

From GROUNDS FOR COGNITION, Erlbaum, 1994

Main points:

- prebiological evolution by natural selection
- a naturalist explanation of goal-directedness
- life defined as goal-directedness
- rethinking teleology

Chapter 2

BY WAY OF MEANS AND ENDS

—

This chapter provides the teleological foundations for our analysis of guidance to goal. Its objective is to ground goal-directedness genetically. The basic suggestion is this. Organisms are small things, with few energy resources and puny physical means, battling a ruthless physical and biological nature. How do they manage to survive and multiply? CLEVERLY, BY ORGANIZING THEMSELVES TO ENGAGE AND EXPLOIT A USER-FRIENDLY ECOLOGICAL NICHE IN A GOAL-DIRECTED MANNER. Organisms evolve structures capable of coopting and directing causal processes originating inside and outside them so as to engage their habitat, and bring about specific, limited, and attainable end states, as goals. Goal satisfaction allows organisms to maintain and replicate their basic design by doing certain things in and to an environment. This is how goal-directedness came about, as an instrument and modality of maintenance

and replication of design, and why it was naturally selected as a successful policy.

Our hypothesis presupposes that evolution precedes, and is more general than, biological life. There are nonliving and nonteleological systems that evolve by natural selection. So the latter cannot explain, or reduce, goal-directedness. In its biological form, natural selection works precisely because, and to the extent to which, its targets, the organisms, are (already) goal-directed systems and their interactions with the environment are themselves patterned goal-directedly. Goal-directedness and its tools (such as guidance to goal) do not reduce to their biological implementations. Yet, since these implementations alone have (so far and naturally) evolved programs causally capable to shape goal-directed processes, including those of guidance, our naturalization of goal-directedness and guidance to goal must necessarily appeal to, and find illustrations in, genetics and bioevolution. The chapter has no ambition beyond providing a pragmatic and limited exposition and defence of goal-directedness, sufficient to establish its materialist credentials and legitimize the notion of guidance to goal. ¹

1. AN UR ROLE FOR GOALS

To flesh out our hypothesis, let us retrogress to ur and speculate. To survive and propagate, organisms must be able to husband their energy

¹ For comprehensive and sustained defenses of teleology, see Taylor 1964; Bennett 1976; and Woodfield 1976 for an early survey of the debates about teleology.

resources by generating or encountering states of affairs that help maintain their internal design, service its needs, and reproduce it. The organisms that fared better over time were those lucky to have evolved internal structures and processes that reliably identify the beneficial states of affairs and guide the behavior to them. I call such structures and processes means-ends (or m/e) structures and processes.

The ability to pursue specific ends by specific means enables an organism not only to specialize and focus its efforts but also to terminate them at some opportune point, and thus save energy and wear, when the results are good, or continue its efforts, or try something else, when they are not. The more efficient the ways to identify and get beneficial results, and the more accurate the information on the results, the better off the organism and its ability to spread its genetic heritage. We can think of the m/e structures and processes that accomplish these things as strategies by which organisms maintain a stable internal organization and replicate it. It is a clever evolutionary ploy to do that by way of pursuing finite and specific states of affairs in a means-ends format. Having and satisfying goals is the strategy of life. Life is maintaining and replicating an internal structure by means of doing something or other. Life is a manner of maintaining and replicating internal organization (the goal-directed manner) but not the only one. There are other forms of such maintenance and replication. This distinction is important to our argument since, by disengaging goal-directed life from other and simpler forms of maintenance and replication of internal organization through natural selection, we weaken the exclusive explanatory grip of natural selection, and therefore make some elbow room for goal-directedness and its (more ultimate) explanatory constraints on life and cognition. The disengagement argument goes as follows.

Simply maintaining and replicating an internal structure, without doing

anything else, is not yet life. There are nonbiological forms of matter (crystals) and concoctions (computer programs) which mutate, develop, maintain, and replicate their internal organization, and do so under the ax of natural selection. Crystals mutate, develop, maintain, inherit, and replicate their organization. They can do that because their organization comes in the form of repetitious physical patterns. The patterns give crystals their structural integrity and allow them to grow by incorporating atoms from the surroundings. Since in many crystals the atomic bonds between layers are weaker than those within a single layer, it is easy for layers to shear along the natural crystalline planes, thus replicating their internal organization (Cairns-Smith 1985; Dawkins, 1986, ch. 6; Casti 1989, ch. 2).

Significantly, crystals enjoy both random mutation of structure and the natural selection of the reproductive effects of such mutations, and so are subject to a prebiological evolution. Crystals often mutate structural defects that can be passed on to successors, as their layers grow and replicate. When the newly formed mutant crystals happen to survive longer and replicate faster than older versions (because of better sticking to the surrounding rocks, for example), thus displacing the latter and coming to dominate their corner of the world, we have a case of genuine crystalline evolution.

For a materialist this is as it should be, if both the emergence and properties of life are to be understood in prebiological yet evolutionary terms. The key parameters of systems that evolve, i.e., (i) the maintenance and (ii) replication of internal organization by means of (iii) random mutation (but not sexual recombination) and (iv) natural selection, antecedently characterize nonliving systems. It follows that these four parameters are not definitive of life, and hence cannot be invoked to explain its goal-directedness, or explain it away. Mere evolutionary explanations in terms of structural mutations and

natural selection cannot account for biological systems, generally, and for their conative and cognitive programs, in particular. This creates elbow room for the teleological difference in its earliest and most fundamental form. Life is teleoevolution.

