1

Meaning and the Means to an Understanding of Ends

1.1 Design in Nature

Biology is unique among the natural sciences in its use of a family of concepts that might seem better suited to the description and explanation of artifacts than the description and explanation of organisms. Artifacts are objects made by intelligent agents; organisms—most of them, at least—owe their construction to no agent. When we think about artifacts of all kinds—shoes, ships, sealing wax—we find it natural to ask what might be their functions, and the functions of their parts, what problems they were made to solve, and so forth. Biologists, and evolutionary biologists in particular, use a similar vocabulary when they describe and approach the organic world. They ask what the function of the stiff-legged jumping behavior (called "stotting") of Thompson's gazelles might be; they conjecture that the bony plates on the back of *Stegosaurus* had the purpose of regulating heat; they suggest that the fragile second penises of male earwigs snap off inside the vagina in order to prevent fertilization from other males; they ask what evolutionary problems our hominid ancestors might have faced in the Pleistocene, and what solutions our species might have found to meet them.

The vocabulary of intelligent design—the vocabulary of problems, solutions, purpose, and function—might seem to presuppose the existence of an intelligent designer. The human sciences may speak of purposes and problems addressed by social institutions; that is no surprise, for the human sciences range over systems that contain intelligent designers. The physical sciences generally admit no intelligent designer into their worldview; correspondingly, physicists do not speak of the purposes of electrons, and chemists do not ask what benzene rings were designed for. Thus biology is in an awkward position: it makes free with a vocabulary of design, even though modern biology recognizes no intelligent designer as the artificer of the species.

In summary, many biologists adopt what I call the *artifact model* of nature: they talk of organisms as though they were designed objects. An examination of the artifact model answers questions that are of as much interest to biologists, students of technology, and philosophers of mind as they are to philosophers keen to understand biological explanation. Much of the debate over adaptationism, for example, has been framed as a question of whether it is right to assume that organisms can be divided into traits each with its own function, in the same way that we might try to draw an exploded diagram of a car that assigns discrete functions to its parts. Developmental biologists have argued that a focus on function has led us to ignore some of the most important factors affecting form. More broadly, the investigation of the analogy between evolution and the design process has been thought by some to yield important insights regarding changes in technology itself. And philosophers of mind have thought that an account of how hearts can be supposed to pump blood, even though they may fail to do so, could yield a wholly unmysterious account of how, for example, beliefs can be supposed to represent cows, even though such a belief may fail to represent its object accurately.

This book addresses what I take to be the most pressing questions raised by the phenomena of artifact talk in biology. Such questions include: what explains the ability of biologists to use such a vocabulary? Are the terms they use mere metaphors that trade on superficial similarities between the appearance of organisms and artifacts, or are there close analogies between the processes that go into the construction of each? Might we be misled by approaching organisms as though they are collections of more or less well-designed solutions to environmental problems? Can such a framework give us a strong predictive engine for the generation of hypotheses about the workings of plants and animals—even of the human mind? Might the kinds of norms that we appear committed to—in speaking of what traits are *supposed to do*—be appropriate to solve problems in the philosophy of mind? How should we explain the appearance of artifact talk in biology and its absence in chemistry and physics? Can the function and design of artifacts themselves be approached from an evolutionary perspective? Most recent work in this area has been concerned with giving an analysis of the concept of biological function as it appears in biological journals. This is certainly an important job, and it forms a part of the work of this book; however, we can already see that such a narrow inquiry into function by no means exhausts the tasks of evaluating and understanding the artifact model.

On the face of it, there are quite simple answers to most of the questions I've just raised. The evolutionary process bears deep similarities to the process of intelligent design. It is these deep similarities that explain and justify the appearance of the same vocabulary in both domains. Just as a designer chooses her materials to fashion an object to meet her problems, so nature selects traits to fashion an organism to meet problems laid down by the environment. Natural selection thereby plays a role analogous to intentional choice, and natural selection is what grounds various claims about function and design in the natural world. Since selection works only on organisms that reproduce themselves, it is selection that explains why artifact talk features in biology alone, and not in the physical sciences. Selection gives traits norms that should be met; hence selection can underpin normative function claims of the sort intended to ground projects to naturalize content in the philosophy of mind.

