of the difficulty posed by the problem.

While there is of course a fair amount of diversity among these various authors, nearly all of them approach our problem from one or the other of two angles, or both. The first (and by far the majority) approach is taxonomic, in which the author(s) sketch in a general typology of physical and living systems, and their relations to one another, in order to home in on the distinctive characteristics of organisms by means of a systematic consideration of relevant similarities and contrasts between living and non-living systems. For example, R. Campbell (2009) elaborates a typology of processes with many fine distinctions, such as "persistence" versus "cohesion," "self-maintenance" versus "recursive self-maintenance," "error detection" versus "flexible learning," and so forth. Similarly, Skewes and Hooker (2009) distinguish between "basic autonomy," "self-

[115] See, for example, Barandiaran & Moreno (2008), Barandiaran et al. (2009), R. Campbell (2009), Chemero (2009), Christensen & Bickhard (2002), Christensen & Hooker (2001), Di Paolo (2005, 2009), Di Paolo & Iizuka (2008), Di Paolo et al. (2010), R.D. Ellis & Newton (2010), Freeman (2001), Hanna & Maiese (2009), Hooker (2009a, 2009b), Juarrero (1999), Kauffman (2000, 2004), Kauffman & Clayton (2006), Kauffman et al. (2008), Kelso (2008), Moreno & Ruiz-Mirazo (1999), Mossio et al. (2009), Ruiz-Mirazo et al. (2010), Skewes & Hooker (2009), E. Thompson (2007), and A. Weber & Varela (2002).

directedness," and "anticipation." In the same vein, Mossio and coworkers (2009) draw an interesting distinction between "organizational closure" and "organizational differentiation," as different forms of "self-maintenance." Finally, Barandiaran and colleagues (2009) distinguish three requirements for autonomous agency: individuality, interactional asymmetry, and normativity.

I do not mean to suggest that such efforts at clearing the conceptual ground are not important. However, in reading this literature, too often one gets the impression that what the authors feel we require above all is the correct perspective on the problem—the right taxonomy for organizing our thoughts on the nature of life. But surely the real difficulty lies not so much in an incomplete or incorrect classification system as in a lack of basic understanding about how living systems are physically possible in the first place. What we require, then, above all, is knowledge of the sort of physical principle that might conceivably help us to distinguish the living state of matter from the nonliving state. On this point, the Autonomous Agency Theorists are decidedly more reticent.

Nevertheless, some of them do attempt to confront this challenge (and this is the second angle of attack mentioned above), at least in a cursory way. Basically, Autonomous Agency Theorists make appeal to the concept of "self-organization," which is then supposed to be explained by reference to the physical discipline of nonequilibrium thermodynamics, as well as the adjunct mathematical discipline of nonlinear dynamics. For example, both Mossio and coworkers (2009) and Skewes and Hooker (2009) discuss the physics of candle flames, and the former throw in Bénard cells for good measure. Even when an Autonomous Agency Theorist does not explicitly discuss these ideas, as is often the case, they are nearly always assumed as a background for the discussion.

Now, once again, I do not wish to be understood to imply that these ideas are not interesting. As we shall see presently, there is indeed much insight to be gleaned from looking upon living things as open thermodynamic systems, nonlinear oscillators, and the like. However, there is a serious problem with the concept of "self-organization" that is usually overlooked in this literature. Partly, the problem is terminological: There is a lack of consistency in the way in which the terms "self-assembly" and "self-organization" are used (Halley & Winkler, 2008). Obviously, linguistic usage cannot be dictated, but it would be a good idea if the following three terms, say, were to be consistently distinguished: "self-assembly," referring to processes involving spontaneous (exergonic) reactions requiring a pre-existing "seed" (i.e., template) and leading to equilibrium; "self-ordering," meaning heterogeneous steady-state systems arising out of a homogeneous medium through fluctuation under constrained energy flow; and "self-organization" properly speaking, indicating functionally coordinated, compensatory action (adaptivity) under the constraints of the Orthogonality Condition and overall viability. Notice that neither self-assembly nor self-ordering can shed very much light on self-organization, according to these definitions. The reason is that the first process is driven by free-energy minimization and the second by the minimization of thermodynamic potentials.[116] In this respect, the nonequilibrium thermodynamic processes in cells are not essentially different from such nonliving steady-state structures as hurricanes, candle flames, Bénard cells,

[116] The precise nature and role of extremal principles in nonequilibrium thermodynamics remains controversial (Grandy, 2008), but for my purposes here it is enough that steady-state processes are ultimately determined by the principle of least action in some form or other, which no one I think disputes, while functionally coordinated behaviors are not.

and the like.[117] All are simply minimizing thermodynamic potentials under given constraints. Therefore, such processes cannot constitute the physical principle underlying the functional coordination of living things.[118] Similar observations apply to explanations based on nonlinear dynamics, networks, and the like. That is to say, nowhere in the writings of the Autonomous Agency Theorists, so far as I am aware, is there to be found a clear recognition of the grave difficulty raised by the Orthogonality Condition, with its implied requirement (on the hypothesis of **TRB**) of a physical coordinating principle transcending least action.

Of course, most of the Autonomous Agency Theorists do show their awareness of this general problem. For example, Mossio and colleagues (2009) acknowledge that:

> Yet, although necessary in order to naturalize teleology and normativity, organizational closure is not a sufficient condition for functional attributions. Minimal self-maintaining systems, such as flames or hurricanes, do instantiate closure, but their components do not have functions. To have functions, self-maintaining systems must belong to a specific class in which different contributions to the self-maintenance of the system can be distinguished. (ibid.; p. 825)

That all sounds eminently forthright and reasonable, until one reflects that nowhere do Mossio and coworkers discuss the physical interpretation of their notion of "organizational differentiation." But that is the very heart of the matter on their view! Without at least a gesture in the direction of an explanation of this phenomenon, one is left with very little by way of a deeper understanding of adaptivity. In short, reading this

[117] See note 105, above. For further discussion of the relevance of nonequilibrium thermodynamics to life, see also Kurzynski (2006), Schneider & Kay (1995), Schneider & Sagan (2005) and Swenson (1998).

[118] For further discussion, see Abel & Trevors (2006), Johnson & Lam (2010), Orgel (2008), and Penzlin (2009).

literature one often has the frustrating feeling that the full depth of the difficulty of the problem under discussion is never properly appreciated.

While I am strongly sympathetic to the goals of the Autonomous Agency Theorists, I believe that their work too often founders on an ambiguity between self-ordering processes and self-organization, properly speaking, as defined above. While there is something deeply right about their general approach to the problem of normative agency, they do not dig down deep enough. They are not much to be blamed for this, seeing that the fault is a general one, shared by most philosophers and biologists alike. As Ó Nualláin (2008) has rightly observed:

> We currently hide our ignorance of the specific processes that obtain through invoking almost as shibboleths "dynamical systems," "emergent behaviours," "complexity," "self-organisation," and so on. The task of unpacking these words into something resembling hard science may take a generation . . . (ibid.; p. 242)

What, then, do I have to offer? Nothing very concrete. But I would like to point to two bodies of theory as evidence that the situation may not be entirely hopeless. Let us remember what it is that we are trying to understand: the normative agency that appears to be an essential attribute of living systems as such. I will divide the problem into two parts: teleology/adaptivity and activity.

Functional stability is a strange sort of property, from the point of view of ordinary physical theory, because it has two characteristics that seem almost magical from a strictly physical point of view: purposiveness (goal-directedness) and adaptivity. But for some time now, theorists have been employing ideas borrowed from the mathematical discipline of nonlinear dynamics to model at least some of the distinctive aspects these properties. Let us briefly review some of the advantages of this perspective.

To begin with, it is of crucial importance to understand that with few exceptions the functional motions performed by biological structures are cyclical (or rhythmic) in character, with similar events recurring according to temporal patterns of varying structure and complexity (Gilbert & Lloyd, 2000; Goldbeter, 1996; Yates & Yates, 2008). This means that functional processes can be modeled as nonlinear oscillators,[119] and their motions as closed phase-space trajectories, or "attractors."[120] Delattre (1986) was the first author, so far as I am aware, to explicitly propose that nonlinear dynamical concepts be used as a means of modeling the purposive or end-directed character of biological functions. Yates (1994, 2008) has done more than anyone else to generalize Delattre's insight, as well as those of others, into an overarching theory of biological function. The philosophical payoff from this approach is twofold.

First, modeling biological functions as nonlinear oscillators solves the infamous "backward causation" problem. The function's phase-space attractor serves as a virtual state corresponding to its goal state. Backward causation is forbidden in science, but virtual states are not. This point also ties into the reason why the Orthogonality Condition is so important. The motions of, and internal to, living systems must be "regulated" or controlled," and not be merely the inevitable result of the playing out of deterministic physical law. That is one of the signal differences between living and nonliving systems.

[119] The nonlinearity of biological oscillators corresponds to the fact that most functional motions occur in response to crossing a critical threshold, to reception of a "signal," to change of state of a "switch," or to some other sort of trigger.

[120] See Glass & Mackey (1988) and Winfree (2010). For a review of dynamical systems theory, see Jackson (2001). For the graph-theoretic (network) approach, see Csermely (2006). For surveys of biological applications, see Camazine et al. (2001), Harrison (2011), Kaneko (2006), Kondo & Miura (2010), and Scott (2007). For the instructive case of the dynamical construction of animal limb coordination, see Bejan & Marden (2006), Frank et al. (2009), Kelso (1995), Turvey (2004), and Warren (2006).

But as Thalos (2007; p. 142) has pointed out, "[s]tructures of control are relations among causes pertaining to potentialities and potential events as well as actualities and actual events." The fact that regulation and control, and thus potentialities, are an essential aspect of life is one reason why dynamical systems theory with its virtual-state "attractors" is such a useful conceptual tool of modeling functional behavior.[121]

Second, nonlinear oscillators have a mathematical property known as "metastability" that models nicely the property of living systems that we have been calling "functional (or dynamical) stability." Recall that we have been referring to the general capacity for adaptive compensatory action that underlies the functional stability of organisms as "adaptivity," and that adaptivity comes in two basic forms: "robustness," which is the ability of a living system to recover from perturbation in such a way that the system's original dynamical equilibrium regime is restored, and "plasticity," which designates the ability of a system, following perturbation, to discover a novel dynamical regime consistent with viability. An example of robustness would be the healing of a dog's broken limb; an example of plasticity would be a dog's learning to walk with a very different gait, if it loses a limb altogether. Nonlinear dynamics is able to model this dual capacity of robustness/plasticity very naturally by means of the notion of metastability. Under many or most perturbations, a metastable system will still find itself within its basin of attraction, and so will spontaneously return to its original equilibrium regime (this mathematical property is called "equifinality"). However, under severe perturbations that carry the system outside of its original basin of attraction altogether, the system may be able nevertheless to find an alternative attractor corresponding to a

[121] The concept of "control" will be discussed in greater detail below.

different way of doing things (this is known as a "bifurcation"). In summary, the mathematical structure of nonlinear dynamics seems tailor-made to represent the two faces of biological adaptivity. In other words, the dynamical-systems viewpoint helps us to see the unity underlying what might otherwise appear to be two very different capacities of living things. This is surely no coincidence, but rather an indication that we are on the right track scientifically. While it is not the whole story, the identification of the goal states of biological functions with phase-space attractors may plausibly be viewed as a distinct step forward in our scientific understanding of living systems. To my mind, Beloussov's (2009) nuanced judgment on the value of nonlinear dynamics for modeling the ubiquitous teleological features of life is just about right:

> I would not like to state that introduction of attractors gives a final solution of an endless problem of teleology, but it is at least a heuristically useful step for combining the elements of teleology with more ubiquitous modes of causality. (ibid.; p. 46)

Whether it is a sufficient step is a question we will consider in due course.

First, however, let us review what it is that we wish these physical theories to help us to explain. In order to go on existing, an organism must be able to choose among states that are indifferent with respect to the laws of physics (the Orthogonality Condition), but which are anything but indifferent with respect to the continued existence of the organism. In other words, the reason that the Orthogonality Condition is a condition on any system that we would count as living is because the ability to act under the constraint of self-preservation, rather than according to the universal principle of least action, is of the essence of life. This observation bring us to the verge of grasping the principle that distinguishes living from nonliving systems, I believe, but it does not yet quite take us all the way. The reason is that manmade artifacts with functions and goal states determined

externally by human intentions must also meet the Orthogonality Condition. An automobile could not get us where we want to go if its internal components did nothing more than minimize free energy or thermodynamic potentials—if, say, a particular configuration of the steering column led to one and only one configuration of the wheels—and could not be physically influenced in accordance with global constraints (in this case, the will of the driver). So, nothing we have said so far has yet grasped hold of the principle that explains how a living system differs from a machine. What is it in the case of the organism that is acting as a global constraint analogous to the driver's will in the case of the automobile? That is the crucial question that cannot be ducked if we are ever to understand teleology, normativity, and agency in a realistic manner.

Abel (2010) has brought a new level of critical acuity to bear on this question. He points out a logical fault which he claims several authors invoking the concept of "self-organization" are guilty of—namely, that of conflating the concept of physical "constraint" with that of "control."[122] In fact, the crux of the problem confronting us is precisely to explain by virtue of what physical principle a "constraint" can be transformed into a "control." To help oneself to the notion of control, on the basis of having invoked the notion of physical constraint, is simply to beg the question.[123]

The problem arises because many authors see that the Orthogonality Condition is somehow at the heart of the problem, but fail to come fully to grips with the fact that the

[122] Abel cites a different literature from the one I have been considering; it is a fine question (which I do not have time to pursue here) whether my Autonomous Agency Theorists are also guilty of this conflation. Kauffman (2000, 2004), at least, is quite aware of the danger lurking here, though he is more optimistic than Abel that we will eventually be able to get safely around the problem (see below).

[123] Other authors who have argued along similar lines in the past include Pattee (1982, 2001) and Rosen (1991, 2000).

Orthogonality Condition by itself does not yet provide us with a means of distinguishing organisms from machines. Abel calls the parts of functionally organized systems that obey the requisite Orthogonality Condition "configurable switches." He writes (ibid.; p. 14) that "[c]onfigurable switch-settings allow the instantiation of formal choice contingency into physicality. While configurable switches are themselves physical, the setting of these switches to achieve formal function is physicodynamically indeterminate—decoupled from and incoherent with physicodynamic causation." This much, all parties can agree upon: parts meeting the Orthogonality Condition (i.e., "configurable switch-settings") are a necessary condition for functionality and normative agency (what Abel calls "formal choice contingency"). The trouble comes when we attempt to slide from the necessity of "configurable switch-settings" for normative agency to their sufficiency, on the basis of some vague notion of "self-organization." This, Abel maintains, is nothing other than to smuggle in surreptitiously the very normative agency that we are attempting to account for naturalistically. As he puts it:

> A hill does not become the simple machine of an "inclined plane" until agency chooses to use the hill to assist in overcoming the formal challenge of outsmarting the agent-perceived problem of gravity. . . . Without the reality of formal choice contingency, physics cannot even distinguish "work" from "wasted energy". The mere transfer of energy from one entity to another often has nothing to do with utility. (ibid.; p. 20)

The challenge that Abel is making to "self-organization" theorists, then, is to explain what physical principle turns mere transfer of energy into utility, or constraints into control. For, how can we possibly give a naturalistic account of agency, if we must

always presuppose agency in order to make the necessary distinction between purely physicodynamic events and formal or functional "processes" (as he puts it)?[124]

Now, Kauffman (2000, 2004), at least, is clearly cognizant of these difficulties, and cannot be accused of simply begging the important questions. For example, following his discussion of the minimum conditions necessary to describe a work cycle in a living cell, he writes: "I said we have no theory of organization, but I have the deep suspicion that this reciprocal linking of work and constraints on the release of energy that constitutes work is part of that theory. If so, notice that this is not part of physics at present, nor of chemistry, nor of biology" (Kauffman, 2004; p. 660). Unfortunately, he does not go on to say where we ought to look for the missing theory.