What is specific of life is the form in which it instantiates the evolutionary parameters (i) - (iv). Goal-directedness offers an opportunity for bioevolution. The maintenance and replication of internal design are attained by specialized mechanisms. It is the latter that genetic mutations and recombinations modify, and it is upon the functional effects of these modifications that natural selection acts. It is because organisms, being goal-directed, do things functionally, not just physically or chemically, that bioevolution can do its job. ²

Bioevolution could not have a grip on crystals precisely because they lack specific means-ends processes whose functionally dedicated effects on the ecology, and then back on the crystals themselves, could be selected for or against. A mutation of crystalline structure is replicated directly, as a matter of physical law. Crystals do not do anything intermediate that could be one day modified by mutation and hence naturally selected for or against. Crystals evolve in virtue of what they are physically, whereas organisms evolve in virtue of what they do, goal-directedly.

The growth, maintenance, replication, and inheritance of internal structure

² This difference between doing things functionally as opposed to physically or chemically is cashed (theoretically) in the types of properties and laws that characterize and explain the doings. The difference is not one of material composition or causation, for everything in our world is compositionally physical or chemical, and everything that happens happens causally. The difference is that some material structures underly causal processes that have (dedicated) effects whose impact on a larger system of such structures and processes warrants their categorization as functional. The biochemistry of organisms is already functional in this sense.

by crystals are the result of mere instantiation of physical properties and laws. Although crystal mutation and replication may be implemented by means of intermediate physical processes, as anything in nature is, the implementation itself is not functional. No intermediate functional processes, served by specialized structures, operate in crystals (or other nonbiological forms) to generate sequences of outcomes whose final effect is to maintain and replicate an internal structure. The same is true of self-replicating organic molecules artificially synthesized,³ and also of computer programs that replicate themselves by passing their own spontaneous internal changes on to their offsprings.⁴ Neither crystals nor synthesized molecules nor self-replicating

³ The MIT chemist Julius Rebek has synthesized an organic molecule which replicates itself. Its constitution and shape are different from those of the biomolecules of life. The molecule is made of an amine (containing nitrogen) and an ester (acid and alcohol), and reproduces in chloroform (and not, as biomolecules do, in water). In a solution of like esters and amines, the synthetic molecule's ester will seize a free amine, and its amine a free ester. The loose ends of the captured compounds then fuse in an exact copy of the parent molecule that brought them together (DISCOVER, December 1990, p. 28). This, I humbly opine, is more than what crystals appear to do but less than what the molecules of life do. The synthetic molecule needs specific ecological assistance (the right solution, like compounds) and some simple functions (recognizing and capturing like compounds) to replicate, which is part of what life is. Yet the molecules of life do more than that, as they build cells and bodies which they must nourish and service and also replicate. The two tasks are connected since it is the bodies thus built and serviced that participate in replication. This is why (I am saying) life is maintenance and replication of something (the ur molecules) by doing something else (building bodies and having them reproduce in an environment by means of development and behavior). More recently (NEW YORK TIMES, August 27, 1991), Rebek's group was reported to have installed occasional errors that are passed on to the descendant molecules. Again, this is not goal-directed life, not because it does not involve the biomolecules typical of life, but because the replication itself does not require pursuing and satisfying intermediate goals by specialized functions (other than the copying functions).

⁴ I find the concept of artificial life coherent and instantiable. I also think that programs other than genetic, once installed in appropriate systems, can display (nonbiological) goal-directedness. Nevertheless, I find the current computer programs, for example, that of Thomas Ray, falling short of goal-directedness, hence life. Like the crystalline and organic

computer programs maintain and multiply their organization by doing something else, not just physically, but in ways whose functional specialty is to have such effects. This is the key difference, the key to goal-directedness.

It does not mean, however, that goal-directedness could not emerge from forms of matter other than biological or programmed other than genetically. On the contrary, this possibility is assumed and exploited by our inquiry. Computer programs could one day build and animate active, gregarious, autonomous and goal-directed robots the way the genes animate bioorganisms as program replication machines. Guidance to goal and its cognitive realization would then be part of the deal, sharing many of the constitutive and regulative properties of their biological counterparts. This was the leading insight of cybernetics and of von Neumann's theory of cellular automata, now revitalized by work on artificial life (Levy 1992). But artificial goal-directedness is only a promise, which is why a naturalization of goal-directedness is still bound to biolife. This is why the generality of my discussion must constantly bow to the specifics of biolife. This is also why, as we note in the next section and often afterwards, we must understand the limited role of bionatural selection very carefully, since inevitably and misleadingly it tends to portray the ecological and competitive coordinates of biolife (and hence the reasons for bioadaptations) as the only shapers of the cognitive and behavioral design of self-reproducing systems.

In its biological embodiment, goal-directedness is a natural product of evolution. It probably emerged as an accidental sequence of processes improvised by the organic compounds that must have exploited crystals into

replicators, Ray's computer program replication is not goal-directed since it does not rely on intermediate specialized functions targeted on states of affairs that secure their maintenance and replication.

becoming genetic formations.⁵ Since goal-directedness turned out to have adaptive results, it must have been naturally selected. Our analysis allows us to say, without circularity, that the evolution that initially generated and selected goal-directed organisms is prebiological; it is the same evolution that was acting on crystals and organic molecules. Once goal-directedness got into the picture, evolution took a biological turn, playing a new natural selection game. It is in this new game that bionatural selection acquires a grip on organisms in virtue of their being goal-directed, hence functional, and not merely in virtue of their having mutating structures that replicate.