I will argue that such a picture is almost completely mistaken. There are deep similarities between the processes that go into the construction of organisms and artifacts; however, although these can help to explain why both types of objects enjoy a gradual accumulation of useful traits over time, it is a mistake to think that natural selection is a good analogue to the intentions of a designer. And it is the internal constitution of biological items, not the fact that selection acts only on biological items, that best explains the appearance of artifact talk in biology alone.

Much of the argument for these propositions turns on a demonstration of just what natural selection is. Natural selection is essentially a population-level, statistical phenomenon. Intentions, on the other hand, can have influence on individual entities. This does not mean that any other element of the evolutionary process yields a better analogue to intention that might instead be used to ground claims about function or design. We have a choice over just how we wish to tighten up function talk in biology, depending on what incongruities with our talk of artifacts we are prepared to tolerate. I suggest that function claims in biology are best understood quite simply as claims about contributions to fitness. However, whatever option we choose, the failure of biological processes to yield a function concept that closely matches the connotations of artifact functions puts limits on the burdens such a concept can bear.

The work of this book also has an impact on debates about creationism. What falls out from its treatment of artifact talk in biology, in terms of the constitution of organic nature and the processes of survival and reproduction, is an explanation of the use of language admittedly laden with connotations of intelligent design. No intelligent designer is needed to make sense of artifact talk.

1.2 Why Is Teleology So Boring?

I will give a map of the structure of this book toward the end of the chapter. First, I should say a little about my choice of topic. There seems to be a feeling among many of the most prominent philosophers of biology that the problem of teleology is a boring one, either because it has already been solved, or because there is no real problem beyond being clear about what one intends when one speaks of "function"; or, because the debate is fruitless, consisting for the main part in the exchange of intuitions about whether one would use the word "function" in certain artificial imaginary scenarios.

Let me give three examples. First, Michael Ruse, in the paragraph that forms one of the epigraphs to this book, tells us that the problem of teleology "is worked out. Natural selection produces designlike objects and so function talk is appropriate" (Ruse 1996, p. 284). The idea that there is little left to say on the subject is supported by the fact that teleology is one of the very few topics in philosophy where there is anything resembling a consensus. Almost all contributions to the functions debate over the past twenty years have consisted in refinements of Wright's (1973) etiological analysis. Examples of such approaches include papers by Neander (1991a,b), Griffiths (1993), Kitcher (1993), and Godfrey-Smith (1993, 1994) to name just a few. Two recent collections—Buller (1999) and Allen, Bekoff, and Lauder (1998)—are dominated by etiological analyses. This said, the most recent work on the topic (Ariew, Cummins, and Perlman 2002; Davies 2001; McLaughlin 2001) shows signs that some are moving away from the etiological consensus.

The basic innovation of recent etiological accounts has been to supplement Wright's analysis with an explicit reference to natural selection, and recent papers tend to argue only over just what the appeal to selection should look like. So while Wright's analysis tells us that the function of some item is what it does to explain why it is there, newer etiological analyses tell us, basically, that a biological item's function is what tokens of that type did in the recent past that caused them to be selected.

For Ruse it seems that the functions question is a significant one, but it has become boring because it has been answered successfully. It is the nontrivial fact that natural selection produces designlike objects that means that function talk is appropriate. Had selection not had this character, function talk would have been a mistake. A comment by Elliott Sober hints at a second type of complaint: "If function is understood to mean adaptation, then it is clear enough what the concept means. If a scientist or a philosopher uses the concept of function in some other way, we should demand that the concept be clarified" (Sober 1993, p. 86). Sober's apparent fatigue is, like Ruse's, partly a result of the thought that the problem has been solved—after all, most philosophers and biologists seem to agree that the analyses of function and adaptation should match—but it also expresses some puzzlement about why we should think there is a serious philosophical problem of functions at all. We need only be clear in saying what we mean by "function" in some context, and that is that. Sober's problem, then, seems quite different to Ruse's. For Ruse, the problem of teleology seems to be the substantive one of vindicating a potentially illegitimate vocabulary. For Sober, it seems to be one of giving clarity to words that are ambiguous.