Abel, I fear, would be unimpressed by Kauffman's efforts to elucidate a necessary connection between physical constraints and the notion of "work." Indeed, he goes on to say that "[w]ork must be defined and pursued formally. That definition must be related to other formalisms such as 'value,' 'economy,' 'usefuleness,' and 'efficiency'" (Abel, 2010; p. 20). This is correct, and the present dissertation has, I hope, been written in full cognizance of this point. Nor, I suspect, would Kauffman fundamentally disagree. He admits, after all, that there something crucially important missing from our current understanding of life.

[124] Abel stipulates that the word "process" ought to mean a controlled event. However, he views organisms as machines, and he would draw the line between the normative "processes" so defined and the actions of agents quite differently than I would (specifically, he would limit the latter to minds), so I do not follow him in this terminology. While Abel puts his finger directly on the crucial problem facing any theory of "self-organization," in my view the problems raised by regarding organisms as machines are even more severe than the problems facing self-organization theory.

But then Abel takes what to me seems a step too far. He writes: "Such formalisms arise only in the minds of agents" (ibid.; p. 20). Farther on, he expands on this point: "The purposeful selection of constraints, not the physicodynamic constraints themselves, constitutes controls. It is only when we manipulate initial conditions or purposefully steer iterations to achieve a desired experimental result that constraints can be considered controls" (ibid.; p. 25). But to identify biological control with mentality is a counsel of despair—and, moreover, one that is not fully warranted by the considerations he has advanced with respect to our problem. It would be warranted, were the following claim that he makes known certainly to be true: "No as-of-yet undiscovered law will ever be able to explain the highly informational organization of living organisms" (ibid.; p. 14). But it is not. Or, rather, the claim is ambiguous, and while on one reading it is very probably true, on the other reading its truth value is unknown at present. If by "physical law" we stipulate the meaning "happening under the constraint of the least-action principle," then we can be fairly certain that no such law could possibly account for living systems, because such a law would violate the Orthogonality Condition, which we have good reason to believe must be met by any physical system we would count as living. However, if by "physical law" we mean merely "happening under some physical constraint," then it is entirely possible that such an as-of-yet undiscovered law may exist.

In short, I believe that Abel's critical analyses are extremely salutary as a corrective to the too-easy slide from talk of physical constraints to talk of controls, and the too-easy invocation of such vague ideas as "dissipative structures" and "self-organization" to cover over the conceptual slide. But that does not mean we must submit to his counsel of despair. It just means that we must accept the challenge of explaining

how the transition from constraints to controls—from physics to life—might be possible by virtue of some physical principle that transcends ordinary physical law in the sense of least action. I will now turn to the task of attempting to say a few brief words about what it would take to meet this challenge head-on.

The principle we are looking for has something to do with the intuitive idea that living systems must act in order to preserve themselves in existence. That is, we need to be able to understand the coordination of events in the cell as the result of the cell's own activity, as opposed to a merely passive, energetically "downhill" slide according to known physical laws. It is the fact that organisms, and only organisms, must act in order to maintain themselves in existence as the kind of system that they are that accounts for our intuitions regarding the proper applicability of normative vocabulary to living systems as such. This fact is also the thing that differentiates organisms from machines. Intuitively, it is obvious that the difference between organisms and machines has to do with the fact that the normativity associated with organisms is somehow endogenously generated, while the normativity associated with machines is exogenous—imposed on the matter of the machine by an external, minded agent. But these intuitions regarding the activity and the endogenous origin of the normativity of living systems need to be connected up with real science. If it were possible to do that, then we might be able to find solutions to the Intrinsicality Problem, the Holism Problem, and the Activity Problem together in one and the same physical principle.

So, what is an organism, if not a machine?

4.6 In Search of the Natural Ground of Normative Agency

In the last section of this chapter, I would like to explore in a preliminary way another approach to the natural ground of adaptivity and normative agency—one that is on the surface quite different from, but in fact may be viewed as complementary to, that of the Autonomous Agency Theorists. Let us begin by reviewing certain well-established empirical results that highlight the physical nature of the cell.

First, it is already well established that the functional integration of the cell depends, at least in certain respects, upon physical, as opposed to biochemical, principles. One example is chemiosmotic coupling and the resulting transmembrane potential (Harold, 1986). As Lane (2010) has noted, the crucial role of chemiosmotic coupling in cellular functioning depends precisely on the fact that it is not just more chemistry. As he puts it (ibid.; pp. 8–9), "chemiosmotic coupling enables metabolism to escape the bounds of chemistry." Another example is the dynamical self-assembly and -disassembly of cellular components mentioned above. Kurakin (2009) argues that the empirical facts require a wholesale reevaluation of the mechanistic view of the cell. For example, he makes this point as follows, in connection with the spatial and temporal organization of cellular traffic:

> In this regard, it should be pointed out that, in reality, the internal resource distribution/ transport systems of biological organisms (at all scales) are not mechanistic pipes built according to a preconceived design, but dynamic and adaptive fluxes of energy/matter in themselves, shaped by both internal and external influences. And their main purpose is not to deliver resources and remove waste—that is the limited interpretation of the mechanistic paradigm—but to integrate energy/matter and space into one scale-free continuum of energy/matter circulation. (ibid.; p. 20)

Next, it is essential to keep in mind that due to extreme molecular crowding, cytoplasm has many of the physical characteristics of the "soft-matter" systems (gels,

liquid crystals, and the like) that are the object of study of condensed-matter physics.[125] Then, as Laughlin and coworkers (2000) point out, nonliving condensed-matter systems that are of the same nanometer-to-micrometer length scale as cells derive their physical properties from such principles as spontaneous symmetry breaking and phase transitions. Therefore, it seems reasonable to suppose that at least some of the distinctive properties of living systems may derive from similar principles.[126] In this connection, Pollack (2001, 2006) makes a number of specific proposals. For example, he has provided evidence that the functioning of such biological systems as secretory vesicles and muscle tissue can best be modeled on the basis, not of diffusion in an aqueous medium or of locally applied forces, but rather of global phase transitions deriving from the gel-like properties of organelles and of cytoplasm generally.[127] In addition, Ho (1997, 2003, 2008) stresses the importance of the active and directed buildup and drawdown of internal ("onboard") energy stores (whether in the form of ATP, chemiosmosis, or some other form of potential) for understanding the difference between organisms and nonliving physical systems, which are passively driven by ambient energy fluxes. She also hypothesizes a deep connection between the requirement for such onboard energy storage and the fundamentally oscillatory nature of biological functions noted previously.[128]

---

[125] See R.J. Ellis (2001), Keighron & Keating (2011), Luby-Phelps (2000), Minton (2001), Wheatley (2003), and Zhou et al. (2008); for a brief overview, see McNiven (2003).

[126] There is unfortunately no space here to give the phenomenon of spontaneous symmetry breaking the consideration it deserves; however, for its relevance to the generation of order in condensed matter, in general, see Anderson (1972), Blasone et al. (2011), Morrison (2006), and Newth & Finnigan (2006), while for its relevance to biology, in particular, see Ho (2008), Pessa (2006), and Sergi (2009).

[127] For further discussion of cytoplasm as a condensed-matter medium, see Nakagaki & Guy (2008), Pollack & Chin (2008), Pollack et al. (2006), and Shepherd (2006).

[128] As she puts it (Ho, 2003; p. 76): "The key to understanding the thermodynamics of organisms is therefore, neither energy flow nor energy dissipation, but energy storage under energy flow."

Moreover, Frauenfelder and coworkers (1999, 2009) emphasize that the functional adaptivity of individual protein molecules is traceable to their highly dynamic and labile (metastable) nature, which in turn derives from the fact that they exist in a state of "frustration" (resulting from myriad competing self-interactions) that precludes relaxation into a clearly defined minimum-energy state, but rather results in a degenerate ground state. In this physical sense (as opposed to the biological sense mentioned above), "degeneracy" refers to the existence of an ensemble of nearly energetically equivalent minimum-energy states, known as "conformational substates," or "conformers."[129] The existence of a diversity of conformational substates may be at the bottom of such recently discovered phenomena as enzyme non-specificity ("promiscuity" or "moonlighting") and the functionality of certain enzymes that do not achieve a folded native state ("intrinsically unstructured [or disordered] proteins")—discoveries which have been declared to "shake this paradigm [genetic determinism] to its roots" and to "shift the explanation of biological specificity from the molecular to the cellular level" (Kupiec, 2009; p. 49).[130] In general, we may say that the physical properties of frustration and energy degeneracy make proteins the inherently dynamic molecules that they are—molecules capable of manifold forms of physical coupling to each other, to various ligands, and to the adjacent solvent.[131] In short, this inherent dynamism is the key to protein functionality. Enzymes now appear much more like active participants in the life

[129] For further discussion of the evidence relating to the dynamics of protein functioning, see Eisenmesser et al. (2005), Henzler-Wildman & Kern (2007), Tokuriki & Tawfik (2009), Vendruscolo & Dobson (2006), and Wolynes (2008).

[130] See Gsponer & Babu (2009), Khersonsky & Tawfik (2010), Uversky (2010), and Yadid et al. (2010). For discussion, see Kurakin (2005, 2009).

[131] On the so-called "slaving" of protein motions to solvent fluctuations, see Fenimore et al. (2002). For full details of our current dynamic view of proteins, see Frauenfelder (2010a); for a succinct statement of this view, and how it was arrived at, see Frauenfelder (2010b).

of the cell than they used to do according to the passive, mechanical, "lock-and-key" model that prevailed formerly.

Finally, Freeman and Vitiello (2006, 2008) apply the concepts of spontaneous symmetry breaking and phase transitions to modeling the large-scale synchronization of nerve-cell assemblies as a key to understanding brain functioning. In short, a number of different research programs are currently under way whose shared aim is to understand biological systems as capable of endogenous activity by virtue of the physical principles associated with various states of condensed matter. All of these analyses provide us with a picture of the cell that is far more active than normally supposed, and one whose activity derives entirely from endogenous physical principles.

There are two major objections to this line of thinking. First, Tegmark (2000) and others have pointed out that various physical features of living systems (especially temperature) ought to preclude quantum coherence effects. The reason is quantum decoherence.[132] Most of the well-known quantum coherence phenomena (superfluidity, superconductivity) occur at temperatures close to absolute zero. Even if similar phenomena could be created in a biological system, such as a protein, by pumping sufficient energy into it, shielding the resulting coherent vibrational modes from interaction with the ambient cytoplasm, which would be at room temperature, would seem to be very difficult if not impossible. This would mean that any quantum coherence effects that existed would be nearly instantaneously destroyed.

[132] For a review, see Stamp (2006); for full details, see Schlosshauer (2007).

There are several possible replies to this worry. First, Mavromatos (2010) has reviewed a number of proposals by means of which effective shielding against decoherence might be effected in biological microenvironments even at room temperature. Second, Aspuru-Guzik and coworkers have observed that partial decoherence may actually enhance quantum transport effects in biological materials, observing that (Rebentrost et al., 2009; p. 2) "at intermediate noise levels coherence and decoherence can collaborate to produce highly efficient transport."[133] Third, Vitiello and coworkers argue that decoherence is a sign of the breakdown of quantum mechanics as a theory, and of the need to use quantum field theory (QFT) in its stead.[134] They point out that QFT is required to explain such condensed-matter phenomena as phase transitions, and thus is the correct theory for explaining the properties of macroscopic systems in general. Moreover, in the sort of effective field theories that one finds in condensed-matter physics, long-range correlations may come into being by such means as spontaneous symmetry breaking, and not solely through the mechanisms of superposition or entanglement.[135] Finally, there is the fact that numerous claims have recently been published of experimental detection of quantum coherence in several biological systems, especially the "antenna" proteins of the light-harvesting complexes of various photosynthetic bacteria and marine algae.[136] These experiments will, of course, have to be replicated in future, but Parson (2007; p. 1439) has already gone so far as to say that

[133] See, also, Plenio & Huelga (2008).

[134] Alfinito et al. (2001).

[135] For further discussion, see Ho (2008), Pessa (2006), and Sergi (2009).

[136] See, e.g., Collini et al. (2010), Engel et al. (2007), Ishizaki & Fleming (2009), Lee et al. (2007), Panitchayangkoon et al. (2010), and Sarovar et al. (2010).

"[i]t seems clear . . . that a complete description of energy migration in photosynthetic complexes will have to include electronic coherence." If such claims are indeed substantiated, then quantum coherence phenomena in biological systems are a reality, whether or not we understand they are possible.[137] It is perhaps not surprising that predictions of the impossibility of such phenomena should fail, seeing that they have been based on knowledge of very different and very much simpler systems. As Stamp (2006; p. 490) has noted: "One should beware of general theorems on decoherence rates for large systems, since they usually make very restrictive (and unrealistic) assumptions about the structure of many-body states." At any rate, what is certain is that the recent experimental work on quantum coherence phenomena in biological systems poses an enormously exciting theoretical challenge for the future.[138]

Whatever one makes of the possible responses to the first objection, there is no doubt that the objection itself is a weighty one. However, the second objection is still more serious. It is this: that nothing I have yet said actually reaches to the heart of the problem. Even if one granted for the sake of argument that quantum coherence existed in living matter, and even if its theoretical underpinnings were fully worked out, we would still not be appreciably better off than we were before. The reason is that coherence per se is not enough to explain adaptivity. In this respect, it is far from clear that quantum coherence would hold out any more promise of such an explanation than do the forms of coherence arising from nonlinear dynamic and nonequilibrium thermodynamic

[137] Another area in which there is experimental evidence of quantum effects (specifically quantum tunneling) in biological materials is protein function; see Gray & Winkler (2003), Hammes-Schiffer (2006), Masgrau et al. (2006), Z.D. Nagel & Klinman (2006), and Sutcliffe & Scrutton (2000a, 2000b).

[138] For further discussion of these and related issues, see Ball (2011), P.C.W. Davies (2004), Leggett (2002), and Wilde et al. (2010).

considerations. While coherence of either the classical or the quantum sort surely represents a step in the right direction toward solving the Intrinsicality, the Holism, and the Activity Problems, it is simply not enough. The reason is that all of the physical ideas mentioned so far are ultimately reducible to least action, and we have seen that no principle of that sort can possibly provide the explanation we are seeking due to the Orthogonality Condition that must also be met.

It is here that we finally come face to face with the heart of the problem—the normativity implicit in adaptivity. Adaptivity is not just coherent activity; it is coherent activity coordinated under the constraint of viability—i.e., controlled, or regulated, coherent activity. Therefore, what we ultimately seek is a physical principle capable of explaining not just the intrinsicality, the holism, and the activity of biological functioning, but also precisely the normativity in normative agency.

Here, I think we are forced to admit that we have reached a dead end. It may still be possible to provide, in general terms, certain criteria that a way forward would have to meet. Thus, for example, Kitano (2007; p. 3) has said that: "The key issue is whether it is possible to find a formalism in which robustness and its trade-offs could be defined so that robustness is a conserved quantity." Similarly, Stelling et al. (2004; p. 681) speak of the mathematical investigation of network topology in systems biology as "somewhat reminiscent of the use of conservation laws in physics," while Chauvet (2004; p. 250) has postulated the existence of a "principle of vital coherence, or the conservation of life, in biology." However, these ruminations are not very helpful. At present, we simply have no idea what concrete form such vague ideas would need to take in order to be of any real

scientific value. Therefore, I am unable to provide a satisfactory resolution to the Ground Problem.

However, it is enough for my purposes in this chapter that the Autonomous Agency Theorists are pursuing their project, and the other ideas canvassed in this section are being discussed. While none of it adds up to a coherent account of adaptivity, or of normative agency, nevertheless, it does show, I think, that the existence of a physical principle underpinning agency in the sui generis dynamics of the living state of matter is not simply inconceivable. And if a ground of normative agency in scientific terms is at least conceivable, then, in light of the other considerations advanced earlier in Chapters 2 and 3, I think I am justified in proposing TRB as a "live option" for consideration alongside the various reductionist and eliminativist approaches to teleology. And that is all that I have set out here to do.

CHAPTER 5:

CONCLUSION

The time has come for stock-taking. What has been shown in this dissertation? By means of what arguments were these results arrived at? And what is their larger significance?