Our evolutionary approach to goal-directedness clarifies another fact, often used against teleology. Prebiological evolution shows that the maintenance and replication of internal structure are not the goals of anything. This applies to the genes, too. Far from being genetic goals, structural integrity and replication are physical dispositions whose exercise has no rationale and no goals of its own. That is true of genes as is of crystals. It is the form of the exercise of those dispositions that differentiates genes from crystals or current versions of artificial life. The DNA replication, unlike the crystalline or artificial ones, is mediated by concrete end states, involving substance and energy, to be attained by the functional means of m/e structures and processes. Likewise, the RNA and proteinic processes involved in this functional mediation do not have their own goals either, and certainly not the goals of preserving and replicating DNA; the RNA and the working proteins have only specific functions (to run causal processes with dedicated effects) which consist in building and

⁵ Cairns-Smith (1985, ch. 14) has a story about crystals binding organic compounds which helped them survive and replicate. At some point, some of the organic structures managed to replicate faster than their crystalline hosts and become dominant by the rules of natural selection. See also Dawkins 1986, ch. 6.

servicing particular organs or running particular types of behaviors.

Still, one may ask, as properties of biological matter, aren't the integrity and replication of internal organization the goals of a large variety of m/e structures and processes? Yes, they are, but no longer in a pure, ur, inert, nonfunctional, and merely physical form. This is a crucial difference. In their biological version, the stability and replication of internal organization acquire specific and diversified expressions as (mainly) so many forms of metabolism and reproduction. The maintenance and replication of internal organization become goals by taking the particularized form of substances to be recognized and digested as food, of colorful and smelly conspecifics to mate with, of situations to avoid, and opportunities to encourage, and so on. The specific goals of life and the strategies to satisfy them are the biological implementers of the maintenance and replication of internal organization. Goal-directedness is an instrumentality, a clever strategy fortuitously evolved to maintain and replicate the design of some classes of physical systems. Biosystems are a subclass of self-organizing and self-reproducing physical systems, with the differentiating property of having goal-directed processes, and not just physical or chemical laws, doing the work of securing the integrity and replication of their design.

To summarize, then, in the beginning there were funny physical propensities for maintenance and replication of internal structure. These are the propensities manifested by any self-organizing and self-reproducing physical system, whether biological or not. Then there was biomatter exercising these physical propensities by way of means and ends, or goal-directedly. And finally, there was bioevolution working on particular biomatter-ecology arrangements. In their exclusively physical form, the maintenance and replication of internal structure are not the goals of anything, since they are not brought about by

any functionally specialized instruments; when, however, such instruments (qua goals and ways to satisfy them) evolve and become operative, the maintenance and replication take the specific forms of metabolism and reproduction associated with concrete designs in concrete interactions with concrete substances found in concrete environments. Biolife, if you will, is crystallization metabolized and reproduced with the help of amino acids, dirt, water, sun, sex, and so much more. ⁶

2. NATURAL SELECTION

Goal-directedness is the result of prebiological evolution, and is constantly

⁶ It is not yet clear how metabolism and reproduction came about, nor which came first. It does not much matter for our speculative story how these issues are going to be settled. Dickerson (1978, p. 73) has a good summary of the issue: "A living cell has two central talents: a capacity for metabolism and a capacity for reproduction. The cell survives in the short run by rearranging the atoms of the compounds it ingests into molecules needed for its own maintenance. It survives vicariously over the long run by being able to reproduce itself and give rise to offspring with similar biochemical talents. Which came first, a functioning metabolism, protected by some kind of membrane against dilution and destruction by its surroundings, or a large molecule that survived by making copies of itself from materials in its surroundings? In other words, which is older, the "protobiont" or the "naked gene"?" Neither Dickerson nor anyone else seems to have an answer. Once the life forms come into being, their metabolic and reproductive functions appear complementary: the instruments of reproduction (the nucleic acids) cannot replicate without the instruments of metabolism (the proteins) and the latter cannot be manufactured without the former. It is reasonable to think of metabolism and reproduction as continuous with and parasitic on the more basic physical propensities (like those attributed to crystals) for maintaining a stable organization and replicating it under exclusively physical constraints. There are of course many hypotheses about these matters, some quite fascinating. The one which I find attractive is the mineral/crystal hypothesis sketched many years ago by Bernal (see Dickerson 1978) and originally developed by Cairns-Smith (1985).

reshaped by natural selection. This relation between goal-directedness and natural selection is a complex and delicate one, which is why we have to work hard, constantly, to keep it in the right perspective. There are two related reasons why the explanatory limitations of natural selection matter to our argument. One (as section 7 below explains) is that biologists and cognitive scientists are uncomfortable with the notion of goal-directedness, and hopeful that natural selection alone could explain (ultimately) what goal-directedness does. What they hope for is either a reduction or an elimination of goal-directedness. I do not think this is either possible or desirable. It is not possible because natural selection presupposes goal-directedness. And it is not desirable because, without a teleological frame of explanation, we are bound to miss some pervasive and systematic patterns governing life and cognition which are not visible from the vantage point of natural selection. The second reason is that, intrinsically, the explanation by natural selection is narrow since it is limited to specific mechanisms and programs that operate in a corner of the environment and propagate in a specific population. On some issues, where the tree-forest difference matters, this limitation encourages either tree- rather than forest- viewing, or else modeling the forest in the image of particular trees. Guidance to goal turns out to be one such issue.