Finally, David Hull (1998) characterizes the debate somewhat differently again. Hull is bored because he thinks of the literature on functions as a form of *conceptual analysis*—a project which he characterizes as the search for the meaning of some phrase like "S knows that P." One philosopher proposes a set of necessary and sufficient conditions to capture the use of the phrase, and other philosophers respond by concocting more or less elaborate scenarios in which the analysis fails to match with their intuitions regarding whether the word should really be used in that scenario or not. The analyst responds by modifying the analysis, or by asserting, like Nissen (1997, p. 215), that the other philosopher's intuitions "are just wrong."

The functions debate certainly has been conducted in this way by some protagonists. One hears what Bigelow and Pargetter (1987) memorably call "the dull thud of conflicting intuitions" in the following passages from Wright and Kitcher, respectively:

> If a small nut were to work itself loose and fall under the valve adjustment screw in such a way as to adjust properly a poorly adjusted valve, it would make an accidental contribution to the smooth running of that engine. We would never call the maintenance of proper valve adjustment the *function* of the nut. (Wright 1973, p. 63)

> Unbeknownst to you, there is a connection that has to be made between two parts if the whole machine is to do its intended job. Luckily, as you were working, you dropped a small screw into the incomplete machine, and it lodged between the two pieces, setting up the required connection. I claim that the screw has a function, the function of making the connection. But its having that function cannot be grounded in your explicit intention that it do that, for you have no intentions with respect to the screw. (Kitcher 1993, p. 380)

These are not peripheral to the philosophers' accounts; the intuitions they express dictate how their theories of function are formulated.

Perhaps some forms of conceptual analysis are legitimate. Thus, on some views of the meaning of scientific terms, we see meaning as deriving from the roles of those terms in the theories in which they feature. Now on this view, to say what the role of terms like "function" is in biology is also to give an account of the meaning of those terms in biology. In this sense it is a conceptual analysis. And the intuitions about use, of those well versed in the theory in question—the intuitions of biologists and well-informed philosophers of biology—could be essential to recovering the role of the term in the theory, and hence its meaning in this sense. Still, these will be intuitions about biological cases; it is hard to see what role there is for intuitions about screws in machines in uncovering the meaning of the *biological* function concept.

So for Hull, the source of frustration with the debate as a whole is different again. Here it seems what is at stake is neither the vindication of a problematic vocabulary, nor the attainment of conceptual perspicacity, but instead the provision of an account of what some concept really means. Comparison with Sober's implicit project will make this clear. For Sober, the project of giving a meaning to the word "function" is achieved just so long as a clear definition is given suitable to biologists' purposes. Giving a clear meaning to a term is not the same as saying what the existing meaning of a term is. At least that is so on an account of meaning according to which we uncover meaning by trying to match actual, rather than recommended, biological use. That is why the two projects respond in different ways to examples of actual use and imaginary scenarios. If our business is giving clear definitions, then considerable tension with both actual use and intuitions concerning imaginary cases can be expected and tolerated. That is true even if we are concentrating on a specifically biological function concept and looking exclusively at biological usage. If, on the other hand, we are trying to outline what the word's existing meaning is, then we should try to match intuitions—and certainly actual use—far more closely. Hull's problem is that the methods available for carrying out this project seem weak. What are we to do when intuitions conflict? Whose intuitions should we respect? Isn't this form of conceptual analysis an *empirical* project?

If philosophers can get frustrated by the functions debate in such different ways, then it suggests that they have quite different conceptions of what the goal of an account of functions is, and what the proper methods are for attaining it. Is it simply a question of bringing clarity to biology? Should we also ask what biologists in fact mean by their terms? Is there any more substantive issue at stake about the nature of design in the organic world?

The project of saying what biologists mean by their terms leads us into thickets that we can happily avoid in our goals of understanding how teleological approaches in biology work, what risks they carry, what forms of teleological content can be grounded by biological processes, and why teleological approaches are found only in biological contexts. These questions elude the complaints of Sober, Hull, and Ruse, for they are substantive, they do not require the idle exchange of intuitions to be answered, and they have not been solved already. But we might now fear that in giving up on the project of exposing existing meaning as too difficult, or subject to idle comparison of intuitions, we are then pushed toward saying our project is merely one of stipulation or construction of meaning suitable for some purpose. Sober suggests that the debate over functions is merely one of what concept is most appropriate for biological use. Millikan (1984, 1989c), notoriously, considers her analyses of function to be exercises in stipulation to be measured by the work they do in philosophy of mind. In neither case does the project seem like an interesting philosophical one for what it says about biology itself.