This dissertation has explored in broad terms what it would mean to adopt a realistic attitude towards teleology in biology, in the immanent sense of goal-seeking or purposive behavior. Furthermore, it has shown that it is rationally defensible to regard teleology as a real property of living things as such, and no mere "projection" of a human conceptual construct upon objects that lack any property objectively corresponding to that construct. In other words, my principal claim is that teleology may be rationally regarded as a real power or capacity possessed by all organisms—one that is on an ontological par with such other biological powers as locomotion, phagocytosis, photosynthesis, and so forth. The overall argument for this principal claim, which I sometimes refer to for convenience as "Teleological Realism in Biology" (**TRB**), has proceeded in three main stages corresponding to the three main academic disciplines dealing with the problems of teleology and normativity.

In stage one (Chapter 2), the argument proceeded primarily by means of reflection upon the way in which we apply the notions of teleology and related concepts. This discussion drew on the literature of the philosophy of action. Specifically, it was shown

that teleology is only one among a number of closely related "elementary normative concepts" (including value, need, well-being, and others). Moreover, it was observed that the elementary normative concepts form a network of mutual implication—that is, they come as a "package deal," and do not make sense in isolation from one another. Furthermore, the main feature that seems to unite them is the fact that they are all involved in one way or another with the concept of action in the normative sense, and hence with agency. In this way, teleology can be seen to be best understood by being brought under the wider concept of normative agency. Finally, it was noted that several of the elementary normative concepts, considered individually, are clearly properly ascribable to organisms as such. It follows, then, that normative agency ought to be properly ascribable to organisms as such, as well. However, it was acknowledged that this conclusion may appear counterintuitive to many, to the point of justifying the outright rejection of the foregoing reasoning. In order to mitigate this sense of paradox, a number of empirical examples were considered, in which we clearly do apply the concept of agency in a way that is consistent with its ascription to organisms as such, including even single-celled creatures. The conclusion of this chapter was that a living system ought to be regarded as a physical system that is under a normative requirement to act in order to preserve itself in existence as the sort of physical system that it is. In this way, **TRB** can be seen to be a "live option" on the conceptual landscape.

The next stage of the overall argument (Chapter 3) attempted to counter the objection that **TRB** is highly implausible, because teleology has already been eliminated from our picture of nature by the advance of scientific knowledge. In particular, so the objection goes, teleology has been "reduced" (i.e., a theoretically and empirically

adequate account of the putatively teleological features of life has been given in wholly non-teleological terms) through the discovery of the microscopic material "mechanisms" underlying biological powers and capacities—that is, in a phrase, through "molecular biology"—as well as by means of the theory of natural selection, which is said to explain how such mechanisms have come into existence. I attempted to cast doubt on this objection—that is, I tried to show that it is far from certain that "teleoreduction" has in fact been successfully accomplished—by means of several arguments, which are based on an examination of the explanatory structure of selection theory, together with a number of empirical considerations. This discussion drew primarily on the literature of the philosophy of biology.

In the third and final stage of the overall argument (Chapter 4), I attempted to further disarm the counterintuitive aura of the dissertation's principal claims—that teleology is an objective feature of living things as such, and that it is best understood as a feature of normative agency also conceived of as an inherent capacity of living things—by surveying and critiquing some contemporary attempts to give a positive scientific account of the physical principles underlying biological agency. Here, I examined a third disciplinary literature, this one lying at the interface between theoretical biology and cognitive science, focusing on the "self-organization" of "autonomous agents." Although I found these and various other accounts of physical principles supposed to provide a "natural ground" of normative agency to be deficient in various respects, I ended by arguing that they do at least support the claim that a scientific account of normative agency as a universal fundamental property of living matter is not simply inconceivable.

And that ought to be enough to help dispel whatever air of paradox may be attached to **TRB**.

In summary, Chapter 2 established the prima facie acceptability of **TRB**, Chapter 3 answered one of the main objections against it, and Chapter 4 showed that it has a right to be considered a "live option" from a scientific point of view.

Although I have not previously discussed the wider implications of the principal claims of this dissertation, I am not unmindful of the manifold connections between **TRB** and other important philosophical problems.

For example, if future empirical research were to confirm the speculation advanced here that the normative agency inherent in all living things is a manifestation of physical principles inherent in a particular phase of condensed matter ("the living state of matter"), then we would have a principled basis for distinguishing between organisms that are really alive and machines that merely simulate various aspects of living. This result would be of interest to workers in such fields as Artificial Intelligence, Robotics, and Artificial Life, as well as to philosophers. It would not mean that life could not be created artificially, of course; it might well turn out to be possible to construct various forms of life artificially, once the physical principles underlying that phase of matter were better understood. It would mean, however, that no artificially constructed system ought to count as being "alive" that did not recreate (and not merely simulate) the real material conditions required for embodying the physical principles that give rise to life.

Another traditional philosophical problem to which **TRB** would be relevant is the mind-body relation. While nothing I have said in this dissertation will be of direct interest

to those primarily concerned with the so-called "Hard Problem" of conscious experience (on the contrary, I have been careful to abstract away from this issue in posing my own problems), nevertheless **TRB** should be of interest to those concerned with the mind-body problem in a more general sense. This is for two reasons.

First, the brain is above all a biological organ, so that any fundamental change in our understanding of life is bound to have an impact upon our understanding of the brain. Second, and more specifically, **TRB** may be viewed as offering a solution to what Perovic (2007) has termed the "life-body problem"—that is, the question of the relation between life and its material substrate—which has a form very similar, if not identical, to the mind-body problem. As such, **TRB** will also be of potential interest to thinkers occupied with foundational issues in the philosophy of action (see Chapter 2, above) and perhaps even to philosophers concerned with the problem of free will (at least, to those such as Balaguer [2010] who construe that problem as largely empirical in nature). Admittedly, though, most of the issues of deepest concern to most philosophers occupied with the mind-body problem are bound up in one way or another with the Hard Problem; therefore, TRB will perhaps be of limited interest to them, at best.

There is one area of traditional philosophical concern, however, to which I believe that **TRB** is directly and importantly relevant, and that is the problem of the "natural ground of normativity." In a sense, the entire dissertation (though especially Chapter 4) is a contribution to this topic. I have indicated this tangentially at various points along the way, but I have not yet said anything about how I think that **TRB** relates to more traditional ways of framing the issues here.

Here is how Cottingham (2006) has recently stated the problem:

> . . . there are certain remarkable properties that truth, beauty and goodness all share. In the first place, they are all what philosophers call normative concepts—they carry with them the sense of a requirement or a demand. The true is that which is worthy of belief—"to be believed"; the beautiful is that which is worthy of admiration; and the good is that which is worthy of choice. They all therefore seem to be rather "queer" properties (as the late Oxford philosopher John Mackie put it). They have this odd, magnetic aspect—they somehow have "to-be-pursuedess" built into them.
>
> Why is this odd? Well, it is a feature that seems incompatible with any purely naturalistic or scientific account of these properties; for it is not easy to see how a purely natural or empirically definable item could have this strange "normativity" or choice-worthiness somehow packed into it. So it starts to look as if thinking about these normative concepts is going to take us beyond the purely natural or empirical domain. (ibid.)

In this dissertation, I have tried to show that Cottingham is mistaken, and that acknowledging the objective existence of both teleology and normativity in a robustly realistic sense does not have to take us beyond the purely natural domain.

As I have noted several times, nothing I have said here rationally requires the normative eliminativist to change his view. To show that holding **TRB** is not just rationally permissible, but positively preferable to the eliminativist view would require a difficult examination of the empirical adequacy of molecular biology, understood within the context of the theory of natural selection, that lies well beyond the horizon of anything I have been able to undertake here. At the end of the day, I am making an empirical claim that will be either confirmed or refuted by the future development of biological science.

However, not everyone is willing to embrace the eliminativist vision of nature in which even we human beings are governed entirely by ordinary physical law, and teleology and normativity are illusions, or at best polite fictions. For those who are

unwilling to do this, I believe that **TRB** does have real relevance. Of course, nothing in this dissertation has anything directly to do with such distinctively human norms as truth, beauty, and (moral) goodness. Nevertheless, by showing how normative requirement in the most general sense may be understood as an objective property of living systems as such, **TRB** seems to me to be in a position to provide a metaphysical foundation for what one might refer to as a "reformed naturalism." And upon such a foundation we may begin the task of constructing a more adequate picture of human nature that both acknowledges the objective reality of the higher norms to which we alone are responsive and is at the same time unequivocally natural.

REFERENCES

Abel, David L. (2010) "Constraints vs. Controls," Open Cybernetics & Systemics Journal **4**: 14–27.

Abel, David L. and Jack T. Trevors (2006) "Self-Organization vs. Self-Ordering Events in Life-Origin Models," Physics of Life Reviews **3**: 211–228.

Agutter, Paul S. and Denys N. Wheatley (2007) About Life: Concepts in Modern Biology. Dordrecht, Netherlands: Springer.

Alberts, Bruce, Alexander Johnson, Julian Lewis, Martin Raff, Keith Roberts, and Peter Walter (2007) Molecular Biology of the Cell, 5th ed. New York: Garland Science.

Albrecht-Buehler, Guenter (2009) "Cell Intelligence" website. http://www.basic.northwestern.edu/g-buehler/FRAME.HTM (last updated 7/24/2009; accessed 11/4/2010)

Alfinito, Eleonora, Rosario G. Viglione, and Giuseppe Vitiello (2001) "The Decoherence Criterion," Modern Physics Letters **B** 15: 127–135.

Allen, Colin (2004) "Animal Pain," Noûs **38**: 617–643.

Allen, Colin and Marc Bekoff (1997) Species of Mind: The Philosophy and Biology of Cognitive Ethology. Cambridge, MA: Bradford Books/MIT Press.

Allen, Colin, Marc Bekoff, and George Lauder, eds. (1998) *Nature's Purposes: Analyses* of Function and Design in Biology. Cambridge, MA: Bradford Books/MIT Press.

Allen, Colin, Perry N. Fuchs, Adam Shriver, and Hilary D. Wilson (2005) "Deciphering Animal Pain," in Murat Aydede, ed., Pain: New Essays on Its Nature and the Methodology of Its Study. Cambridge, MA: Bradford Books/MIT Press, pp. 351–366.

Alvarez, Maria (2010) Kinds of Reasons: An Essay in the Philosophy of Action. Oxford: Oxford University Press.

Amundson, Ron (2005) The Changing Role of the Embryo in Evolutionary Thought: Roots of Evo-Devo. Cambridge: Cambridge University Press.

Amundson, Ron and George V. Lauder (1994) "Function without Purpose: The Uses of Causal Role Function in Evolutionary Biology," Biology and Philosophy **9**: 443–469.

Amzallag, G.N. and H.R. Lerner (1995) "Physiological Adaptation of Plants to Environmental Stresses," in Mohammad Pessarakli, ed., Handbook of Plant and Crop Physiology. New York: Marcel Dekker, pp. 557–576.

Anderson, Philip W. (1972) "More Is Different: Broken Symmetry and the Nature of the Hierarchical Structure of Science," Science **177**: 393–396.

Ariew, André and Richard C. Lewontin (2004) "The Confusions of Fitness," British Journal for the Philosophy of Science **55**: 347–363.

Ariew, André, Robert Cummins, and Mark Perlman, eds. (2002) Functions: New Essays in the Philosophy and Psychology of Biology. Oxford: Oxford University Press.

Auffray, Charles, Sandrine Imbeaud, Magali Roux-Rouquié, and Leroy Hood (2003) "Self-Organized Living Systems: Conjunction of a Stable Organization with Chaotic Fluctuations in Biological Space-Time," Philosophical Transactions of the Royal Society of London A **361**: 1125–1139.

Bach-y-Rita, Paul (2004) "Tactile Sensory Substitution Studies," in Mihail C. Roco and Carlo D. Montemagno, eds., The Coevolution of Human Potential and Converging Technologies (=Annals of the New York Academy of Sciences, vol. 1013). New York: New York Academy of Sciences, pp. 83–91.

Badyaev, Alexander V. (2005) "Stress-Induced Variation in Evolution: From Behavioural Plasticity to Genetic Assimilation," Proceedings of the Royal Society of London B **272**: 877–886.

Badyaev, Alexander V. (2009) "Evolutionary Significance of Phenotypic Accommodation in Novel Environments: An Empirical Test of the Baldwin Effect," Philosophical Transactions of the Royal Society of London B **364**: 1125–1141.

Badyaev, Alexander V. (2010) "The Beak of the Other Finch: Coevolution of Genetic Covariance Structure and Developmental Modularity during Adaptive Evolution," Philosophical Transactions of the Royal Society of London B **365**: 1111–1126.

Badyaev, Alexander V. (2011) "Origin of the Fittest: Link between Emergent Variation and Evolutionary Change as a Critical Question in Evolutionary Biology," Proceedings of the Royal Society of London B, published online before print, doi:10.1098/rspb.2011.0548.

Balaguer, Mark (2010) Free Will as an Open Scientific Problem. Cambridge, MA: Bradford Books/MIT Press.

Balcombe, Jonathan (2010) Second Nature: The Inner Lives of Animals. New York: Palgrave Macmillan.

Ball, Philip (2011) "The Dawn of Quantum Biology," Nature **474**: 272–274.

Baluška, František and Stefano Mancuso, eds. (2009) Signaling in Plants. Berlin: Springer.

Barandiaran, Xabier and Álvaro Moreno (2008) "Adaptivity: From Metabolism to Behavior," Adaptive Behavior **16**: 325–344.

Barandiaran, Xabier E., Ezequiel Di Paolo, and Marieke Rohde (2009) "Defining Agency: Individuality, Normativity, Asymmetry, and Spatio-temporality in Action," Adpative Behavior **17**: 367–386.

Barbieri, Marcello (2003) The Organic Codes: An Introduction to Semantic Biology. Cambridge: Cambridge Univesity Press.

Barbieri, Marcello (2007) "Is the Cell a Semiotic System?," in idem, ed., Introduction to Biosemiotics: The New Biological Synthesis. Dordrecht, Netherlands: Springer, pp. 179–208.

Barlow, Peter W. (2009) "Reflections on 'Plant Neurobiology,'" Biosystems **92**: 132–147.

Battail, Gérard (2009) "Living Versus Inanimate: The Information Border," Biosemiotics **2**: 321–341.

Bechtel, William (2007) "Biological Mechanisms: Organized to Maintain Autonomy," in Fred C. Boogerd, Frank J. Bruggeman, Jan-Hendrik S. Hofmeyr, and Hans V. Westerhoff, eds., Systems Biology: Philosophical Foundations. Amsterdam: Elsevier, pp. 269–302.

Beckerman, Martin (2005) Molecular and Cellular Signaling. New York: Springer/AIP Press.

Bedau, Mark A. (1990) "Against Mentalism in Teleology," American Philosophical Quarterly **27**: 61–70.

Bedau, Mark A. (1992a) "Goal-Directed Systems and the Good," Monist **75**: 34–51.

Bedau, Mark A. (1992b) "Where's the Good in Teleology?," Philosophy and Phenomenological Research, 1992, 52: 781–805. (Reprinted in Allen et al. 1998.)

Bedau, Mark A. (1993) "Naturalism and Teleology," in Steven J. Wagner and Richard Warner, eds., Naturalism: A Critical Appraisal. Notre Dame, IN: University of Notre Dame Press, pp. 23–51.

Bedau, Mark A. (1996) "The Nature of Life," in Margaret A. Boden, ed., The Philosophy of Artificial Life. Oxford: University of Oxford Press, pp. 332–357.

Bedau, Mark A. and Carol E. Cleland, eds. (2010) The Nature of Life: Classical and Contemporary Perspectives from Philosophy and Science. Cambridge: Cambridge University Press.

Bedau, Mark A. and Paul Humphreys, eds. (2008) Emergence: Contemporary Readings in Philosophy and Science. Cambridge, MA: Bradford Books/MIT Press.

Bejan, Adrian and James H. Marden (2006) "Constructing Animal Locomotion from New Thermodynamics Theory," American Scientist **94**: 342–349.

Bekoff, Marc, Colin Allen, and Gordon M. Burghardt, eds. (2002) The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition. Cambridge, MA: Bradford Books/MIT Press.

Beloussov, Lev V. (2009) "Dialogue," in Joseph Seckbach and Richard Gordon, eds., Divine Action and Natural Selection: Science, Faith, and Evolution. Hackensack, NJ: World Scientific, pp. 45–46. (Commentary on Edward T. Oakes, "The Enigma of Final Causality: Biological Causality in Aristotle and Neo-Darwinism, ibid., pp. 31–44.)