Before we unpack this theme, let me note vigorously that I am concerned here only with the explanatory limitations of natural selection in relation to goal-directedness and its subordinate phenomena such as guidance to goal. Natural selection has other limitations as well, for which it has taken a domestic beating in evolutionary biology and a foreign beating in cognitive science. The general charge has been that the explanatory range of natural selection has important limits because not all biological phenomena are touched by its shaping hand. There are forces other than natural selection which drive bioevolution: random

genetic drift, meiotic drive, exaptation (old adapted structures getting new functions), accidents that isolate organisms in an adaptive ecology, and so on. Also, various properties may spread in a population not because of their own adaptive value but because they are by-products of other adaptive structures; there was a selection for the latter but only of the former, to use Sober's distinction (Sober 1984a, section 3.2). Finally, as advanced organisms and their social and cultural habitats become more complex and versatile in their mutual interactions, it gets ever harder to explain their current patterns of performance in terms of the architectures originally sanctioned by natural selection.⁷

These explanatory limits of natural selection bear in many ways on the evolutionary story of cognition. (For a sample of the debates, see Dennett 1983, and ensuing commentaries; Dennett 1987, pp. 277-286; Lewontin 1990; Pinker and Bloom 1990, and ensuing commentaries; Stich 1990, ch. 3; Barkow, Cosmides and Tooby 1992). What is important for our argument is that all the forces of evolution, including natural selection, cannot fail to design organisms that operate goal-directedly and cognize because they must guide their behaviors to goals. Our argument is independent of which and how many evolutionary forces happen to be historically and causally responsible for the design of cognition, and of how adapted that design is. Our question is what it takes for a system to exercise guidance, given general constraints on its goals,

⁷ Just think of the fact that when last naturally selected (more than ten thousand years ago) the key cognitive programs of humans were not acquainted with either writing or mathematical figures and symbols. Then think how much your thought, consciousness and memory depend every minute upon these relatively recent accessories. Although a natural selection account could reveal what programs are used in thinking or memorizing, with what systematic tendencies and biases, which is no small accomplishment, it is quite unlikely that the account could identify and explain a variety of patterns of performance that the programs had picked up only later, when hooked up to the new accessories (for a stimulating review of these matters, see Donald 1991).

design, behavior, and habitat. Our answer will sketch several possible versions of guidance. How real organisms manage their guidance in real habitats then becomes a question about which of these possible versions, or which mixture of versions, they instantiate and why.

It is this latter question (which is not ours) that natural selection theories can answer, most of the time. It was Darwin's revolutionary insight to construe natural selection as the causal shaper and optimizer of functional design. Far from excluding genetic goal-directedness, as is often assumed, Darwin's insight about natural selection presupposes it. We recall that Darwin's view and the entire "evolutionary synthesis" of the 1930's and 40's, which is still its most popular reading, were completed before the major advances made in genetics, and the realization that there is prebiological and molecular evolution (Eldredge 1985, ch. 1; Mayr 1982, ch. 13; Dawkins 1986, ch. 6).

These advances and realizations begin to suggest that goal-directedness, genetic programming, and natural selection are compatible (Mayr 1982; Granit 1977), and that their compatibility has explanatory value. Both the work of natural selection and that of the genes are causal, although at vastly different scales of space and time. What goal-directedness brings to this causal picture is the format of functional configurations in which the genetic instructions are executed. In causally shaping genetic programs, natural selection is bound to act on their goal-directed policies already in place. Natural selection works on fully developed individuals and their reproductive capabilities. (As Trivers (1985, p. 20) notes, "[i]t is individuals that are purposively organized, and they are organized to leave surviving offspring"). Therefore, the purposive organization of individuals, including their propensity to reproduce, must be presupposed by the work of natural selection. Since, as noted in the previous section, natural selection also works on self-maintaining and self-replicating but purposeless

crystals or (current) computer programs, we must infer that natural selection is both blind and indifferent to how (by what means) the individuals subject to its attentions maintain and replicate themselves. This, again, is why natural selection cannot account for goal-directedness, which, we argued, is precisely a matter of the means by which organisms maintain and replicate themselves.

In the biological world, unlike that of crystals, natural selection designs programs that are good at doing something. Since the functional character of being-good-at-doing-something is (on the present account) a consequence of goal-directedness, the outcome of natural selection must be constrained by the antecedent goal-directed format and exercise of the programs in question. The constraining goes as follows: being good at doing something means capable of discharging a function; capable of discharging a function means capable of having a causal process yield a specific type of effect, with certain implications for the system's well being and operation; for reasons to be explored shortly, some of these types of effects can be best characterized as goals. This is why we can regard the genes as causally executing sequences of functions with goals as their end effects. This is also why we can regard natural selection as causally designing genetic programs that execute their goal-oriented functions adaptively. Nothing in this story puts goal-directedness and natural selection on a collision course.

3. TELEOGENETICS

The next step in our argument on behalf of genetic goal-directedness is to show that genes do operate goal-directedly. We can think of the DNA as a programming architecture, and of the RNA, the functional proteins it transcribes

from DNA instructions, and the proteinic outputs (tissues, organs, processes, behaviors) as an executive architecture. The genes are of two sorts, structural and regulator. The former work in development and behavior, the latter in timing and regulating the developmental and behavioral expression of the structural genes.

Together, the structural and regulator genes and the transcription controls are responsible for configuring the goal-directedness we find in organisms. The configuration of goal-directed mechanisms and processes emerges from a vast and intricate functional interaction initiated at the genetic and transcriptional levels (Beardsley 1991). There are no individual "master genes" for any particular goal-directed processes just as, and because, there are no individual "master genes" for anything significant in an organism, whether tissue, organ, function, or behavior. The latter are all the cumulative outcomes of numerous interactions among genes, input signals, and transcription proteins, and so is goal-directedness itself, given that the developmental goals are nothing but specific types of tissues and organs, and the behavioral goals emerge only from the execution of various functions and behaviors relative to ecological opportunities.

The configuration of goal-directedness is material and causal in virtue of the nature and organization of the (chemical) elements involved. I need the notion of configuring to identify the formation of a mechanism, or the deployment of a process, through the execution of genetic instructions. Goal-directedness is a matter of genetic and transcriptional programs causally/functionally configuring structures and processes along certain dimensions (the teleological parameters to be described in section 5 below).⁸

⁸ Mayr (1982, especially pp. 48, 56, 68) is the biologist perhaps most clearly committed to the notion that genetic programs work goal-directedly. He also takes the view, which I share,

The causal configuring along teleological dimensions takes place in several stages during gene expression: a DNA structure is copied or transcribed into a messenger or mRNA structure; the mRNA structure is translated into particular proteins, which then form tissues and organs, and run the organism. In the DNA to mRNA transcription the causal direction is to, not from, the DNA. It is not a DNA structure that is being produced as a causal effect. A DNA structure is already in place, but not in a form propitious for transcription and forward-looking causation. So processes must be marshalled to open the DNA structure up and ready it for transcription.