In fact, so long as our project is understood broadly enough, there is a way of approaching the phenomena of artifact talk in biology that reduces to neither a dull exercise in the comparison of intuitions nor a dull exercise in stipulation. In the next section I give an imaginary example that helps us see how our questions should be tackled, and also why we can be silent on the question of meaning.

1.3 Meaning, Metaphor, and Methodology

Our problem is to understand why biology makes use of a vocabulary that seems ill suited to it, to understand whether that vocabulary is genuinely helpful, and to understand why biology, and not the physical sciences, tends to make use of that vocabulary. I would like to use a parable to explain how we should go about answering these questions.

For many centuries scholars thought that all of nature was invested with spirits who controlled the movements of rocks, trees, clouds, and so forth. They would speak of rocks, trees, and clouds in human terms, reflecting the intentions of the agents who were thought to reside within them. So trees would strive to attain the sunniest spots in the forest, rocks would race each other downhill in landslides, and clouds would chase each other across the sky. With the development of physical theory, most of the sciences abandoned this animist paradigm. Now geologists would no longer talk of rocks racing each other down hills, only of some falling faster than others. And botanists gave up speaking of trees striving to attain the light, preferring instead to understand their motions as the result of mechanical tropisms. Only the meteorologists persisted with the old vocabulary. Still they would talk of clouds chasing each other across the skies. And the models they produced, framed in that vocabulary, had great predictive success. Thinking of clouds as chasing each other allowed them to predict successfully that the faster chasers would succeed in their goal of catching and swallowing the slower clouds. The chasing paradigm in meteorol-

ogy seemed to work perfectly well as a model for predicting how cloud positions and cloud conformations would change in the skies.

Philosophers of science found the meteorologists' success puzzling. How was it that they were able to persist with a vocabulary that had been discredited by the rest of science? Could it be that clouds alone really did harbour some kind of intentional or pseudointentional states that explained the survival of the animist paradigm in meteorology alone? And what should they say about the small group of renegade meteorologists who argued that the animist paradigm was misleading, that meteorology should grow up, that it should go the way of geology and begin speaking of clouds in the sterile ways the enlightened geologists speak of rocks?

It will be helpful to keep this parable in mind throughout this book, for it makes clear how philosophers stumble into analytical dead-ends in thinking about function and design that they would recognize quite clearly in the case of our meteorologists. How should our imaginary philosophers of science proceed?

If our goal is to explain how meteorologists have been able to continue to use the language of "chasing" with predictive success, then we need to look to similarities between the motions of clouds and the motions of people to explain that success. It is because cloud positions are covariant in ways that resemble the positions of people who chase each other, that "chasing talk" remains in meteorology.

A philosopher who maps these similarities does most of the work in explaining the persistence of chasing talk in meteorology. Yet he has not provided any kind of analysis of "chasing talk." He has not given us some short formula of the form "Two clouds chase each other iff conditions C obtain." Suppose he looks at the relations of covariance between clouds and decides that meteorologists typically use the language of chasing when certain kinds of covariance relations apply. He then supplies an analysis of the concept "meteorological chasing," which gives these covariance relations as necessary and sufficient conditions. How should we understand this analysis? Does it tell us what the meteorologists mean by "chases"?

This is the kind of question that has led philosophers of biology into a dead-end. Once we know what kind of chasing concept meteorological processes are able to support, we can then compare that concept with the concept we apply to humans chasing each other. We will conclude that they differ in some respects, and that they are similar in others. Thus, we can answer the questions of how meteorology is able to make use of the chasing vocabulary, how the chasing vocabulary might mislead through associations with human chasing, and whether there are deep similarities or mere vague resemblances between the meteorological and human chasing concepts. We do not need to say what meteorologists mean in order to answer these questions. The answers we give are compatible with the thought that on the lips of meteorologists, the word "chases" has a strict technical meaning that should not be confused with the vernacular concept, or that it is a metaphor that happens to be useful, or that their talk is wholly mistaken and they really believe clouds to be invested with spirits. Trying to say which of these views about meaning is correct adds very little to our understanding of chasing talk in meteorology.