Ben-Jacob, Eshel (2009a) "Learning from Bacteria about Natural Information Processing," in Günther Witzany, ed., Natural Genetic Engineering and Natural Genome Editing (=Annals of the New York Academy of Sciences, vol. 1178). Boston, MA: Blackwell Publishers on behalf of the New York Academy of Sciences, pp. 78–90.

Ben-Jacob, Eshel (2009b) "Bacterial Complexity: More Is Different on All Levels," in Shigetada Nakanishi, Ryoichiro Kageyama, and Dai Watanabe, eds., Systems Biology: The Challenge of Complexity. Tokyo: Springer, pp. 25–35.

Ben-Jacob, Eshel and Herbert Levine (2006) "Self-Engineering Capabilities of Bacteria," Journal of the Royal Society Interface **3**: 197–214.

Bentolila, Simone (2005) "'Live Memory' of the Cell, the Other Hereditary Memory of Living Systems," BioSystems **80**: 251–261.

Bermúdez, José Luis (2003) Thinking without Words. Oxford: Oxford University Press.

Beurton, Peter J., Raphael Falk, and Hans-Jörg Rheinberger, eds. (2000) The Concept of the Gene in Development and Evolution: Historical and Epistemological Perspectives. Cambridge: Cambridge University Press.

Biesecker, David (1999) "The Importance of Being Erroneous: Prospects for Animal Intentionality," Philosophical Topics **27**(1), Spring: 281–308.

Bird, Adrian (2007) "Perceptions of Epigenetics," Nature **447**: 396–398.

Blasone, Massimo, Petr Jizba, and Giuseppe Vitiello (2011) Quantum Field Theory and Its Macroscopic Manifestations: Boson Condensation, Ordered Patterns, and Topological Defects. London: Imperial College Press.

Boden, Margaret A. (1999) "Is Metabolism Necessary?," British Journal for the Philosophy of Science **50**: 231–248.

Boniolo, Giovanni (2003) "Biology without Information," History and Philosophy of the Life Sciences **25**: 255–273.

Boorse, Christopher (1976) "Wright on Functions," Philosophical Review **85**: 70–86.

Boorse, Christopher (2002) "A Rebuttal on Functions," in André Ariew, Robert Cummins, and Mark Perlman, eds., Functions: New Essays in the Philosophy and Psychology of Biology. Oxford: Oxford University Press, pp. 63–112.

Bouchard, Frédéric (2011) "Darwinism without Populations: A More Inclusive Understanding of the 'Survival of the Fittest,'" Studies in History and Philosophy of Biological and Biomedical Sciences **42**: 106–114.

Bouchard, Frédéric and Alex Rosenberg (2004) "Fitness, Probability and the Principles of Natural Selection," British Journal for the Philosophy of Science **55**: 693–712.

Boyle, Matthew and Douglas Lavin (2010) "Goodness and Desire," in Sergio Tenenbaum, ed., Desire, Practical Reason, and the Good. Oxford: Oxford University Press, pp. 161–201.

Braithwaite, Victoria (2010) Do Fish Feel Pain? Oxford: Oxford University Press.

Brandon, Robert N. (1990) Adaptation and Environment. Princeton, NJ: Princeton University Press.

Brandon, Robert N. (2006) "The Principle of Drift: Biology's First Law," Journal of Philosophy **103**: 319–335.

Brandon, Robert N. and Grant Ramsey (2007) "What's Wrong with the Emergentist Statistical Interpretation of Natural Selection and Random Drift?," in David L. Hull and Michael Ruse, eds., The Cambridge Companion to the Philosophy of Biology. Cambridge: Cambridge University Press, pp. 66–84.

Bray, Dennis (2009) Wetware: A Computer in Every Living Cell. New Haven, CT: Yale University Press.

Brembs, Björn (2010) "Towards a Scientific Concept of Free Will as a Biological Trait: Spontaneous Actions and Decision-Making in Invertebrates," Proceedings of the Royal Society of London B, **278**: 930–939.

Brenner, Eric D., Rainer Stahlberg, Stefano Mancuso, Jorge Vivanco, František Baluška, and Elizabeth Van Volkenburgh (2006) "Plant Neurobiology: An Integrated View of Plant Signaling," Trends in Plant Science **11**: 413–419.

Bruggeman, Frank J., Hans V. Westerhoff, and Fred C. Boogerd (2002) "BioComplexity: A Pluralistic Research Strategy Is Necessary for a Mechanistic Explanation of the 'Live' State," Philosophical Psychology **15**: 411–440.

Bruni, Luis Emilio (2007) "Cellular Semiotics and Signal Transduction," in Marcello Barbieri, ed., Introduction to Biosemiotics: The New Biological Synthesis. Dordrecht, Netherlands: Springer, pp. 365–407.

Brunswik, Egon (1957) "Scope and Aspects of the Cognitive Problem," in Jerome Bruner, Leon Festinger, and Fritz Heider, eds., Contemporary Approaches to Cognition. Cambridge, MA: Harvard University Press, pp. 5–31.

Buller, David J. (1999a) "Natural Teleology," in idem, ed., Function, Selection, and Design. Albany: State University of New York Press, pp. 1–27.

Buller, David J., ed. (1999b) Function, Selection, and Design. Albany: State University of New York Press.

Bullwinkel, Jörn, Bettina Baron-Lühr, Anja Lüdemann, Claudia Wohlenberg, Johannes Gerdes, and Thomas Scholzen (2006) "Ki-67 Protein Is Associated with Ribosomal RNA Transcription in Quiescent and Proliferating Cells," Journal of Cellular Physiology **206**: 624–635.

Burge, Tyler (2003) "Perceptual Entitlement," Philosophy and Phenomenological Research **67**: 503–548.

Burge, Tyler (2009) "Primitive Agency and Natural Norms," Philosophy and Phenomenological Research **79**: 251–278.

Burge, Tyler (2010) Origins of Objectivity. Oxford: Clarendon Press.

Burks, A.W. (1988) "Teleology and Logical Mechanism," Synthese **76**: 333–370.

Cabanac, Michel, Arnaud J. Cabanac, and André Parent (2009) "The Emergence of Consciousness in Phylogeny," Behavioural Brain Research **198**: 267–272.

Callebaut, Werner, Gerd B. Müller, and Stuart A. Newman (2007) "The Organismic Systems Approach: Evo-Devo and the Streamlining of the Naturalistic Agenda," in Roger Sansom and Robert B. Brandom, eds., Integrating Evolution and Development: From Theory to Practice. Cambridge, MA: Bradford Books/MIT Press, pp. 25–92.

Calvo, Paco and Fred Keijzer (2009) "Cognition in Plants," in František Baluška, ed., Plant-Environment Interactions: From Sensory Plant Biology to Active Plant Behavior. Berlin: Springer, pp. 247–266.

Camazine, Scott, Jean-Louis Deneubourg, Nigel R. Franks, James Sneyd, Guy Theraulaz, and Eric Bonabeau (2001) Self-Organization in Biological Systems. Princeton: Princeton University Press.

Cameron, Richard (2004) "How to Be a Realist about Sui Generis Teleology Yet Feel at Home in the Twenty-First Century," Monist **87**: 72–95.

Campbell, Iain D. (2008) "Croonian Lecture 2006: Structure of the Living Cell," Philosophical Transactions of the Royal Society of London B **363**: 2379–2391.

Campbell, Richard (2009) "A Process-Based Model for an Interactive Ontology," Synthese **166**: 453–477.

Carroll, Sean B. (2005) "Evolution at Two Levels: On Genes and Form," PLoS Biology **3**(7): article e245.

Carruthers, Peter (1989) "Brute Experience," Journal of Philosophy **86**: 258–269.

Chauvet, Gilbert (2004) The Mathematical Nature of the Living World: The Power of Integration. Singapore: World Scientific.

Chemero, Anthony (2009) Radical Embodied Cognitive Science. Cambridge, MA: Bradford Books/MIT Press.

Christensen, Wayne D. and Mark H. Bickhard (2002) "The Process Dynamics of Normative Function," Monist **85**: 3–28.

Christensen, Wayne D. and Clifford A. Hooker (2001) "Self-Directed Agents," Canadian Journal of Philosophy, **Supp. Vol. 27**: 19–51.

Churchland, Paul M. (2007) Neurophilosophy at Work. Cambridge: Cambridge University Press.

Clayton, Philip and Paul Davies, eds. (2006) The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion. Oxford: Oxford University Press.

Cleland, Carol E. and Christopher F. Chyba (2002) "Defining 'Life,'" Origins of Life and Evolution of the Biosphere **32**: 387–393.

Collini, Elisabetta, Cathy Y. Wong, Krystyna E. Wilk, Paul M.G. Curmi, Paul Brumer, and Gregory D. Scholes (2010) "Coherently Wired Light_harvesting in Photosynthetic Marine Algae at Ambient Temperature," Nature **463**: 644–648.

Conway Morris, Simon (2003) *Life's Solution: Inevitable Humans in a Lonely Universe.* Cambridge: Cambridge University Press.

Cooper, John M. (2004) "Hypothetical Necessity," in idem, Knowledge, Nature, and the Good: Essays on Ancient Philosophy. Princeton, NJ: Princeton University Press, pp. 130–147. (Originally published in Allan Gotthelf, ed., Aristotle on Nature and Living Things: Philosophical and Historical Studies. Pittsburgh, PA/Bristol, UK: Mathesis Publications/Bristol Classical Press, 1985, pp. 151–167.)

Cope, Edward Drinker (1887) The Origin of the Fittest: Essays on Evolution. New York: D. Appleton and Co.

Cornish-Bowden, Athel, María Luz Cárdenas, Juan-Carlos Letelier, and Jorge Soto-Andrade (2007) "Beyond Reductionism: Metabolic Circularity as a Guiding Vision for a Real Biology of Systems," Proteomics **7**: 839–845.

Corradini, Antonella and Timothy O'Connor, eds. (2010) Emergence in Science and Philosophy. New York: Routledge.

Cottingham, John (2006) "Philosophers Are Finding Fresh Meaning in Truth, Beauty and Goodness," The Times (London), June 17: "Features" section, p. 68.

Crick, Francis (1966) Of Molecules and Men. Seattle: University of Washington Press.

Csermely, Péter (2006) Weak Links: The Universal Key to the Stability of Networks and Complex Systems. Berlin: Springer.

Cummins, Robert (1975) "Functional Analysis," Journal of Philosophy **72**: 741–765. (Reprinted in Allen et al. 1998, and in Buller 1999b.)

Cummins, Robert (2002) "Neo-Teleology," in André Ariew, Robert Cummins, and Mark Perlman, eds., Functions: New Essays in the Philosophy and Psychology of Biology. Oxford: Oxford University Press, pp. 157–172.

Davidson, Donald (1987) "Knowing One's Own Mind," Proceedings and Addresses of the American Philosophical Association **60**(3): 441–458.

Davidson, Donald (2001a) "Mental Events," in idem, Essays on Actions and Events, 2nd ed. Oxford: Clarendon Press, pp. 207–224. (Originally published in Lawrence Foster and J.W. Swanson, eds., Experience and Theory. Amherst: University of Massachusetts Press, 1970, pp. 79–101.)

Davidson, Donald (2001b) "Rational Animals," in idem, Subjective, Intersubjective, Objective. Oxford: Clarendon Press, pp. 95–105. (Originally published in Dialectica, 1982, **36**: 317–327.)

Davies, Paul C.W. (2004) "Does Quantum Mechanics Play a Non-Trivial Role in Life?," Biosystems **78**: 69–79.

Davies, Paul Sheldon (2009) Subj*ects of the World: Darwin's Rhetoric and the Study of* Agency in Nature. Chicago: University of Chicago Press.

Dawkins, Marian Stamp (1993) Through Our Eyes Only? The Search for Animal Consciousness. Oxford: W.H. Freeman/Spektrum.

Deacon, Terrence W. (2010) "What Is Missing from Theories of Information?," in Paul Davies and Niels Henrik Gregersen, eds., Information and the Nature of Reality: From Physics to Metaphysics. Cambridge: Cambridge University Press, pp. 146–169.

De Caro, Mario and Alberto Voltolini (2010) "Is Liberal Naturalism Possible?," in Mario De Caro and David Macarthur, eds., Naturalism and Normativity. New York: Columbia University Press, pp. 69–86.

Delancey, Craig (2006) "Action, the Scientific Worldview, and Being-in-the-World," in Hubert L. Dreyfus and Mark A. Wrathall, eds., A Companion to Phenomenology and Existentialism. Malden, MA: Blackwell, pp. 356–376.

Delattre, Pierre (1986) "An Approach to the Notion of Finality according to the Concepts of Qualitative Dynamics," in S. Diner, D. Fargue, and G. Lochak, eds., Dynamical Systems: A Renewal of Mechanism; Centennial of George David Birkhoff. Singapore: World Scientific, pp. 149–154.

Dennett, Daniel C. (1987) The Intentional Stance. Cambridge, MA: Bradford Books/MIT Press.

Depew, David J. (2011) "Adaptation as Process: The Future of Darwinism and the Legacy of Theodosius Dobzhansky," Studies in History and Philosophy of Biological and Biomedical Sciences **42**: 89–98.

De Waal, Frans (1999) "Anthropomorphism and Anthropodenial: Consistency in Our Thinking about Humans and Other Animals," Philosophical Topics **27**(1), Spring: 255–280.

Di Paolo, Ezequiel A. (2005) "Autopoiesis, Adaptivity, Teleology, Agency," Phenomenology and the Cognitive Sciences **4**: 429–452.

Di Paolo, Ezequiel (2009) "Extended Life," Topoi **28**: 9–21.

Di Paolo, Ezequiel A. and Hiroyuki Iizuka (2008) "How (Not) to Model Autonomous Behaviour," BioSystems **91**: 409–423.

Di Paolo, Ezequiel A., Marieke Rohde, and Hanne De Jaegher (2010) "Horizons for the Enactive Mind: Values, Social Interactions, and Play," in John Stewart, Olivier Gapenne, and Ezequiel A. Di Paolo, eds., Enaction: Toward a New Paradigm for Cognitive Science. Cambridge, MA: Bradford Books/MIT Press, pp. 33–87.

Dretske, Fred I. (1981) Knowledge and the Flow of Information. Cambridge, MA: Bradford Books/MIT Press.

Dretske, Fred (1986) "Misrepresentation," in Radu J. Bogdan, ed., Belief: Form, Content, and Function. Oxford: Clarendon Press, pp. 17–36.

Dretske, Fred (2000) "Norms, History, and the Constitution of the Mental", in Fred Dretske, Perception, Knowledge and Belief: Selected Essays. Cambridge: Cambridge University Press, pp. 242–258.

Dretske, Fred (2006) "Representation, Teleosemantics, and the Problem of Self-Knowledge," in Graham Macdonald and David Papineau, eds., Teleosemantics. Oxford: Clarendon Press, pp. 69–84.

Dupré, John (2010) "Postgenomic Darwinism," in William Brown and Andrew C. Fabian, eds., Darwin. Cambridge: Cambridge University Press, pp. 150–200.

Dürr, Hans-Peter, Fritz-Albert Popp, and Wolfram Schommers, eds. (2002) What Is Life? Scientific Approaches and Philosophical Position. Singapore: World Scientific.

Edelman, Gerald M. and Joseph A. Gally (2001) "Degeneracy and Complexity of Biological Systems," Proceedings of the National Academy of Sciences, USA **98**: 13763–13768.

Eisenmesser, Elan Z., Oscar Millet, Wladimir Labeikovsky, Dmitry M. Korzhnev, Magnus Wolf-Watz, Daryl A. Bosco, Jack J. Skalicky, Lewis E. Kay, and Dorothee Kern (2005) "Intrinsic Dynamics of an Enzyme Underlies Catalysis," Nature **438**: 117–121.

El Hani, Charbel Niño, João Queiroz, and Claus Emmeche (2009) Genes, Information, and Semiosis. Tartu, Estonia: Tartu University Press.

Ellis, Ralph D. and Natika Newton (2010) How the Mind Uses the Brain (to Move the Body and Image the World). Chicago: Open Court Publishing Company.