This is a critical point in our case for genetic goal-directedness. The DNA is not a causal initiator. It is a structurer, a shaper, a configurer, in the guise of a set of instructions about what to do when a causal initiative is taken by the environment or by other genes or internal processes.⁹ The program instructions have effects in configurations that render biological causation goal-directed.

Here is an example of teleogenetic configuration that yields guidance to goal. Plants live on solar energy. To satisfy its metabolic goals a plant must grow in the direction of metabolically helpful conditions. Solar light is both an

that the goal-directedness of genetic programs offers the ultimate evolutionary explanations of why organisms are configured as they are, and operate as they do.

⁹ Intuitively, the DNA may be likened to the laws of (say) physics which (to indulge in a Platonic fancy) are "consulted" or "read" as instructions by physical interactions as they shape themselves one way or another. The physical laws themselves are not causal in any sense; they are a blueprint for causation. It is the physical elements and structures in interaction, organized and animated by the laws, which have real causal potency. Just as the character of physical causation resides in the laws governing it, so the character of genetically controlled causation resides in the program instructions coded in the organism's DNA.

input to photosynthetic metabolism and an informational guide to growth. Photomorphogenesis is the process by which solar light acts on the genes and guides plant growth. Photomorphogenesis has three parts: registration of light by means of pigments; transduction of the light signal from pigment to gene; and induction of development through genetic regulation. The gene activation involves the usual transcription of DNA into an mRNA, which is then translated into proteins. It is through the modulation of the mRNA transcription that the light input guides the growth process by affecting both the choice of the goals of development (e.g., leaves, stem) and their actual form and size. The proteins executing the genetic instructions can either become structural parts of the plant or else perform functions that organize the form of the plant. This is how the DNA configures the whole process and its products.

The registration phase marks the beginning of guidance. It involves a pigment, called phytochrome, which activates the genes controlling development. The phytochrome operates as a sensing device that allows the plant to detect and measure variations in the quality, intensity, and duration of the light input. The light acts as a stimulus on the regulatory ambiance of a light-responsive gene, governs the expression of the gene, and modulates its transcription into mRNA. The pigmental guidance (registration and transduction) by the phytochrome steers the developmental functions to their goals. In so doing, the pigmental guidance aligns development to the relevant variables of the environment (Moses and Chua 1988).

The process just sketched is both causal and goal-directed. It is causal because it displays the right physical interactions in the right order; it is goal-directed because the functional interactions that are genetically configured converge in a dedicated and terminal fashion on an end state. This means that only certain types of causal processes (those functionally dedicated) execute

genetic instructions, and that the processes are called off (are terminal) when the end state obtains. The causal interactions involved can be said to implement a goal-directed configuration, as (2.1) reveals:

(2.1) environmental conditions --> light input --> registration (by phytochrome) ==> transduction (phytochrome conversion) ==> genetic induction (activation of developmental genes) ==> developmental processes by proteins under pigmental guidance ==> part developed (goal).

The occasion for the genetic processing is provided by an external (light) input. This is a triggering causation symbolized with simple arrows. While the input triggers the internal causal processes resulting in the development of some part (or, in other organisms, in behavior), it does not configure these processes. The latter task belongs to the genetic program. The DNA is a causal configurer in virtue of its structure and functional consequences. Its role consists in organizing the causal processes in certain structural and functional formats.¹⁰ This role of configurational causation is symbolized above with double arrows.

4. TELEOLOGICAL PARAMETERS

What is goal-directedness? Which are its dimensions? The teleological

¹⁰ In the mind naturalization literature, the distinction between triggering and configurational causation has been evoked by Fred Dretske (1987, pp. 42-43) and Jerry Fodor (1987, ch. 2; also his 1990a, ch. 5).

literature has identified a few that will do for our purposes.¹¹ I call them teleological parameters. I do not mean them to explain goal-directedness. That explanation must emerge from an evolutionary theory of the genetic expression, transcription and control, in the biological domain, and from a more general theory of self-organizing and self-replicating systems, in other domains. Rather, the teleological parameters should be construed as jointly characterizing a type of program expression, genetic or of some other sort. Nor do I see the teleological parameters as jointly constituting goal-directedness, since there are goal-directed organisms (or systems, generally) that fail to exemplify one or more of the parameters. There are, for example, simple organisms, such as bacteria, or simple systems, such as heat seeking missiles, which reach their goals by fixed functional routes, thus failing to instantiate the plasticity and versatility parameters. Having thus clarified the nature of our analysis, I will take a quick look at the key teleological parameters needed by our project, acknowledge a few notions and distinctions that will be useful later, and then allow the parameters implicitly define our notion of goal-directed system. The biological illustrations remain prominent, although the resulting definition transcends the biological domain.

A. Systemic function. It is important to distinguish two concepts of

¹¹ See, in particular, the comprehensive surveys by Nagel (1961; 1977) and Woodfield (1976). The latter (p. 40) mentions E.S. Russell (who in turn inspired Braithwaite) as the biologist who treated bioprocesses as goal-directed and analyzed goal-directedness in terms of terminality of effect, cessation of process when goal is achieved, persistence when it is not, and of alternative causal/functional routes in achieving the goal. These are key parameters of goal-directedness. None of these writers, with the exception of Mayr (1982), grounds goal-directedness in genetics, although Nagel (1977) considers yet rejects this possibility.

function (Nagel 1961, pp. 522-526; Lycan 1981; Enç 1982; Papineau 1984; Sober 1985). There is the formal concept according to which a function can be understood as a mapping or transformation of some input into some output in virtue of their intrinsic structural properties. Such is a function that takes an expression into another in virtue of the form of the expressions. Addition is an example. The formal function redescribes a local input-output correlation irrespective of a larger environment and an instantiating hardware. This is not the notion we need.