What is more, if our view of meaning is that "chases" is metaphorical, then the tightened analysis of chasing is of great value, since it gives the scientists a cleaned up concept that tells them what steps they need to go through to test claims about chasing, and so forth. If we think the concept is a technical one, then the analysis has value in a similar way; it helps tighten any looseness in discourse. So, whatever our view of meaning, the proposed analysis has value, and for similar reasons. The question of the status of the analysis—whether it cleans up actual meaning, whether it exposes actual meaning, or whether it creates a technical meaning to replace a metaphor—need not be answered. Note, however, that if all we do is give a stipulative analysis of "meteorological chasing" that is intended to be useful to meteorologists, philosophers, or whoever, then we fall short of answering many of the questions that interested us at the outset about what the relationship might be between this kind of chasing and normal human chasing, or how the chasing vocabulary in the meteorological realm might mislead through inappropriate connotations. That is why the stipulative project on its own is of limited value, unless supplemented by a contrastive exercise that ranges across the two domains.

Without some kind of comparison between artifact functions and biological functions we have no guarantee that biological functions have more than the most distant relationship to their artifact cousins. Compare: we notice that physicists use the word "charm" in connection with quantum particles, and decide to stipulate some meaning for what charm means there. It is clear that what we might call quantum charm has nothing whatsoever to do with romantic charm, and can support few of its connotations. What is more, unless we can show that biological function and artifact function are closely related, we cannot claim to have provided anything like a "naturalized" account of functions. In giving an account of the meaning of quantum charm we certainly do not give a naturalized account of romantic charm—this is so precisely because romantic charm and quantum charm have nothing to do with each other. To give an analysis of "biological function" in terms of wholly natural processes does not consist in a successful naturalization of function unless one can demonstrate that the new concept merits the name it bears.

The methodological stance that I have developed in the context of meteorology and physics goes for biology also. Whether we think that the meaning of "function" is technical or metaphorical, we need to examine the similarities between the processes that underlie the production of organisms and artifacts. This exercise tells us what kinds of function concepts biology can support, and how close they are to the function concepts we apply to artifacts. By looking at the roles teleological terms play in biological research and biological theorizing, we will also be able to construct a clear analysis of terms like "function" that will be beneficial to biologists. This is a valuable exercise whether we think the analysis exposes actual meaning or fashions a new meaning.

Note, finally, that my skepticism about the value of the exchange of intuitions, and about some forms of conceptual analysis, does not deny that thought experiments—even quite outlandish ones—may have value in understanding the significance of some set of concepts. Thought experiments can play a role in teasing out the similarities and differences between the kinds of concepts we apply to artifacts, and the kinds of concepts we can apply to organisms. Abstract thought experiments will play a role in my discussion of functions in chapters 5 and 6.

1.4 Metaphorical or Technical?

It is just as well that we do not need to say what artifact terms mean in the mouths of biologists, if only because this project would be so difficult to carry out conclusively. A decent case can be made, for example, for saying that when biologists use terms like "function," "design," and so forth, they mean exactly what the rest of us mean when we apply these terms to artifacts. This is what one might say who espouses a Davidsonian theory of metaphor, and who thinks that all artifact language in biology is metaphorical. On Davidson's (1978) theory of metaphor there is no distinct metaphorical meaning beyond the usual meaning of the words contained in the metaphorical sentence. Most metaphorical sentences, like "Cesare is a wily fox," are therefore false; however, they are used to draw attention to certain similarities—in this case, between Cesare and a fox. If those similarities themselves run deep enough, then they might explain why, for the most part, treating Cesare as though he were a fox might be a useful way to approach him. In biology also, we might explain the continued usefulness of metaphors of purpose, function, and design by reference to the deep similarities between the processes that go into the construction of organisms and artifacts. So this appeal to metaphor is one way in which Nissen's (1997) analysis of biological function statements, according to which biological functions are what some agent intends a trait to perform, might be able to make sense of the success of function talk within biology.