Ellis, R. John (2001) "Macromolecular Crowding: Obvious but Underappreciated," Trends in Biochemical Sciences **26**: 597–604.

Engel, Gregory S., Tessa R. Calhoun, Elizabeth L. Read, Tae-Kyu Ahn, Tomáš Mančal, Yuan-Chung Cheng, Robert E. Blankenship, and Graham R. Fleming (2007) "Evidence for Wavelike Energy Transfer through Quantum Coherence in Photosynthetic Systems," Nature **446**: 782–786.

Enoch, David (2006) "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action," Philosophical Review **115**: 169–198.

Ereshefsky, Marc (2009) "Defining 'Health' and 'Disease,'" Studies in History and Philosophy of Biological and Biomedical Sciences **40**: 221–227.

Fenimore, Paul W., Hans Frauenfelder, Benjamin H. McMahon, and Fritz G. Parak (2002) "Slaving: Solvent Fluctuations Dominate Protein Dynamics and Functions," Proceedings of the National Academy of Sciences, USA **99**: 16047–16051.

Fine, Kit (2002) "The Question of Realism," in Andrea Bottani, Massimiliano Carrara, and Pierdaniele Giaretta, eds., Individuals, Essence and Identity: Themes of Analytic Metaphysics. Dordrecht, Netherlands: Kluwer Academic, pp. 3–48.

Finkelstein, David H. (2007) "Holism and Animal Minds," in Alice Crary, ed., Wittgenstein and the Moral Life: Essays in Honor of Cora Diamond. Cambridge, MA: Bradford Books/MIT Press, pp. 251–278.

Fodor, Jerry (2000) *The Mind Doesn't Work That Way: The Scope and Limits of* Computational Psychology. Cambridge, MA: Bradford Books/MIT Press.

Fontana, Walter and Leo W. Buss (1994) "'The Arrival of the Fittest': Toward a Theory of Biological Organization," Bulletin of Mathematical Biology **56**: 1–64.

Foot, Philippa (2001) Natural Goodness. Oxford: Clarendon Press.

Forber, Patrick (2005) "On the Explanatory Roles of Natural Selection," Biology and Philosophy **20**: 329–342.

Ford, Brian J. (2000) The Secret Language of Life: How Animals and Plants Feel and Communicate. New York: Fromm International.

Ford, Brian J. (2009) "On Intelligence in Cells: The Case for Whole Cell Biology," Interdisciplinary Science Reviews **34**: 350–365.

Frank, Till D., Michael J. Richardson, Stacy M. Lopresti-Goodman, and Michael T. Turvey (2009) "Order Parameter Dynamics of Body-Scaled Hysteresis and Mode Transitions in Grasping Behavior," Journal of Biological Physics **35**: 127–147.

Frankena, William K. (1976) "Obligation and Motivation in Recent Moral Philosophy," in idem, Perspectives on Morality: Essays by William Frankena, ed. Kenneth E. Goodpaster. Notre Dame, IN: University of Notre Dame Press, pp. 49–73.

Frankfurt, Harry G. (1997) "The Problem of Action," in Alfred R. Mele, ed., The Philosophy of Action. Oxford: Oxford University Press, pp. 42–52. (Originally published in American Philosophical Quarterly, 1978, **15**: 157–162.)

Frauenfelder, Hans (2010a) The Physics of Proteins: An Introduction to Biological Physics and Molecular Biophysics, edited by Shirley S. Chan and Winnie S. Chan. New York: Springer.

Frauenfelder, Hans (2010b) "Myoglobin as an Example of Protein Complexity," Chemical Physics **375**: 612–615.

Frauenfelder, Hans, Peter G. Wolynes, and Robert H. Austin (1999) "Biological Physics," Reviews of Modern Physics **71** (special issue): S419–S430.

Frauenfelder, Hans, Guo Chen, Joel Berendzen, Paul W. Fenimore, Helen Jansson, Benjamin H. McMahon, Izabela R. Stroe, Jan Swenson, and Robert D. Young (2009) "A Unified Model of Protein Dynamics," Proceedings of the National Academy of Sciences, USA **106**: 5129–5134.

Freeman, Walter J. (2001) How Brains Make Up Their Minds. New York: Columbia University Press.

Freeman, Walter J. and Giuseppe Vitiello (2006) "Nonlinear Brain Dynamics as Macroscopic Manifestation of Underlying Many-Body Field Dynamics," Physics of Life Reviews **3**: 93–118.

Freeman, Walter J. and Giuseppe Vitiello (2008) "The Dissipative Quantum Model of Brain and Laboratory Observations," in Ignazio Licata and Ammar Sakaji, eds., Physics of Emergence and Organization. Hackensack, NJ: World Scientific, pp. 233–252.

Garner, Richard T. (1990) "On the Genuine Queerness of Moral Properties and Facts," Australasian Journal of Philosophy **68**: 137–146.

Gayon, Jean, Christophe Malaterrre, Michel Morange, Florence Raulin-Cerceau, and Stéphane Tirard, eds. (2010) "Special issue: Defining Life," Origins of Life and Evolution of the Biosphere **40**(2): 119–244.

Gerhart, John and Marc Kirschner (2007) "The Theory of Facilitated Variation," Proceedings of the National Academy of Sciences, USA **104**: 8582–8589.

Gierasch, Lila M. and Anne Gershenson (2009) "Post-Reductionist Protein Science, or Putting Humpty Dumpty Back Together Again," Nature Chemical Biology **11**: 774–777.

Gilbert, Don and David Lloyd (2000) "The Living Cell: A Complex Autodynamic Multi-Oscillator System?," Cell Biology International **24**: 569–580.

Gillett, Carl (2007) "Understanding the New Reductionism: The Metaphysics of Science and Compositional Reduction," Journal of Philosophy **104**: 193–216.

Gilson, Etienne (2009) From Aristotle to Darwin and Back Again: A Journey in Final Causality, Species, and Evolution. San Francisco: Ignatius Press. (Originally published as *D'Aristote à Darwin et retour.* Paris: Vrin, 1971; English translation originally published by University of Notre Dame Press in 1984.)

Glass, Leon and Michael C. Mackey (1988) From Clocks to Chaos: The Rhythms of Life. Princeton, NJ: Princeton University Press.

Glennan, Stuart (2009) “Productivity, Relevance, and Natural Selection,” Biology and Philosophy **24**: 325–339.

Glock, Hans-Johann (2009) “Can Animals Act for Reasons?,” Inquiry **52**: 232–254.

Godfrey-Smith, Peter (2007) “Conditions for Evolution by Natural Selection,” Journal of Philosophy **104**: 489–516.

Goldbeter, Albert (1996) Biochemical Oscillations and Cellular Rhythms: The Molecular Bases of Periodic and Chaotic Behaviour. Cambridge: Cambridge University Press.

Gould, Stephen Jay (1989) Wonderful Life: The Burgess Shale and the Nature of History. New York: W.W. Norton.

Grandy, Walter T., Jr. (2008) Entropy and the Time Evolution of Macroscopic Systems. Oxford: Oxford University Press.

Gray, Harry B. and Jay R. Winkler (2003) “Electron Tunneling through Proteins,” Quarterly Reviews of Biophysics **36**: 341–372.

Griffin, Donald R. (2001) Animal Minds: Beyond Cognition to Consciousness, revised & expanded ed. Chicago: University of Chicago Press.

Griffiths, Paul E. (2001) “Genetic Information: A Metaphor in Search of a Theory,” Philosophy of Science **68**: 394–403.

Griffiths, Paul E. and Karola Stotz (2006) “Genes in the Postgenomic Era,” Theoretical Medicine and Bioethics **27**: 499–521.

Gsponer, Jörg and M. Madan Babu (2009) “The Rules of Disorder or Why Disorder Rules,” Progress in Biophysics and Molecular Biology **99**: 94–103.

Gunderson, Keith (1964) “Descartes, La Mettrie, Language, and Machines,” Philosophy **39**: 193–222.

Hacker, Peter M.S. (2007) Human Nature: The Categorial Framework. Malden, MA: Blackwell Publishing.

Hacker-Wright, John (2009) “What Is Natural about Foot’s Ethical Naturalism?,” Ratio (n.s.) **22**: 308–321.

Halley, Julianne D. and David A. Winkler (2008) “Consistent Concepts of Self-Organization and Self-Assembly,” Complexity **14**(2): 10–17.

Hammes-Schiffer, Sharon (2006) “Hydrogen Tunneling and Protein Motion in Enzyme Reactions,” Accounts of Chemical Research **39**: 93–100.

Hanna, Robert and Michelle Maiese (2009) Embodied Minds in Action. Oxford: Oxford University Press.

Harold, Franklin M. (1986) The Vital Force: A Study of Bioenergetics. San Francisco: W.H. Freeman & Company.

Harold, Franklin M. (2001) The Way of the Cell: Molecules, Organisms, and the Order of Life. Oxford: University of Oxford Press.

Harrison, Lionel G. (2011) The Shaping of Life: The Generation of Biological Pattern. Cambridge: Cambridge University Press.

Hauser, Marc D. (2000) Wild Minds: What Animals Really Think. New York: Owl Books/Henry Holt and Company.

Heisenberg, Martin (2009) "Is Free Will an Illusion?," Nature **459**: 164–165.

Heisenberg, Martin and Reinhard Wolf (1984) Vision in Drosophila: Genetics of Microbehavior. Berlin: Springer-Verlag.

Heisenberg, Martin, Reinhard Wolf, and Björn Brembs (2001) "Flexibility in a Single Behavioral Variable of Drosophila," Learning and Memory **8**: 1–10.

Helm, Bennett W. (2001) Emotional Reason: Deliberation, Motivation, and the Nature of Value. Cambridge: Cambridge University Press.

Henzler-Wildman, Katherine A. and Dorothee Kern (2007) "Dynamic Peronalities of Proteins," Nature **450**: 964–972.

Heschl, Adolf (1990) "L=C: A Simple Equation with Astonishing Consequences," Journal of Theoretical Biology **145**: 13–40.

Heyes, Cecilia and Anthony Dickinson (1990) "The Intentionality of Animal Action," Mind and Language **5**: 87–104.

Heyes, Cecilia and Ludwig Huber, eds. (2000) The Evolution of Cognition. Cambridge, MA: Bradford Books/MIT Press.

Ho, Mae-Wan (1997) "Towards a Theory of Organism," Integrative Physiological and Behavioral Science **32**: 343–363.

Ho, Mae-Wan (2003) "From 'Molecular Machines' to Coherent Organisms," in Francesco Masumeci, Larissa S. Brizhik, and Mae-Wan Ho, eds., Energy and Information Transfer in Biological Systems: How Physics Could Enrich Biological Understanding. Singapore: World Scientific, pp. 63–81.

Ho, Mae-Wan (2008) The Rainbow and the Worm: The Physics of Organisms, 3rd ed. Singapore: World Scientific.

Hofmeyr, Jan-Hendrik S. (2007) "The Biochemical Factory That Autonomously Fabricates Itself: A Systems Biological View of the Living Cell," in Fred C. Boogerd, Frank J. Bruggeman, Jan-Hendrik S. Hofmeyr, and Hans V. Westerhoff, eds., Systems Biology: Philosophical Foundations. Amsterdam: Elsevier, pp. 217–242.

Hohwy, Jakob and Jesper Kallestrup, eds. (2008) Being Reduced: New Essays on Reduction, Explanation, and Causation. New York: Oxford University Press.

Hooker, Clifford A. (2009a) "On Fundamental Implications of Systems and Synthetic Biology," in Pengcheng Fu and Sven Panke, eds., Systems Biology and Synthetic Biology. Hoboken, NJ: John Wiley & Sons, pp. 593–613.

Hooker, Clifford A. (2009b) "Interaction and the Bio-Cognitive Order," Synthese **166**: 513–546.

Hornsby, Jennifer (1997) Simple Mindedness: In Defense of Naïve Naturalism in the Philosophy of Mind. Cambridge, MA: Harvard University Press.

Hurley, Susan (2003) "Animal Action in the Space of Reasons," Mind and Language **18**: 231–256. (Reprinted with revisions as "Making Sense of Animals," in idem and Matthew Nudds, eds., Rational Animals? Oxford: Oxford University Press, pp. 139–171.)

Hurley, Susan and Matthew Nudds, eds. (2006) Rational Animals? Oxford: Oxford University Press.

Ishizaki, Akihito and Graham R. Fleming (2009) "Theoretical Examination of Quantum Coherence in a Photosynthetic System at Physiological Temperature," Proceedings of the National Academy of Sciences, USA **106**: 17255–17260.

Jablonka, Eva (2002) "Information: Its Interpretation, Its Inheritance, and Its Sharing," Philosophy of Science **69**: 578–605.

Jablonka, Eva and Gal Raz (2009) "Transgenerational Epigenetic Inheritance: Prevalence, Mechanisms, and Implications for the Study of Evolution," Quarterly Review of Biology **84**: 131–176.

Jackson, E. Atlee (2001) *Exploring Nature's Dynamics.* New York: Wiley-Interscience/John Wiley & Sons.

Jacobs, Jonathan (1986) "Teleology and Reduction in Biology," Biology and Philosophy **1**: 389–399.

Janicki, Susan M. and David L. Spector (2003) "Nuclear Choreography: Interpretations from Living Cells," Current Opinion in Cell Biology **15**: 149–157.

Jennings, Herbert Spencer (1962) Behavior of the Lower Organisms. Bloomington: Indiana University Press. (Originally published: New York: Columbia University Press, 1906)

Johnson, Brian R. and Sheung Kwan Lam (2010) "Self-Organization, Natural Selection, and Evolution: Cellular Hardware and Genetic Software," BioScience **60**: 879–885.

Juarrero, Alicia (1999) Dynamics in Action: Intentional Behavior as a Complex System. Cambridge, MA: Bradford Books/MIT Press.

Kaneko, Kunihiko (2006) Life: An Introduction to Complex Systems Biology. Berlin: Springer.

Kaplan, Jonathan (2008) "Evolutionary Innovations and Developmental Resources: From Stability to Variation and Back Again," Philosophy of Science **75**: 861–873.

Karsenti, Eric (2008) "Self-Organization in Cell Biology: A Brief History," Nature Reviews Molecular and Cell Biology **9**: 255–262.

Kauffman, Stuart A. (2000) Investigations. Oxford: Oxford University Press.

Kauffman, Stuart (2004) "Autonomous Agents," in John D. Barrow, Paul C.W. Davies, and Charles L. Harper, Jr., eds., Science and Ultimate Reality: Quantum Theory, Cosmology, and Complexity. Cambridge: Cambridge University Press, pp. 654–666.

Kauffman, Stuart and Philip Clayton (2006) "On Emergence, Agency, and Organization," Biology and Philosophy **21**: 501–521.

Kauffman, Stuart, Robert K. Logan, Robert Este, Randy Goebel, David Hobill, and Ilya Shmulevich (2008) "Propagating Organization: An Inquiry," Biology and Philosophy **23**: 27–45.

Keighron, Jacqueline D. and Christine D. Keating (2011) "Towards a Minimal Cytoplasm," in Pier Luigi Luisi and Pasquale Stano, eds., The Minimal Cell: The Biophysics of Cell Compartment and the Origin of Cell Functionality. Dordrecht, Netherlands: Springer, pp. 3–30.

Keijzer, Fred (2001) Representation and Behavior. Cambridge, MA: Bradford Books/MIT Press.

Kelso, J.A. Scott (1995) Dynamic Patterns: The Self-Organization of Brain and Behavior. Cambridge, MA: Bradford Books/MIT Press.

Kelso, J.A. Scott (2008) "An Essay on Understanding the Mind," Ecological Psychology **20**: 180–208.

Kempner, Ellis S. and Jay H. Miller (2003) "The Molecular Biology of Euglena gracilis. XV. Recovery from Centrifugation-Induced Stratification," Cell Motility and the Cytoskeleton **56**: 219–224.

Khersonsky, Olga and Dan S. Tawfik (2010) "Enzyme Promiscuity: A Mechanistic and Evolutionary Perspective," Annual Review of Biochemistry **79**: number 11.

Kim, Jaegwon (1998) Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation. Cambridge, MA: Bradford Books/MIT Press.

Kirschner, Marc (2010) "Cell Biology as a World View," Molecular Biology of the Cell **21**: 3803.