There is also a systemic (or system-sensitive) concept of function. This is the notion we want. Biological processes are systemic because so are the gene expressions responsible for them. A systemic function redescribes a causal relation in the light of the role of its effects in a larger system of parts and processes. Such a function is executed when the implementing causal processes have dedicated effects in the system. For example, the light impacting the retina affects causally its cellular membranes, blood circulation, and other physical and physiological conditions of the retina, yet none of these are visually dedicated effects, for they do not feed into the subsequent processes that produce visual images. There is no cognitively systematic relation between, say, blood pressure and visual representation. As a result, the latter is not a (systemic) function of the former, and the former does not have a function with respect to the latter -- although there could be various causal interactions between the two in both directions (the blood pressure may affect the chemistry of the visual process, and what one sees may affect one's blood pressure).

Historically, a dedicated effect of the systemic function of a structure can be construed as the sort of effect for which the structure was selected. Once goal-directedness is fully characterized below, the notion of systemic function

will also be called teleofunction.¹²

B. Dedicated, cumulative, and orthogonal patterns of functional contribution. We know that (and when) the effect of the execution of a systemic function is dedicated when we know how and why it fits into a larger sequence of effects which display some pattern of cooperation and cumulative contribution in a system, relative to the execution of a task. Various functional events in the visual cortex have cognitive significance only to the extent to which their contribution converges on the formation of an image. Systemic functions are organized in well patterned sequences. Cells differentiate at specific points as a result of the execution (causal process) of genetic instructions. As a cumulative result of many cells doing their thing, specific organs (dedicated effects) are built. We can draw a path that picks up these effects only in the dedicated pattern suggested, and ignores all the other numerous and concomitant or subsequent but undedicated effects produced along the way.

The patterns of functional contribution are not only dedicated relative to the effects to be attained, but also orthogonal. This means that the different causal paths that implement a particular function are independent of each other in that (within limits) the values of one path do not systematically covary with those of another (Nagel 1977, p. 273). An example, to be discussed in more detail later, is that of temperature stability. Several mechanisms could achieve

¹² A function can be systemic without being a teleofunction if it works in a system that is not goal-directed. The general distinction between function, including systemic function, and goal-directedness is essential. Many teleologists fail to make it, and take goal-directedness to apply wherever functionality applies. This failure weakens considerably the case for goal-directedness, as is evident from Woodfield's (1976, pp. 60-63) effective criticisms of Somerhoff and Nagel.

it. Their job is not temperature stability as such, but something more specific (perspiration, vessel dilation and contraction, etc.), and to that extent the mechanisms work independently of each other. At any particular point one of them is in charge, and the others are back-ups. Temperature stability is the dedicated effect that matters; the means to get it are many and independent.

13

C. Terminality. The execution of developmental and behavioral functions often terminates in a final dedicated effect that may take the form of either a structure (a tissue, an organ), a state (having the right temperature), a behavior or event (food grabbed, enemy avoided). A causal sequence of processes, viewed as the execution of a systemic function, is said to have a terminal outcome whenever (i) the process persists as long as it takes to maintain or attain a state or structure as dedicated effect, and (ii) upon bringing about that effect, the process is called off by a control or feedback program.

A terminal effect is not terminal absolutely (nothing in nature is) but only in the context of interactive systemic functions. Among the many (often uninstructed and distant) effects of genetic instructions, the terminal types stand out as those which, when produced, turn off the instructions or

¹³ As Ernest Nagel (1977) notes, orthogonality may be the most critical mark of goal-directedness. If we ask why a pendulum is not goal-directed, given that it always seeks to reach a stable state of equilibrium, the answer is that "the controlling variables of the ball's motion are not independent of each other, since the restoring force is proportional to the magnitude of the displacement force, though oppositely directed" (p. 274). Orthogonality and its systemic effects (plasticity and persistence of implementation) explain why goal-directedness as such is not exhausted by physical or chemical laws (which are nonorthogonal), although its specific implementations are. As Nagel notes, orthogonality is relative to a frame of goal-directedness: what is orthogonal in one may fail to be so in another.

deactivate interim functional links, thus making evident the systematic correlation between instructions and causal termini. Grown wings terminate genetically initiated and controlled developmental processes; further effects of having and using the wings no longer concern the "growing-wings genes," although they concern other sorts of behavior-instructing genes. Not only development but behaviors also have outcomes (getting to see something, grabbing some food to eat, or running away from predator) which are terminal with respect to some internal pattern of systemic functions.

D. Plasticity and versatility of means. Biological dispositions often show plasticity or versatility in how (i.e., by means of which causal/functional processes) they deliver their outcomes. A cell in movement to some destination (say, to build an organ) finds alternative ways to get there if confronted with obstacles. Likewise, animals often try various and sometimes novel behavioral maneuvers to avoid a predator, protect their offspring, or find food. We would not understand what the cell's or the animal's movements signify unless we relate them to an invariant terminal effect. The more complex an organism, the more plastic and versatile its behavior relative to some target of interest. There is another, equally familiar, and widespread manifestation of the invariance of dedicated effects. Many developmental and behavioral dispositions show persistence in deploying (often with versatility) processes needed to obtain assigned effects, often for as long as it takes. Not only are organisms persistent in some behavior, the way hungry predators persist in chasing their prey, and the prey persists in running away, but so are, on a smaller scale, the soldiers of the immune system or the developmental cells. Other things being equal, the processes involved tend to persist until an appropriate outcome state is reached, after which they are called off.