There are at least three prima facie reasons for thinking that much of the teleological language used in biology may be metaphorical, yet none of them is conclusive. First, as I have already remarked, there has been remarkably little change in biologists' use of teleological language over the past two centuries or more. In *The Blind Watchmaker,* for example, Dawkins (1986) demonstrates through both his choice of title and style of exposition that he regards Paley's *Natural Theology* (1802) with admiration, primarily for Paley's ability to expose the quality of design within nature.

Other natural theologians use language remarkably in tune with modern biology. The following passage from the fourth *Bridgwater Treatise* would not sound too unusual on the lips of a contemporary adaptationist:

> Shell fish have their covering for a double purpose: to keep them at the bottom of the sea, and to protect them when drifted by the tide against rocks. Animals of the molluscous division, which inhabit the deep sea, and float singly, or in groups, as the genus scalpa, have a leathern covering only: because they are not liable to the rough movements to which the others are subject, in the advancing and returning tides. (Bell 1837, p. 280)

Presumably, when organisms were considered to be artifacts made by God, function language had the same meaning regardless of whether one was talking about the function of a fork or a frog's leg. And function language continues to be used in the same way. It is a prima facie strength of the metaphor theory that it explains continuity of use in a simple way.

One who thinks that biological function instead has a quite technical meaning within the science, and who therefore argues against the metaphor theorist, can also explain resilience of use in spite of change of meaning through the mechanism of the *dead metaphor*. Dead metaphors, importantly, are not metaphors. When a metaphor dies, a word that was previously metaphorical loses its old meaning and acquires a new one. Wright (1976) gives numerous examples: when we speak of a jackknifed lorry, we do not speak metaphorically. It is testimony to the death of the metaphor that one can know what a jackknifed lorry is without knowing what a jackknife is. We could argue that terms like "function" are able to retain the same pattern of use in biology, because as biologists grow to realize that the nature of systems to which these terms are applied is different to what they had thought, they adjust the meaning of those terms to reflect that realization.

The second reason for thinking that teleological language might be metaphorical is that it is most often found in biologists' popular works. This might suggest that the primary function of teleological terms is to illustrate the makeup and history of organisms and parts of organisms for nonspecialist readers. The following two passages appear in *Nature*—a journal intended for a wide scientific audience—and in a work for popular consumption, respectively.

> If there are ways in which mutation can increase the probability of survival within cells without effect on organismal phenotype, then sequences whose only "function" is self-preservation will inevitably arise and be maintained by what we call "non-phenotypic selection." (Doolittle and Sapienza 1980, p. 601)

> Natural selection may build an organ "for" a specific function or group of functions. But the "purpose" need not fully specify the capacity of the organ. Objects designed for definite purposes can, as a result of their structural complexity, perform many other tasks as well.... (Gould 1980, p. 57)

These passages support the thought that any account which tries to give a tight analysis of what biologists mean when using terms like "function," "purpose," and "design" is misconceived. We could argue on the basis of these passages that such terms are not serious scientific terms at all. The words "function" and "purpose" are often placed in inverted commas. This suggests that such language is not intended to be understood literally, or that the users are suspicious of the propriety of their own terms.

There may be a good deal of truth in all this, but again, we do not need to drop the idea that "function," at least, is a technical term in biology. It is quite natural to assume that when writing for popular audiences, biologists would use words like "function" and especially "purpose" and "design" with caution precisely because they would not want them to be confused with the common language namesakes, and they would certainly not want their readers to think that they are committed in any way to the view that the organic world is the product of conscious design. Even if "function" is a respectable technical term within biology, one would not want a lay reader thinking one had the intentions of a creator in mind when speaking of the function of the panda's thumb. What is more, in the first passage the word "function" is used in scare quotes, in part because it reflects a nonstandard use in biology itself. The idea that selfish DNA has any function at all will sound odd to some biologists who are perfectly happy with the idea that other traits do have genuine functions.