Kirschner, Marc and John Gerhart (1998) "Evolvability," Proceedings of the National Academy of Sciences, USA **95**: 8420–8427.

Kirschner, Marc W. and John C. Gerhart (2005) The Plausibility of Life: Resolving *Darwin's Dilemma.* New Haven, CT: Yale University Press.

Kirschner, Marc and Tim Mitchison (1986) "Beyond Self-Assembly: From Microtubules to Morphogenesis," Cell **45**: 329–342.

Kirschner, Marc, John Gerhart, and Tim Mitchison (2000) "Molecular 'Vitalism,'" Cell **100**: 79–88.

Kitano, Hiroaki (2007) "Towards a Theory of Biological Robustness," Molecular Systems Biology **3**: article number 137.

Kitcher, Philip (1993) "Function and Design," Midwest Studies in Philosophy **18**: 379–397. (Reprinted in Allen et al. 1998, and in Buller 1999b.)

Kolb, Vera M. (2007) "On the Applicability of the Aristotelian Principles to the Definition of Life," International Journal of Astrobiology **6**: 51–57.

Kolnai, Aurel (1980) "The Ghost of the Naturalistic Fallacy," Philosophy **55**: 5–16.

Kolnai, Aurel (2008) "Morality and Practice I: The Ambiguity of Good," in idem, Ethics, Value and Reality. New Brunswick, NJ: Transaction Publishers, pp. 63–94. (Originally published: Indianapolis: Hackett Publishing Co., 1978.)

Kondepudi, Dilip and Ilya Prigogine (1998) Modern Thermodynamics: From Heat Engines to Dissipative Structures. Chichester, UK: John Wiley & Sons.

Kondo, Shigeru and Takashi Miura (2010) "Reaction-Diffusion Model as a Framework for Understanding Biological Pattern Formation," Science **329**: 1616–1620.

Koonin, Eugene V. and Yuri I. Wolf (2009) "Is Evolution Darwinian or/and Lamarckian?," Biology Direct **4**: article number 42.

Korsgaard, Christine M. (2009) Self-Constitution: Agency, Identity, and Integrity. Oxford: Oxford University Press.

Kraut, Richard (2007) What Is Good and Why: The Ethics of Well-Being. Cambridge, MA: Harvard University Press.

Kupiec, Jean-Jacques (2010) "On the Lack of Specificity of Proteins and Its Consequences for a Theory of Biological Organization," Progress in Biophysics and Molecular Biology **102**: 45–52.

Kurakin, Alexei (2005) "Self-Organization versus Watchmaker: Stochastic Dynamics of Cellular Organization," Biological Chemistry **386**: 247–254.

Kurakin, Alexei (2009) "Scale-Free Flow of Life: On the Biology, Economics, and Physics of the Cell," Theoretical Biology and Medical Modelling **6**: number 6.

Kurzynski, Michal (2006) The Thermodynamic Machinery of Life. Berlin: Springer.

Lane, Nick (2010) "Chance or Necessity? Bioenergetics and the Probability of Life," Journal of Cosmology **10**(August/September): 3286–3304.

Lange, Marc (1996) "Life, 'Artificial Life,' and Scientific Explanation," Philosophy of Science **63**: 225–244.

Laubichler, Manfred D. and Jane Maienschein, eds. (2009) Form and Function in Developmental Evolution. Cambridge: Cambridge University Press.

Laughlin, R.B., D. Pines, J. Schmalian, B.P. Stojkovic, and P. Wolynes (2000) "The Middle Way," Proceedings of the National Academy of Sciences, USA **97**: 32–37.

Lee, Hohjai, Yuan-Chung Cheng, and Graham R. Fleming (2007) "Coherence Dynamics in Photosynthesis: Protein Protection of Excitonic Coherence," Science **316**: 1462–1465.

Leggett, Anthony J. (2002) "Testing the Limits of Quantum Mechanics: Motivation, State of Play, Prospects," Journal of Physics: Condensed Matter **14**: R415–R451.

Lenman, James (2005) "The Saucer of Mud, the Kudzu Vine and the Uxorious Cheetah: Against Neo-Aristotelian Naturalism in Metaethics," European Journal of Analytical Philosophy **1**(2): 37–50.

Lennox, James G. (1992) "Teleology," in Evelyn Fox Keller and Elisabeth A. Lloyd, eds., Keywords in Evolutionary Biology. Cambridge, MA: Harvard University Press, pp. 324–333.

Lewens, Tim (2010) "The Natures of Selection," British Journal for the Philosophy of Science **61**: 313–333.

Lippincott-Schwartz, Jennifer, Theresa H. Roberts, and Koret Hirschberg (2000) "Secretory Protein Trafficking and Organelle Dynamics in Living Cells," Annual Review of Cell and Developmental Biology **16**: 557–589.

Lloyd, David and Anthony J. Hayes (1995) "Vigour, Vitality and Viability of Microorganisms," FEMS Microbiology Letters **133**: 1–7.

Loewenstein, Werner R. (1999) The Touchstone of Life: Molecular Information, Cell Communication, and the Foundations of Life. New York: Oxford University Press.

Lorenz, Konrad (1977) Behind the Mirror: A Search for a Natural History of Human Knowledge. New York: Harcourt Brace Jovanovich. (Originally published as Die Rückseite des Spiegels: Versuch einer Naturgeschichte menschlichen Erkennens. Munich: Piper, 1973.)

Love, Alan C. (2008) "Explaining Evolutionary Innovations and Novelties: Criteria of Explanatory Adequacy and Epistemological Prerequisites," Philosophy of Science **75**: 874–886.

Lowe, E.J. (2008) Personal Agency: The Metaphysics of Mind and Action. Oxford: Oxford University Press.

Luby-Phelps, Katherine (2000) "Cytoarchitecture and Physical Properties of Cytoplasm: Volume, Viscosity, Diffusion, Intracellular Surface Area," in Harry Walter, Donald E. Brooks, and Paul A. Srere, eds., International Journal of Cytology, Vol. 192: Microcompartmentation and Phase Separation in Cytoplasm. San Diego, CA: Academic Press, pp. 189–221.

Lurz, Robert W., ed. (2009) The Philosophy of Animal Minds. Cambridge: Cambridge University Press.

Lyon, Pamela (2006) "The Biogenic Approach to Cognition," Cognitive Processing **7**: 11–29.

Macdonald, Graham and David Papineau, eds. (2006) Teleosemantics: New Philosophical Essays. Oxford: Oxford University Press.

Mackie, John L. (1990) Ethics: Inventing Right and Wrong: London: Penguin Books. (Originally published by Pelican Books in 1977.)

Macphail, Euan M. (1998) The Evolution of Consciousness. Oxford: Oxford University Press.

Malcolm, Norman (1977) "Thoughtless Brutes," in idem, Thought and Knowledge. Ithaca, NY: Cornell University Press, pp. 40–57.

Mameli, Matteo (2004) "Nongenetic Selection and Nongenetic Inheritance," British Journal for the Philosophy of Science **55**: 35–71.

Manning, Richard N. (1997) "Biological Function, Selection, and Reduction," British Journal for the Philosophy of Science **48**: 69–82.

Maselko, Jerzy and Maciej Maselko (2009) "Nonlinear Chemical Dynamics and the Origin of Life: The Inorganic-Physical Chemist Point of View," in Steen Rasmussen, Mark A. Bedau, Liaohai Chen, David Deamer, David C. Krakauer, Norman H. Packard, and Peter F. Stadler, eds., Protocells: Bridging Nonliving an Living Matter. Cambridge, MA: Bradford Books/MIT Press, pp. 537–562.

Masgrau, Laura, Anna Roujeinikova, Linus O. Johannissen, Parvider Hothi, Jaswir Basran, Kara E. Ranaghan, Adrian J. Mulholland, Michael J. Sutcliffe, Nigel S. Scrutton, and David Leys (2006) "Atomic Description of an Enzyme Reaction Dominated by Proton Tunneling," Science **312**: 237–241.

Massey, Gerald J. and Deborah A. Boyle (1999) "Descartes's Tests for (Animal) Mind," Philosophical Topics **27**(1), Spring: 87–146.

Matthen, Mohan (1999) "Evolution, Wisconsin Style: Selection and the Explanation of Individual traits," British Journal for the Philosophy of Science **50**: 143–150.

Matthen, Mohan and André Ariew (2002) "Two Ways of Thinking about Fitness and Natural Selection," Journal of Philosophy **99**: 55–83.

Matthen, Mohan and André Ariew (2009) "Selection and Causation," Philosophy of Science **76**: 201–224.

Mattick, John S. (2009) "Deconstructing the Dogma: A New View of the Evolution and Genetic Programming of Complex Organisms," in Günther Witzany, ed., Natural Genetic Engineering and Natural Genome Editing (=Annals of the New York Academy of Sciences, vol. 1178). Boston, MA: Blackwell Publishers on behalf of the New York Academy of Sciences, pp. 29–46.

Maund, Barry (2000) "Proper Functions and Aristotelian Functions in Biology," Studies in History and Philosophy of Biological and Biomedical Sciences **31**: 155–178.

Mavromatos, Nick E. (2010) "Quantum Mechanical Aspects of Cell Microtubules: Science Fiction or Realistic Possibility?" http://arxiv.org/PS_cache/arxiv/pdf/1011/1011.6270v1.pdf

Maye, Alexander, Chih-hao Hsieh, George Sugihara, and Björn Brembs (2007) "Order in Spontaneous Behavior," PLoS ONE **2**(5): article e443.

Mayer, Bruce J., Michael L. Blinov, and Leslie M. Loew (2009) "Molecular Machines or Pleiomorphic Ensembles: Signaling Complexes Revisited," Journal of Biology **8**: number 81.

Maynard Smith, John (2010) "The Concept of Information in Biology," in Paul Davies and Niels Henrik Gregersen, eds., Information and the Nature of Reality: From Physics to Metaphysics. Cambridge: Cambridge University Press, pp. 123–145. (Originally published in Philosophy of Science, 2000, **67**: 177–194.)

Mayr, Ernst (1982) The Growth of Biological Thought: Diversity, Evolution, and Inheritance. Cambridge, MA: Belknap Press/Harvard University Press.

Mayr, Ernst (1988) "The Multiple Meanings of Teleological," in idem, Toward a New Philosophy of Biology. Cambridge, MA: Harvard University Press, pp. 38–66.

McDowell, John (1998) "Two Sorts of Naturalism," in idem, Mind, Value, and Reality. Cambridge, MA: Harvard University Press, pp. 167–197. (Originally published in Rosalind Hursthouse, Gavin Lawrence, and Warren Quinn, eds., Virtues and Reasons: Philippa Foot and Moral Theory. Oxford: Clarendon Press, 1996, pp. 149–179.)

McLaughlin, Peter (2001) What Functions Explain: Functional Explanation and Self-Reproducing Systems. Cambridge: Cambridge University Press.

McLaughlin, Peter (2009) "Functions and Norms," in Ulrich Krohs and Peter Kroes, eds., Functions in Biological and Artificial Worlds: Comparative Philosophical Perspectives. Cambridge, MA: MIT Press, pp. 93–102.

McLaughlin, Peter (2011) "The Arrival of the Fittest," in Dennis Dieks, Wenceslao J. Gonzalez, Stephan Hartmann, Thomas Uebel, and Marcel Weber, eds., Explanation, Prediction, and Confirmation. Dordrecht, Netherlands: Springer, pp. 203–222.

McNiven, M.A. (2003) "The Solid State Cell," Biology of the Cell **94**: 555–556.

McShea, Daniel W. and Robert N. Brandon (2010) *Biology's First Law: The Tendency* for Diversity and Complexity to Increase in Evolutionary Systems. Chicago: University of Chicago Press.

Mehta, Pankaj, Sidhartha Goyal, Tao Long, Bonnie L. Bassler, and Ned S. Wingreen (2009) "Information Processing and Signal Integration in Bacterial Quorum Sensing," Molecular Systems Biology **5**: article number 325.

Millikan, Ruth Garrett (1984) Language, Thought, and Other Biological Categories: New Foundations for Realism. Cambridge, MA: Bradford Books/MIT Press.

Millikan, Ruth Garrett (1989) "In Defense of Proper Functions," Philosophy of Science **56**: 288–302. (Reprinted in Allen et al. 1998.)

Millikan, Ruth Garrett (1996) "On Swampkinds," Mind and Language **11**: 103–117.

Millikan, Ruth Garrett (2002) "Biofunctions: Two Paradigms," in André Ariew, Robert Cummins, and Mark Perlman, eds., Functions: New Essays in the Philosophy and Psychology of Biology. Oxford: Oxford University Press, pp. 113–143.

Millstein, Roberta L. (2006) "Natural Selection as a Population-Level Causal Process," British Journal for the Philosophy of Science **57**: 627–653.

Minton, Allen P. (2001) "The Influence of Macromolecular Crowding and Macromolecular Confinement on Biochemical Reactions in Physiological Media," Journal of Biological Chemistry **276**: 10577–10580.

Misteli, Tom (2001) "The Concept of Self-Organization in Cellular Architecture," Journal of Cell Biology **155**: 181–185.

Misteli, Tom (2007) "Beyond the Sequence: Cellular Organization of Genome Function," Cell **128**: 787–800.

Mogil, Jeffrey S. (2009) "Animal Models of Pain: Progress and Challenges," Nature Reviews Neuroscience **10**: 283–294.

Monod, Jacques (1972) Chance and Necessity: An Essay in the Natural Philosophy of Modern Biology. New York: Vintage Books. (Originally published as Le hasard et la nécessité: Essai sur la philosophie naturelle de la biologie moderne. Paris: Seuil, 1970.)

Monteiro, Antónia and Ondrej Podlaha (2009) "Wings, Horns, and Butterfly Eyespots: How Do Complex Traits Evolve?," PLoS Biology **7**(2): article e1000037.

Moreno Bergareche, Álvaro and Kepa Ruiz-Mirazo (1999), "Metabolism and the Problem of Its Universalization", BioSystems 49: 45-61.

Morrison, Margaret (2006) "Emergence, Reduction, and Theoretical Principles: Rethinking Fundamentalism," Philosophy of Science **73**: 876–887.

Moscona, Aron A. (1959) "Tissues from Dissociated Cells," Scientific American **200**(5): 132–144.

Moss, Lenny (2003) *What Genes Can't Do.* Cambridge, MA: Bradford Books/MIT Press.

Mossio, Matteo, Cristian Saborido, and Álvaro Moreno (2009) "An Organizational Account of Biological Functions," British Journal for the Philosophy of Science **60**: 813–841.

Müller, Gerd B. (2007) "Evo-Devo: Extending the Evolutionary Synthesis," Nature Reviews Genetics **8**: 943–949.

Müller, Gerd B. and Stuart A. Newman, eds. (2003) Origination of Organismal Form: Beyond the Gene in Developmental and Evolutionary Biology. Cambridge, MA: Bradford Books/MIT Press.

Müller, Gerd B. and Stuart A. Newman (2005) "The Innovation Triad: An Evo-Devo Agenda," Journal of Experimental Zoology Part B: Molecular and Developmental Evolution **304B**: 487–503.

Mundale, Jennifer and William Bechtel (1996) "Integrating Neuroscience, Psychology, and Evolutionary Biology through a Teleological Conception of Function," Minds and Machines **6**: 481–505.

Nagel, Ernest (1979) "Teleology Revisited," in idem, Teleology Revisited and Other Essays in the Philosophy and History of Science. New York: Columbia University Press, pp. 275–316**.** (Originally published in Journal of Philosophy, 1977, **74**: 261–301.)

Nagel, Zachary D. and Judith P. Klinman (2006) "Tunneling and Dynamics in Enzymatic Hydride Transfer," Chemical Reviews **106**: 3095–3118.

Nakagaki, Toshiyuki, Hiroyasu Yamada, and Ágotha Tóth (2000) "Maze-Solving by an Amoeboid Organism," Nature **407**: 470.

Nakagaki, Toshiyuki, Atsushi Tero, Ryo Kobayashi, Isamu Onishi, and Tomoyuki Miyaji (2009) "Computational Ability of Cells based on Cell Dynamics and Adaptability," New Generation Computing **27**: 57–81.