E. Control and Guidance. In both development and behavior, cellular formations, from the simplest to the most complex, display functions of control and guidance which provide registration, tracking, feedback, and evaluation, and thus steer dedicated processes to their anticipated outcomes, signal when such an outcome has been produced, ascertain how good it is, check with control or some other internal standard, make adjustments, and call off the functional processes involved, if everything is right. Why would such control, guidance and feedback mechanisms be selected for (or installed in some other manner) unless they govern processes that must reach recognizable outcomes? How would such control and guidance mechanisms know the identity of those outcomes and distinguish them from other undedicated concomitant or subsequent effects, unless the mechanisms are programmed to match and measure such outcomes against some internal standard or reference? The presence of the latter is perhaps the clearest indication that the system has specific functions whose performances in the outside world must be tracked and adjusted internally; if, moreover, the system is naturally programmed, as organisms are, to have such internal standards, that fact in turn indicates that its functions are likely to converge on definite classes of effects, its goals.

Functionally speaking, the simpler biological versions of internal control and feedback standards (variously called Sollwert, efference copies, or set points in the ethological literature) can be viewed as precursors of internal representations of goals (intentions, desires) in advanced species (McFarland 1983). Those simpler versions find their mechanical echoes in the control and feedback mechanisms whose design has been the main practical preoccupation of the form of engineering that culminated in cybernetics. For the philosophically minded cyberneticists (like Wiener, Ashby or Somerhoff), those

mechanisms provide the evidence that in principle both organisms and servomechanisms could be goal-directed in the same internal sense, by being appropriately configured architecturally and programmed functionally¹⁴ -- other things (i.e., other teleological parameters and additional conditions) being equal. But other things are not always equal, which is why we have to weave the teleological parameters into a conceptual framework that approximates goal-directedness in realistic terms.

F. The Hard Frame. Putting things together, we can say that goal-directed are those systems that (i) autonomously maintain and replicate their internal organization, or that of their components, by (ii) implementation means that (iii) jointly instantiate most (though not necessarily all) the teleological parameters. This characterization, which I call 'the hard frame', is admittedly strong since it rules out most of the usual suspects (thermostats, missiles, computers, robots, bodily organs by themselves) and some unusual suspects (occasionally entire organisms whose goal-directedness, either in a context or in general, is social, as is the case of ants or bees) while allowing in systems larger than we usually think of (societies). (i) and (ii) are the hard (explanatory) conditions of the frame, while (iii) is the soft (explicative) condition.

There are certain features of the hard frame that are worth emphasizing. As anticipated by our earlier discussion, the hard frame does not entail that

¹⁴ See Wiener 1948; Somerhoff 1974; also Woodfield 1976, ch. 11, and Nagel 1977, for surveys. As Bennett (1976, p. 61) points out, control by feedback is not necessary for goal-directedness. In real life, control by feedback, like plasticity of behaviors, results from increased complexity of design; yet simpler organisms lacking such capabilities are nevertheless goal-directed. I take the importance of control and guidance to reside not in their constitutive relation to goal-directedness but rather in their manifesting, clearer than other teleological parameters, the identity of the goals of the system or organism.

goal-directedness must be biological. The hardware does not have to be made of amino acids or crystals or chips or any other particular form of matter, although evidently it must be made of some matter with appropriate capabilities. This posture is compatible with the fact that in our universe only amino acids might physically implement genetic instructions, and also with the fact that a different form of matter might conceivably maintain and replicate its structure in ways other than genetic.

To be goal-directed, aggregations of matter must do things functionally to maintain and replicate an internal organization in some ecological niche in ways characterized by the teleological parameters. It is worth stipulating that they do so autonomously, by internal programming, since otherwise we may have to assume that there are larger goal-directed systems that designed and/or run them -- an implication which, from religion to philosophy of mind, has a way of obscuring the intrinsic properties of life and cognition. And in order to get a consistently autonomous operation by internal programming, we have to stipulate further that the systems under consideration enjoy some form (not necessarily biological) of life and evolution. This is the point of the hard conditions, (i) and (ii), of the frame: to ensure that goal-directedness is a natural phenomenon in the world. Most teleologists are happy just with the soft condition (iii), which allows goal-directedness without life and evolution. This leaves goal-directedness ungrounded naturally, and invites a behaviorist and exclusively analytic reading of teleology, which I find neither realistic nor very explanatory.¹⁵

¹⁵ The soft condition had provided the standard frame of analysis in the most exuberant behaviorist period in recent teleology, 1940's to 1970's, in the works of Braithwaite, Somerhoff, Nagel, Taylor, Dennett, Bennett, and others. The quest for necessary and sufficient conditions favored a symptomatic (soft) rather than explanatory (hard) analysis. Examples and counterexamples of isolated teleological parameters were thought to settle

The hard frame does not go into the specifics history or natural selection, and does not entail an adaptive (or indeed, any particular) implementation of goal-directedness. All of these factors, crucial particularly for the biological forms of goal-directedness, are needed to settle two critical questions about goal-directed systems which the hard frame does not address. One is the origin of the system. The other question is the identity of the goals of the system, and by implication the identity of the targets of the system's behavior and cognition. These identities cannot be determined without information about the history of the system and the specific ecology in which it operates. That is not the concern of the hard frame. The hard frame simply characterizes goal-directedness, without saying how and when it works, and what it works on. In our world, history and natural selection characterize the biochemical implementation of goal-directedness, and to that extent they specify how an organism's design was shaped in and by a given ecology, hence what its goals are, and what teleofunctions, including cognitive, service the goals.