The third reason for suspicion that function talk is metaphorical is, that although biologists may often speak informally of function and design, these terms are seen quite rarely in technical journals. One seldom finds straightforward claims about the functions of specific traits in such publications. Terms like "design" and "purpose" feature even less often. In a technical article by Kingsolver and Koehl (1985), often cited by philosophers in support of the claim that the functions of traits can change over time, the authors use the word "function" rarely, and they decline to make any explicit claims about the function of the insect wing. Instead they prefer to discuss the evolution of the wing in terms of its "adaptive value." Moreover, in those cases where the word "function" is used in this paper, it is most often found in locutions like "functions as," or "serves the function of." The authors thus distance themselves from paradigm statements of function of the form "The function of the wing is. . . ," and instead make claims like the following: "Elongation of the wings first evolved in small insects as a result of selection for thermoregulatory capacity, followed by an isometric increase—either gradual or abrupt—in body size, after which wings could function as aerodynamic structures" (Kingsolver and Koehl 1985, p. 503). Even here, the proponent of the technical term view can make a number of responses. First, as Allen and Bekoff (1995) note, the English language offers many ways of speaking of functions without using the word itself; technical papers often contain claims that are equivalent to function statements, such as "Parents remove white eggshells to protect their young" (Drickamer and Vesey 1992, p. 23). What is more, many research papers explicitly ask questions about function, even when the explicit answer given does not use the word. Allen and Bekoff cite a paper by Gordon et al. (1993) whose title is "What Is the Function of Encounter Patterns in Ant Colonies?" Gordon et al. make no explicit claim about what the function is; yet they do make a clear implicit function claim (p. 1099): "An ant that suddenly encounters alien ants may be in danger . . . the sudden increase in [antenna] contact rate, though short lived, may be sufficient to generate a defensive response to intruders." If function claims are not made explicitly, that is in part because of epistemic caution; in the paper by Gordon et al., the investigators simply are not confident enough to make firm claims. Yet it is clear that inquiry after function is central to many biological disciplines—most notably ethology and behavioral ecology. The question "What is the function of a trait or behavior?" is one of Tinbergen's (1963) famous "four whys?," and probably the one that behavioral ecologists, and more recently evolutionary psychologists, have been most strongly motivated to answer. A search through recent scientific journals yields a range of titles like *Functional Ecology, Cell Structure, and Function,* and so forth, all suggesting that the concept of function has a central role even in the technical practice of biology.

In summary, both the metaphor account and the account of function as a technical concept can be made plausible. What is more, it is not clear to me what methods one would need to choose between them. Since they are both intended as accounts of what a particular group of people means by some term, the best methods for adjudication would seem to be empirical. We would need to undertake interviews with biologists of varying types, subject a range of journals to textual analysis, observe language use closely in the context of lectures, seminars, day-to-day fieldwork, and informal discussion around the laboratory. And there is no guarantee that the meanings attributed would be univocal. The picture that might well emerge is one of a set of terms with different meanings in the mouths of different biologists. For some they are metaphorical, for others they have some more-or-less technical meaning so that "function," like "altruism," is a biological concept with a rather elastic connection to its common language namesake. The difficulties in answering the question of meaning, and the fact that answering this question is unnecessary to our main project, mean that I will remain agnostic on exactly what terms like "function," "design," and so forth mean.

1.5 What Lies Ahead

As the preceding discussion makes clear, the first task for understanding the presence and limitations of the artifact model is an investigation of the nature of the processes that explain organic form. This task is undertaken in chapter 2, where the received view of selection is outlined, together with its supposed relation to the phenomenon of adaptation. I argue, first, that the view of evolutionary theory as a theory of forces needs to be handled with care, lest we lose sight of the fact that natural selection and drift must be understood as population-level statistical phenomena. Second, I show that natural selection should be distinguished from selective forces, and that these selective forces can explain the emergence of adaptation only when they range over suitably organized entities. The upshot is that development plays as much of a role in the explanation of adaptation as selection.

Chapter 3 introduces the artifact model, and begins an assessment of the use made of it by the adaptationist program. I outline the adaptationist framework that conceives selection pressures by analogy with design problems, and traits by analogy with the parts of artifacts that are designed to meet such problems. I argue that the most common criticisms leveled against adaptationists do not, in fact, threaten the artifact model in general, for they highlight methodological difficulties in explaining and predicting the form of artifacts themselves. That said, there are a number of crucial disanalogies between selection and intention—most obviously in terms of the population-level nature of one, and the individual-level nature of the other. These disanalogies mean that artifact thinking can lead us to ignore drift, and also to underestimate the functional interconnectedness of organic, as opposed to artificial, design. I also tackle the more practical problem of whether artifact thinking—especially in the guises of reverse-engineering and adaptive thinking—faces epistemic difficulties. Here I suggest that evolutionary psychologists, in particular, should not expect the strategy of predicting adaptive solutions to problems laid down in ancestral environments to be of much use in uncovering the workings of the mind.