Nakagaki, Toshiyuki and Robert D. Guy (2008) "Intelligent Behaviors of Amoeboid Movement Based on Complex Dynamics of Soft Matter," Soft Matter **4**: 57–67.

Nanay, Bence (2005) "Can Cumulative Selection Explain Adaptation?," Philosophy of Science **72**: 1099–1112.

Nanay, Bence (2010) "A Modal Theory of Function," Journal of Philosophy **107**: 412–431.

Neander, Karen (1995a) "Pruning the Tree of Life," British Journal for the Philosophy of Science **46**: 59–80.

Neander, Karen (1995b) "Explaining Complex Adaptations: A Reply to Sober's 'Reply to Neander,'" British Journal for the Philosophy of Science **46**: 583–587.

Nédélec, François J., Thomas Surrey, A.C. Maggs, and Stanislas Leibler (1997) "Self-Organization of Microtubules and Motors," Nature **389**: 305–308.

Neumann-Held, Eva M. and Christoph Rehmann-Sutter, eds. (2006) Genes in Development: Re-reading the Molecular Program. Durham, NC: Duke University Press.

Newth, David and John Finnigan (2006) "Emergence and Self-Organization in Chemistry and Biology," Australian Journal of Chemistry **59**: 841–848.

Newton, Jessica R. and Mriganka Sur (2005) "Rewiring Cortex: Functional Plasticity of the Auditory Cortex during Development," in Josef Syka and Michael M. Merzenich, eds., Plasticity and Signal Representation in the Auditory System. Berlin: Springer, pp. 127–138.

Nissen, Lowell (1997) Teleological Language in the Life Sciences. Lanham, MD: Rowman & Littlefield Publishers.

Nordenskiöld, Erik (1928) The History of Biology: A Survey. New York: Tudor Publishing Co. (Originally published in three volumes as Biologins historia, in Stockholm, 1920–1924; numerous English editions.)

O'Connor, Timothy (2000) Persons and Causes: The Metaphysics of Free Will. New York: Oxford University Press.

Okrent, Mark (2007) Rational Animals: The Teleological Roots of Intentionality. Athens, OH: University of Ohio Press.

O'Malley, Maureen A. and John Dupré (2007) "Size Doesn't Matter: Towards a More Inclusive Philosophy of Biology," Biology and Philosophy **22**: 155–191.

Ó Nualláin, Seán (2008) "Subjects and Objects: Metaphysics, Biology, Consciousness, and Cognition," Biosemiotics **1**: 239–251.

Orgel, Leslie E. (2008) "The Implausibility of Metabolic Cycles on the Prebiotic Earth," PLoS Biology **6**(1): article e18.

Palmer, A. Richard (2004) "Symmetry Breaking and the Evolution of Development," Science **306**: 828–933.

Panitchayangkoon, Gitt, Dugan Hayes, Kelly A. Fransted, Justin R. Caram, Elad Harel, Jianzhong Wen, Robert E. Blankenship, and Gregory S. Engel (2010) "Long-Lived Quantum Coherence in Photosynthetic Complexes at Physiological Temperatures," Proceedings of the National Academy of Sciences, USA **107**: 12766–12770.

Panksepp, Jaak (2005) "Affective Consciousness: Core Emotional Feelings in Animals and Humans," Consciousness and Cognition **14**: 30–80.

Parfit, Derek (forthcoming) On What Matters, 2 vols. Oxford: Oxford University Press.

Parson, William W. (2007) "Long Live Electronic Coherence!," Science **316**: 1438–1439.

Pattee, Howard H. (1982) "Cell Psychology: An Evolutionary Approach to the Symbol-Matter Problem," Cognition and Brain Theory **5**: 325–341.

Pattee, Howard H. (2001) "The Physics of Symbols: Bridging the Epistemic Cut," BioSystems **60**: 5–21.

Penzlin, Heinz (2009) "The Riddle of 'Life,' a Biologist's Critical View," Naturwissenschaften **96**: 1–23.

Perovic, Slobodan (2007) "The Limitations of Kim's Reductive Physicalism in Accounting for Living Systems and an Alternative Nonreductionist Ontology," Acta Biotheoretica **55**: 243–267.

Pessa, Eliano (2006) "Physical and Biological Emergence: Are They Different?," in Gianfranco Minati, Eliano Pessa, and Mario Abram, eds., Systems of Emergence: Research and Development. New York: Springer, pp. 355–374.

Phillips, David (2010) "Mackie on Practical Reason," in Richard Joyce and Simon Kirchin, eds., *A World Without Values: Essays on John Mackie's Moral* Error Theory. Berlin: Springer, pp. 87–99.

Phillips, Rob and Stephen R. Quake (2006) "The Biological Frontier of Physics," Physics Today **59**(5): 38–43.

Pichot, André (2004) Histoire de la notion de vie. Paris: Gallimard.

Piersma, Theunis and Jan A. van Gils (2011) The Flexible Phenotype: A Body-Centered Integration of Ecology, Physiology, and Behaviour. Oxford: Oxford University Press.

Pigliucci, Massimo (2009) "An Extended Synthesis for Evolutionary Biology," in Carl D. Schlichting and Timothy A. Mousseau, eds., The Year in Evolutionary Biology 2009 (=Annals of the New York Academy of Sciences, vol. 1168). Malden, MA: Wiley-Blackwell, on behalf of the New York Academy of Sciences, pp. 218–228.

Pigliucci, Massimo and Jonathan Kaplan (2006) Making Sense of Evolution: The Conceptual Foundations of Evolutionary Theory. Chicago: University of Chicago Press.

Pigliucci, Massimo and Gerd B. Müller, eds. (2010) Evolution: The Extended Synthesis. Cambridge, MA: MIT Press.

Pigliucci, Massimo, Courtney J. Murren, and Carl D. Schlichting (2006) "Phenotypic Plasticity and Evolution by Genetic Assimilation," Journal of Experimental Biology **209**: 2362–2367.

Piro, Rosario Michael (2011) "Are All Genes Regulatory Genes?," Biology and Philosophy **26**: 595–602.

Pittendrigh, Colin S. (1958) “Adaptation, Natural Selection, and Behavior,” in Anne Roe and George Gaylord Simpson, eds., Behavior and Evolution. New Haven, CT: Yale University Press, pp. 390–416.

Plenio, Martin B. and Susana F. Huelga (2008) “Dephasing-Assisted Transport: Quantum Networks and Biomolecules,” New Journal of Physics **10**: number 113019.

Pokorný, Jiří and Tsu-ming Wu (1998) Biophysical Aspects of Coherence and Biological Order. Berlin: Springer-Verlag.

Pollack, Gerald H. (2001) Cells, Gels, and the Engines of Life: A New, Unifying Approach to Cell Function. Seattle: Ebner & Sons.

Pollack, Gerald H. (2006) “Cells, Gels, and Mechanics,” in Mohammad R. K. Mofrad and Roger D. Kamm, eds., Cytoskeletal Mechanics: Models and Measurements. Cambridge: Cambridge University Press, pp. 129–151.

Pollack, Gerald H. and Wei-Chun Chin, eds. (2008) Phase Transitions in Cell Biology. Dordrecht, Netherlands: Springer.

Pollack, Gerald H., Ivan L. Cameron, and Denys N. Wheatley, eds. (2006) Water and the Cell. Dordrecht, Netherlands: Springer.

Popa, Radu (2004) Between Necessity and Probability: Searching for the Definition and Origin of Life. New York: Springer.

Portin, Petter (2009) “The Elusive Concept of the Gene,” Hereditas **146**: 112–117.

Povinelli, Daniel J. and Steve Giambrone (1999) “Inferring Other Minds: Failure of the Argument by Analogy,” Philosophical Topics **27**(1), Spring: 167–201.

Prete, Frederick R., ed. (2004) Complex Worlds from Simpler Nervous Systems. Cambridge, MA: Bradford Books/MIT Press.

Ptito, Maurice, Solvey M. Moesgaard, Albert Giedde, and Ron Kupers (2005) “Cross-Modal Plasticity Revealed by Electrotactile Stimulation of the Tongue of the Congenitally Blind,” Brain **128**: 606–614.

Pust, Joel (2001) “Natural Selection Explanation and Origin Essentialism,” Canadian Journal of Philosophy **31**: 201–220.

Queiroz, João and Charbel Niño El Hani (2006) “Towards a Multi-Level Approach to the Emergence of Meaning Processes in Living Systems,” Acta Biotheoretica **54**: 179–206.

Radner, Daisie and Michael Radner (1989) Animal Consciousness. Buffalo, NY: Prometheus Books.

Railton, Peter (2009) “Practical Competence and Fluent Agency,” in David Sobel and Steven Wall, eds., Reasons for Action. Cambridge: Cambridge University Press, pp. 81–115.

Reiss, John O. (2009) *Not by Design: Retiring Darwin’s Watchmaker.* Berkeley: University of California Press.

Rensberger, Boyce (1996) Life Itself: Exploring the Realm of the Living Cell. New York: Oxford University Press.

Reznikova, Zhanna (2007) Animal Intelligence: From Individual to Social Cognition. Cambridge: Cambridge University Press.

Ristau, Carolyn A., ed. (1991) Cognitive Ethology: The Minds of Other Animals: Essays in Honor of Donald R. Griffin. Hillsdale, NJ: Lawrence Erlbaum Associates.

Rizzotti, Martino, ed. (1996) Defining Life: The Central Problem in Theoretical Biology. Padua, Italy: University of Padova.

Robert, Jason Scott (2004) Embryology, Epigenesis, and Evolution: Taking Development Seriously. Cambridge: Cambridge University Press.

Roederer, Juan G. (2005) Information and Its Role in Nature. Berlin: Springer.

Rosati, Connie S. (2003) “Agency and the Open Question Argument,” Ethics **113**: 490–527.

Rosen, Robert (1991) Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life. New York: Columbia University Press.

Rosen, Robert (2000) Essays on Life Itself. New York: Columbia University Press.

Rosenberg, Alex (2006) Darwinian Reductionism, Or, How to Stop Worrying and Love Molecular Biology. Chicago: University of Chicago Press.

Rosenblueth, Arturo, Norbert Wiener, and Julian Bigelow (1943) “Behavior, Purpose, and Teleology,” Philosophy of Science **10**: 18–24.

Ruiz-Mirazo, Kepa, Juli Peretó, and Álvaro Moreno (2004) “A Universal Definition of Life: Autonomy and Open-Ended Evolution,” Origins of Life and Evolution of the Biosphere **34**: 323–346.

Ruiz-Mirazo, Kepa, Juli Peretó, and Álvaro Moreno (2010) “Defining Life or Bringing Biology to Life,” Origins of Life and Evolution of the Biosphere **40**: 203–213.

Rundle, Bede (1997) Mind in Action. Oxford: Clarendon Press.

Sachse, Christian (2007) Reductionism in the Philosophy of Science. Frankfurt: Ontos Verlag.

Saigusa, Tetsu, Atsushi Tero, Toshiyuki Nakagaki, and Yoshiki Kuramoto (2008) "Amoebae Anticipate Periodic Events," Physical Review Letters **100**: number 018101.

Sarkar, Sahotra (2005) "Biological Information: A Skeptical Look at Some Central Dogmas of Molecular Biology," in idem, Molecular Models of Life: Philosophical Papers on Molecular Biology. Cambridge, MA: Bradford Books/MIT Press, pp. 205–260. (Originally published in Sahotra Sarkar, ed., The Philosophy and History of Molecular Biology: New Perspectives. Dordrecht, Netherlands: Kluwer Academic, 1996, pp. 187–231.)

Sarovar, Mohan, Akihito Ishizaki, Graham R. Fleming, and K. Birgitta Whaley (2010) "Quantum Entanglement in Photosynthetic Light-Harvesting Complexes," Nature Physics **6**: 462–467.

Savellos, Elias E. and Ümit D. Yalçin, eds. (1995) Supervenience: New Essays. Cambridge: Cambridge University Press.

Schlichting, Carl D. and Massimo Pigliucci (1998) Phenotypic Evolution: A Reaction Norm Perspective. Sunderland, MA; Sinauer Associates.

Schlosshauer, Maximilian (2007) Decoherence and the Quantum-to-Classical Transition. Berlin: Springer.

Schneider, Eric D. and James J. Kay (1995) "Order from Disorder: The Thermodynamics of Complexity in Biology," in Michael P. Murphy and Luke A.J. O'Neill, eds., What Is Life? The Next Fifty Years: Speculations on the Future of Biology. Cambridge: Cambridge University Press, pp. 161–174.

Schneider, Eric D. and Dorion Sagan (2005) Into the Cool: Energy Flow, Thermodynamics, and Life. Chicago: University of Chicago Press.

Scholzen, Thomas and Johannes Gerdes (2000) "The Ki-67 Protein: From the Known to the Unknown," Journal of Cellular Physiology **183**: 311–322.

Schueler, G.F. (2003) Reasons and Purposes: Human Rationality and the Teleological Explanation of Action. Oxford: Clarendon Press.

Scott, Alwyn C. (2007) The Nonlinear Universe: Chaos, Emergence, Life. Berlin: Springer.

Searle, John (1994) "Animal Minds," Midwest Studies in Philosophy **19**: 206–219.

Sehon, Scott (2005) Teleological Realism: Mind, Agency, and Explanation. Cambridge, MA: Bradford Books/MIT Press.

Sergi, Alessandro (2009) "Quantum Biology," Atti della Accademia Peloritana dei Pericolanti; Classe di Scienze Fisiche, Matematiche e Naturali **87**: number

C1C0901001.
http://cab.unime.it/journals/index.php/AAPP/article/viewFile/449/352

Seth, Anil K., Bernard J. Baars, and David B. Edelman (2005) "Criteria for Consciousness in Humans and Other Mammals," Consciousness and Cognition **14**: 119–139.

Shapiro, James A. (2007) "Bacteria Are Small but Not Stupid: Cognition, Natural Genetic Engineering and Socio-Bacteriology," Studies in History and Philosophy of Biological and Biomedical Sciences **38**: 807–819.

Shapiro, James A. (2009) "Revisiting the Central Dogma in the 21st Century," in Günther Witzany, ed., Natural Genetic Engineering and Natural Genome Editing (=Annals of the New York Academy of Sciences, vol. 1178). Boston, MA: Blackwell Publishers on behalf of the New York Academy of Sciences, pp. 6–28.

Shapiro, James A. (2011) Evolution: A View from the 21st Century. Upper Saddle River, NJ: FT Press.

Shapiro, Larry and Elliott Sober (2007) "Epiphenomenalism: The Dos and the Don'ts," in Peter Machamer and Gereon Wolters, eds., Thinking about Causes: From Greek Philosophy to Modern Physics. Pittsburgh, PA: University of Pittsburgh Press, pp. 235–264.

Sharma, Jitendra, Alessandra Angelucci, and Mriganka Sur (2000) "Induction of Visual Orientation Modules in Auditory Cortex," Nature **404**: 841–847.

Shepherd, Virginia A. (2006) "The Cytomatrix as a Cooperative System of Macromolecular and Water Networks," in Gerald P. Schatten, ed., Current Topics in Developmental Biology, Vol. 75. San Diego, CA: Academic Press/Elsevier, pp. 171–223.

Shields, Christopher (2012) "The Dialectic of Life," Synthese **185**: 103–124.

Shimizu, Thomas Simon and Dennis Bray (2003) "Modelling the Bacterial Chemotaxis Receptor Complex," in Gregory Bock and Jamie A. Goode, eds., *'*In Silico*'* Simulation of Biological Processes (Novartis Foundation Symposium 247). Chichester, UK: John Wiley & Sons, pp. 162–177.

Shubin, Neil, Cliff Tabin, Sean Carroll (2009) "Deep Homology and the Origins of Evolutionary Novelty," Nature **457**: 818–823.

Simons, Paul (1992) The Action Plant: Movement and Nervous Behaviour in Plants. Oxford: Blackwell.

Skewes, Joshua C. and Clifford A. Hooker (2009) "Bio-agency and the Problem of Action," Biology and Philosophy **24**: 283–300.

Slijper, E.J. (1942) "Biologic-anatomical Investigations on the Bipedal Gait and Upright Posture in Mammals, with Special Reference to a Little Goat, born without Forelegs," Proceedings of the Koninklijke Nederlandsche Akademie van Wetenschappen **45**: 288–295, 407–415.