Chapter 4 looks to more radical challenges to artifact thinking from what we might term *constructivist* and *structuralist* camps. The first group, of whom Lewontin is the archetype, argues that no sense can be made of the crucial concept of an adaptive problem to which solutions might be developed, with the result that the concept of adaptation should be dropped in favour of recognizing a dialectical, constructive relationship between organism and environment. In response, I construct a concept of an environmental problem that can serve the purposes of the adaptationist program, while taking Lewontin's legitimate concerns into account. The structuralist camp—exemplified by Bryan Goodwin, but with allies in David Wake and others—argues for an elimination of teleological styles of argument altogether in favor of mechanistic explanations of form alone. I show that while a structuralist research program that looks to the explanation of form independently of adaptation may have considerable value, it is unlikely to wholly supplant functional biology.

Chapters 5 and 6 look to more traditional problems in the philosophy of biology regarding the nature of function statements and functional explanations. It is important to distinguish sharply between two questions. First, What is the best analysis of function claims in biology? Second, What explains why biologists make function claims but physicists and chemists do not? The parable of the clouds shows how such questions can come apart. An analysis of "chasing" in terms of covariance tells us the best way of cleaning up meteorologists' chasing talk. Such an analysis won't do to explain why only meteorologists talk about chasing. That question might be best answered historically, or by reference to the fact that clouds look a little like creatures, or by reference to the usefulness of the approach.

In chapter 5 I argue that the best analysis of function statements in biology is simply to think of the function of a trait as the contribution that tokens of that trait make to fitness. The concept is nonhistorical. Some (even most) philosophers say that only a historical function concept that ties functions of traits to their selection history can make sense of the normative and explanatory connotations of function claims. To these people I reply by suggesting that selection does not, in fact, meet these connotations particularly well, that the nonhistorical concept meets them well enough, and that in any case, it is not clear that all of the connotations that have been thought to be marks of teleological function claims should really be accepted.

Chapter 6 addresses the comparative question of why we find artifact talk in biology but not in physical sciences like chemistry. One might think that it is because only biological items are subject to a special force—natural selection—that gives rise to purposive states. Here I argue, on the contrary, that selection is neither necessary nor sufficient for the appearance of artifact talk. Inorganic "sorting" processes—the kinds of processes that sort ions bonding to the surface of a metal catalyst, or nuts in muesli, or stones on a beach—might also give rise to such talk. What is more, in cases where selection does not act, and where we might encounter artifact talk all the same, one cannot argue that such talk is mere "as-if" function talk, whereas biological functions are more genuine, purposive features of organisms. That is so because sorting processes support the same connotations—the connotations typically alleged to be the marks of bona fide functions—as selection does. The result, then, is that the account of functions in this book should be regarded as deflationary regarding the normative status of biological functions.

The final chapter turns the organism/artifact analogy on its head to look at the prospects for an informative evolutionary theory of technology change. Most commentators in this debate think either that the evolutionary view is false, even obviously so, or failing that they believe the successful application of evolutionary theory to technology will revolutionize the way we think about design, or marketing, or history, or economics. Neither of these views seems right to me. Artifacts do evolve, yet only a very abstract version of evolutionary theory that declines to comment about the broad character of selection pressures and the nature of cultural inheritance systems can be made to fit. The price for this abstraction is a corresponding lack of explanatory and predictive power when we try to apply evolutionary models to specific technological changes. In spite of all this, I end by outlining some possible lines for future investigation for technological evolutionists, and I show why looking seriously at how selection models can explain intelligent design will give discomfort to those creationists who want to contrast explanations of natural design that look to selection with those that look to intelligence.

The principles that we need to investigate to show us how the design of artifacts should be explained take us right back to the first theme addressed in this book—the relation between adaptation, selection, and development. Developmental organization itself is instrumental in generating complex adaptation. Hence an inquiry into the general principles of development and heredity may yield insights for the study of both organisms and artifacts. Our central analogy will remain ripe for investigation even when the book is done. Let us begin, then, at the end, with adaptation and development.