Sober, Elliott (1984) The Nature of Selection: Evolutionary Theory in Philosophical Focus. Cambridge, MA: Bradford Books/MIT Press, 1984.

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

Sober, Elliott (2000) The Philosophy of Biology, 2nd ed. Boulder, CO: Westview Press. (1st ed. published in 1993.)

Sober, Elliott (2005) "Comparative Psychology Meets Evolutionary Biology: Morgan's Canon and Cladistic Parsimony," in Lorraine Daston and Gregg Mitman, eds., Thinking with Animals: New Perspectives on Anthropomorphism. New York: Columbia University Press, pp. 85–99.

Sommerhoff, Gerd (1969) "The Abstract Characteristics of Living Systems," in F.E. Emergy, ed., Systems Thinking. Harmondsworth, Middlesex, UK: Penguin Education, pp. 147–202.

Sommerhoff, Gerd (1990) Life, Brain, and Consciousness: New Perceptions through Targeted Systems Analysis. Amsterdam: North-Holland.

Sommers, Tamler and Alex Rosenberg (2003) "Darwin's Nihilistic Idea: Evolution and the Meaninglessness of Life," Biology and Philosophy **18**: 653–668.

Stamp, Philip C.E. (2006) "The Decoherence Puzzle," Studies in History and Philosophy of Modern Physics **37**: 467–497.

Stegmann, Ulrich E. (2010) "What Can Natural Selection Explain?," Studies in History and Philosophy of Biological and Biomedical Sciences **41**: 61–66.

Stelling, Jörg, Uwe Sauer, Zoltan Szallasi, Francis J. Doyle, III, and John Doyle (2004) "Robustness of Cellular Functions," Cell **118**: 675–685.

Stephens, Christopher (2004) "Selection, Drift, and the 'Forces' of Evolution," Philosophy of Science **71**: 550–570.

Stephens, Christopher (2010) "Forces and Causes in Evolutionary Theory," Philosophy of Science **77**: 716–727.

Sterelny, Kim and Paul E. Griffiths (1999) Sex and Death: An Introduction to Philosophy of Biology. Chicago: University of Chicago Press.

Steward, Helen (2009a) "Animal Agency," Inquiry **52**: 217–231.

Steward, Helen (2009b) "Sub-intentional Actions and the Over-mentalization of Agency," in Constantine Sandis, ed., New Essays on the Explanation of Action. London: Palgrave/Macmillan, pp. 295–312.

Stewart, John (1996) "Cognition = Life: Implications for Higher-Level Cognition," Behavioural Processes **35**: 311–326.

Stock, Jeffry B. and Michael G. Surette (1996) "Chemotaxis," in Frederick C. Neidhardt, ed., Escherichia coli and Salmonella: Cellular and Molecular Biology, 2nd ed. Washington, DC: ASM Press, Vol. I, pp. 1103–1129.

Stotz, Karola (2006) "With 'Genes' Like That, Who Needs an Environment? Postgenomics's Argument for the 'Ontogeny of Information,'" Philosophy of Science **73**: 905–917.

Strand, Anders and Gry Oftedal (2009) "Functional Stability and Systems Level Causation," Philosophy of Science **76**: 809–820.

Stratton, George M. (1897) "Vision without Inversion of the Retinal Image," Psychological Review **4**: 341–360, 463–481.

Street, Sharon (2006) "A Darwinian Dilemma for Realist Theories of Value," Philosophical Studies **127**: 109–166.

Sultan, Sonia E. (1992) "Phenotypic Plasticity and the Neo-Darwinian Legacy," Evolutionary Trends in Plants **6**: 61–71.

Sutcliffe, Michael J. and Nigel S. Scrutton (2000a) "Enzyme Catalysis: Over-the-Barrier or Through-the-Barrier?," Trends in Biochemical Sciences **25**: 405–408.

Sutcliffe, Michael J. and Nigel S. Scrutton (2000b) "Enzymology Takes a Quantum Leap Forward," Philosophical Transactions of the Royal Society of London A **358**: 367–386.

Swenson, Rod (1998) "Spontaneous Order, Evolution, and Autocatakinetics: The Nomological Basis for the Emergence of Meaning," in Gertrudis Van de Vijver, Stanley N. Salthe, and Manuela Delpos, eds., Evolutionary Systems: Biological and Epistemological Perspectives on Selection and Self-Organization. Dordrecht, Netherlands: Kluwer Academic, pp. 155–180.

Tegmark, Max (2000) "Importance of Quantum Decoherence in Brain Processes," Physical Review E **61**: 4194–4206.

Tero, Atsushi, Ryo Kobayashi, and Toshiyuki Nakagaki (2005) "A Coupled-Oscillator Model with a Conservation Law for the Rhythmic Amoeboid Movements of Plasmodial Slime Molds," Physica D **205**: 125–135.

Tero, Atsushi, Seiji Takagi, Tetsu Saigusa, Kentaro Ito, Dan P. Bebber, Mark D. Fricker, Kenji Yumiki, Ryo Kobayashi, and Toshiyuki Nakagaki (2010) "Rules for Biologically Inspired Adaptive Network Design," Science **327**: 439–442.

Terzis, George and Robert Arp, eds. (2011) Information and Living Systems: Philosophical and Scientific Perspectives. Cambridge, MA: Bradford Books/MIT Press.

Thalos, Mariam (2007) "The Sources of Behavior: Toward a Naturalistic, Control Account of Agency," in Don Ross, David Spurrett, Harold Kincaid, and G. Lynn Stephens, eds., Distributed Cognition and the Will: Individual Volition and Social Context. Cambridge, MA: Bradford Books/MIT Press, pp. 123–167.

Thompson, Evan (2007) Mind in Life: Biology, Phenomenology, and the Sciences of Mind. Cambridge, MA: Belknap Press/Harvard University Press.

Thompson, Michael (2008) Life and Action: Elementary Structures of Practice and Practical Thought. Cambridge, MA: Harvard University Press.

Tokuriki, Nobuhiko and Dan S. Tawfik (2009) "Protein Dynamism and Evolvability," Science **324**: 203–207.

Trewavas, Anthony J. (1999) "The Importance of Individuality," in H.R. Lerner, ed., Plant Responses to Environmental Stresses: From Phytohormones to Genome Reorganization. New York: Marcel Dekker, pp. 27–42.

Trewavas, Anthony (2003) "Aspects of Plant Intelligence," Annals of Botany **92**: 1–20.

Trewavas, Anthony (2005) "Plant Intelligence," Naturwissenschaften **92**: 401–413.

Trewavas, Anthony (2009) "What Is Plant Behaviour?," Plant, Cell and Environment **32**: 606–616.

Trewavas, Anthony (2010) "The Green Plant as an Intelligent Organism," in František Baluška, Stefano Mancuso, and Dieter Volkmann, eds., Communication in Plants: Neuronal Aspects of Plant Life. Berlin: Springer, pp. 1–18.

Tropea, Daniela, Audra Van Wart, and Mriganka Sur (2009) "Molecular Mechanisms of Experience-Dependent Plasticity in Visual Cortex," Philosophical Transactions of the Royal Society B **364**: 341–355.

Turner, J. Scott (2007) *The Tinkerer's Accomplice: How Design Emerges from Life Itself.* Cambridge, MA: Harvard University Press.

Turvey, Michael T. (2004) "Impredicativity, Dynamics, and the Perception-Action Divide," in Viktor K. Jirsa and J.A. Scott Kelso, eds., Coordination Dynamics: Issues and Trends. Berlin: Springer, pp. 1–20.

Uller, Tobias (2008) "Developmental Plasticity and the Evolution of Parental Effects," Trends in Ecology and Evolution **23**: 432–438.

Uversky, Vladimir N. (2010) "The Mysterious Unfoldome: Structureless, Underappreciated, Yet Vital Part of Any Given Proteome," Journal of Biomedicine and Biotechnology **2010**: article ID 568068.

van Duijn, Marc, Fred Keijzer, and Daan Franken (2006) "Principles of Minimal Cognition: Casting Cognition as Sensorimotor Coordination," Adaptive Behavior **14**: 157–170.

van Gelder, Timothy (1995) "What Might Cognition Be, If Not Computation?," Journal of Philosophy **92**: 345–381.

Vendruscolo, Michele and Christopher M. Dobson (2006) "Dynamic Visions of Enzymatic Reactions," Science **313**: 1586–1587.

von Melchner, Laurie, Sarah L. Pallas, and Mriganka Sur (2000) "Visual Behaviour Mediated by Retinal Projections Directed to the Auditory Pathway," Nature **404**: 871–876.

von Wright, Hans Georg (1971) Explanation and Understanding. Ithaca, NY: Cornell University Press.

Vries, Hugo de (1912) Species and Varieties: Their Origin by Mutation, 3rd ed, corrected and revised. Chicago: Open Court. (1st ed. published by Open Court in 1904; 3rd ed. reprinted by Garland Publishing, Inc., New York and London, 1988.)

Wachbroit, Robert (1994) "Normality as a Biological Concept," Philosophy of Science **61**: 579–591.

Wadham, George H. and Judith P. Armitage (2004) "Making Sense of It All: Bacterial Chemotaxis," Nature Reviews: Molecular & Cell Biology **5**: 1024–1037.

Wagner, Andreas (2005) Robustness and Evolvability in Living Systems. Princeton, NJ: Princeton University Press.

Wagner, Günter P., Chi-Hua Chiu, and Manfred Laubichler (2000) "Developmental Evolution as a Mechanistic Science: The Inference from Developmental Mechanisms to Evolutionary Processes," American Zoologist **40**: 819–831.

Walker, Stephen (1983) Animal Thought. London: Routledge & Kegan Paul.

Walsh, Denis M. (1998) "The Scope of Selection: Sober and Neander on What Natural Selection Explains," Australasian Journal of Philosophy **76**: 250–264.

Walsh, Denis M. (2003) "Fit and Diversity: Explaining Adaptive Evolution," Philosophy of Science **70**: 280–301.

Walsh, Denis M. (2006) "Organisms as Natural Purposes: The Contemporary Evolutionary Perspective," Studies in History and Philosophy of Biological and Biomedical Sciences **37**: 771–791.

Walsh, Denis M. (2007) "The Pomp of Superfluous Causes: The Interpretation of Evolutionary Theory," Philosophy of Science **74**: 281–303.

Walsh, Denis M., Tim Lewens, and André Ariew (2002) "The Trials of Life: Natural Selection and Random Drift," Philosophy of Science **69**: 452–473.

Warren, William H. (2006) "The Dynamics of Perception and Action," Psychological Review **113**: 358–389.

Waters, Christopher M. and Bonnie L. Bassler (2005) "Quorum Sensing: Cell-to-Cell Communication in Bacteria," Annual Review of Cell and Developmental Biology **21**: 319–346.

Weber, Andreas and Francisco J. Varela (2002) "Life after Kant: Natural Purposes and the Autopoietic Foundations of Biological Individuality," Phenomenology and the Cognitive Sciences **1**: 97–125.

Weber, Bruce H. and David J. Depew, eds. (2003) Evolution and Learning: The Baldwin Effect Reconsidered. Cambridge, MA: Bradford Books/MIT Press.

Webre, Daniel J., Peter M. Wolanin, and Jeffry B. Stock (2003) "Bacterial Chemotaxis," Current Biology **13**(2): R47–R49.

Wedgwood, Ralph (2007) The Nature of Normativity. Oxford: Clarendon Press.

West-Eberhard, Mary Jane (2003) Developmental Plasticity and Evolution. Oxford: Oxford University Press.

West-Eberhard, Mary Jane (2005) "Phenotypic Accommodation: Adaptive Innovation due to Developmental Plasticity," Journal of Experimental Zoology B: Molecular and Developmental Evolution **304B**: 610–618.

West-Eberhard, Mary Jane (2008) "Toward a Modern Revival of Darwin's Theory of Evolutionary Novelty," Philosophy of Science **75**: 899–908.

Wheatley, Denys N. (2003) "Diffusion, Perfusion and the Exclusion Principles in the Structural and Functional Organization of the Living Cell: Reappraisal of the Properties of the 'Ground Substance,'" Journal of Experimental Biology **206**: 1995–1961.

Whitacre, James and Axel Bender (2010) "Degeneracy: A Design Principle for Achieving Robustness and Evolvability," Journal of Theoretical Biology **263**: 143–153.

Whitesides, George M. (2008) “Foreword: The Improbability of Life,” in John D. Barrow, Simon Conway Morris, Stephen J. Freeland, and Charles L. Harper, Jr., eds., Fitness of the Cosmos for Life: Biochemistry and Fine-Tuning. Cambridge: Cambridge University Press, pp. xi–xix.

Whitesides, George M. and Bartosz Grzybowski (2002) “Self-Assembly at All Scales,” Science **295**: 2418–2421.

Wilde, Mark M., James M. McCracken, and Ari Mizel (2010) “Could Light Harvesting Complexes Exhibit Non-Classical Effects at Room Temperature?,” Proceedings of the Royal Society of London A **466**: 1347–1363.

Williams, Bernard (1995) “Evolution, Ethics, and the Representation Problem,” in idem, Making Sense of Humanity and Other Philosophical Papers, 1982–1993. Cambridge: Cambridge University Press, pp. 100–110. (Originally published in D.S. Bendall, ed., Evolution from Molecules to Men. Cambridge: Cambridge University Press, 1983, pp. 554–566.)

Wilson, George M. (1989) The Intentionality of Human Action, revised and enlarged ed. Stanford CA: Stanford University Press.

Wilson, Margaret Dauler (1995) “Animal Ideas,” Proceedings and Addresses of the American Philosophical Society **69**(2), November: 7–25.

Winfree, Arthur T. (2010) The Geometry of Biological Time, 2nd ed. Berlin: Springer.

Wolynes, Peter G. (2008) “Protein Folding and Beyond: Energy Landscapes and the Organization of Living Matter in Time and Space,” in Ahmed H. Zewail, ed., Physical Biology: From Atoms to Medicine. London: Imperial College Press, pp. 267–288.

Wouters, Arno (1995) “Viability Explanation,” Biology and Philosophy **10**: 435–457.

Wouters, Arno (2003) “Four Notions of Biological Function,” Studies in History and Philosophy of Biological and Biomedical Sciences **34**: 633–668.

Wouters, Arno (2005) “The Function Debate in Philosophy,” Acta Biotheoretica **53**: 123–151.

Wright, Larry (1973) “Functions,” Philosophical Review **82**: 139–168. (Reprinted in Allen et al. 1998, and in Buller 1999b.)

Wynne, Clive D. (2004) Do Animals Think? Princeton, NJ: Princeton University Press.

Yadid, Itamar, Noam Kirshenbaum, Michal Sharon, Orly Dym, and Dan S. Tawfik (2010) “Metamorphic Proteins Mediate Evolutionary Transitions of Structure,” Proceedings of the National Academy of Sciences, USA **107**: 7287–7292.

Yates, F. Eugene (1994) “Order and Complexity in Dynamical Systems: Homeodynamics as a Generalized Mechanics for Biology,” Mathematical and Computer Modelling **19**(6–8): 49–74.

Yates, F. Eugene (2008) “Homeokinetics/Homeodynamics: A Physical Heuristic for Life and Complexity,” Ecological Psychology **20**: 148–179.

Yates, F. Eugene and L.B. Yates (2008) “Ultradian Rhythms as the Dynamic Signature of Life,” in David Lloyd and Ernest L. Rossi, eds., Ultradian Rhythms from Molecules to Mind: A New Vision of Life. Berlin: Springer, pp. 249–260.

Zallen, Richard (1983) The Physics of Amorphous Solids. New York: John Wiley & Sons.

Zammito, John (2006) “Teleology Then and Now: The Question of Kant’s Relevance for Contemporary Controversies over Function in Biology,” Studies in History and Philosophy of Biological and Biomedical Sciences **37**: 748–770.

Zhou, Huan-Xiang, Germán Rivas, and Allen P. Minton (2008) “Macromolecular Crowding and Confinement: Biochemcial, Biophysical, and Potential Physiological Consequences,” Annual Review of Biophysics **37**: 375–397.

Zimmer, Carl (2008) Microcosm: E. coli and the New Science of Life. New York: Pantheon.