5. GOAL AS OUTCOME PROGRAMMED

The notion of goal is commonsensical, not scientific. Our commonsense talk

the matter. Not surprisingly for conceptual teleobehaviorism, the central issue was the evidence for goal-directedness, the reasons for attributing it, on the basis of observation and analytic imagination. In spite of having bravely promoted the cause of teleology, identifying most of the relevant parameters, and working out the methodological constraints on the attribution of teleological properties, the analytic teleobehaviorists failed to make the strongest case for teleology. This failure in turn opened their enterprise to needless, often obscure, and eventually detrimental conceptual skirmishes (see Woodfield's (1976, chs. 3-6) meticulous diagnoses of the failures of analytic teleobehaviorism). The enterprise was ultimately ungrounded because of genetic (or internal programming) neglect.

individuates goals either in terms of an end or outcome condition (action, its object, a state of the environment, or a state of the body) generated by some internal and/or behavioral processes, or in terms of some internal indicators of goals, such as urges, desires, or emotions. I will have more to say about our commonsense teleology in chapter 8, section 1. Now I need its form of individuation to approximate the teleogenetic notion of goal implicit in the hard frame characterization of goal-directedness.

The suggestion is to count as goals those types of conditions that are the terminal effects of functional processes which (i) are configured (though not necessarily initiated) by internal programs, and (ii) jointly satisfy most if not all of the teleological parameters. Goals are the sorts of conditions that place a system (or a part of it) in a position to do something, undergo some process, exercise some function, operate some program. An organism undergoes some internal process (perspires) or does something (runs) in order to be in a position to undergo (lower its temperature) or do (eat) something else which it is internally programmed to seek and bring about.¹⁶

Goals, so understood, explain the interface of various programs: the operation of one program (say, cognitive) places another (behavioral) in an operating position, and so forth. As long as the programs are chained in a rigid functional sequence and directly express genetic instructions, as in the simplest organisms, we can talk of basic goals. In the biological world, these are outcome conditions with immediate metabolic or reproductive import. For example, unicellular organisms have cognitive and behavioral programs exclusively

¹⁶ 'In order to' is an explanatory locution indicating that this is how we explain the organism's design: that it is so configured that it can position itself in a behavioral-ecological condition to undergo or do something. Causally, of course, it is always some input that activates some program that discharges its function and steers behavior.

dedicated to the identification, tracking, and appropriation of substances for metabolism. The operation of these programs relies on functionally specialized proteins that directly execute genetic instructions.¹⁷

A basic goal is an outcome type of state internally programmed (genetically instructed in biolife) -- that is, a type of state at the receiving end of dedicated functional processes (internal and behavioral) and at the initiating end of informational transactions that eventuate in feedback, cessation of effort or adjustment. This is not a commonsense notion by any stretch of the imagination, although its formulation exploits a commonsense gambit. We individuate concepts or actions in a relational sense by indicating their referents or results, respectively. We do the same with the notion of basic goal: instead of referents of concepts or results of actions, we take goals as outcomes that individuate the basic and intermediate programs (genetic or otherwise) responsible for them. This is, for example, how we redescribe the genetic instructions in teleological terms as 'being for wings, legs, or other organs and behaviors.' We can schematize the notion of goal as outcome programmed as follows:

(2.2) input --> programs ==> internal processes ==> behavior ==> (basic) goal

Complex organisms reach their basic goals by first reaching intermediate or instrumental goals. The latter feat requires the interaction of various programs. To be in a position to eat, an organism must first find food, say, by identifying and pursuing a prey, which in turn requires its own cognitive and behavioral

¹⁷ What I mean here by direct execution of genetic instructions is that the proteins (the key functional executors in all organisms, and the only ones in the simplest of them) are fabricated as soon as the RNA reads the appropriate DNA instructions, given the stimuli.

programs. Most organisms execute the programs for prey identification and prey pursuit with teleological brio: functional dedication, persistence, versatility of means, control, guidance and feedback, terminality of effort when success is achieved or when failure is inevitable. This is evidence that the prey programs have their own instrumental goals. Once the instrumental goals are reached, other basic or instrumental goals and their programs take center stage, playing their own teleological scripts. Instrumental goals are defined by specialized programs that implement specific cognitive or behavioral policies, only ultimately conducive to maintenance and replication. Instrumental goals indicate a higher functional complexity. Their individuation scheme looks as follows:

(2.3) input --> specialized (cognitive, conative, behavioral) programs ==> behavior ==> instrumental goal

6. THE RELATIVE UNIMPORTANCE OF CAUSAL AND FUNCTIONAL REALIZATIONS

Several of its parameters suggest that goal-directedness is causally and functionally underdetermined, essentially and often on a massive scale. I have in mind the orthogonality, plasticity and versatility of goal-directedness, the relativity of what counts as a goal-directed system, and, consequently, the higher explanatory importance of the dedicated effect (goal) over how it came about (by what causal/functional means). As a result, causes and functions, and their laws and patterns, lose their eminent explanatory power in the domain of goal-directedness, even at the proximate level. (That will be clear in the case of

mental cognition, in chapter 7). Teleology, notoriously, downgrades the role of functional causation. Being underdetermined by triggering and implementing causes and functions, a goal-directed system can be best individuated and explained in terms of the dedicated effects (goals) it is configured and programmed to produce. The partisan implication is that goal-directedness cannot be reduced to or eliminated by nonteleological facts and laws. This implication motivates our opposition to the reductive policies of psychological Newtonianism and psychosemanticism.

I consider first the orthogonality, plasticity, and versatility of implementation, and the resulting explanatory importance of the outcome, and then the relativity of framing goal-directedness.

Think of an animal fleeing a predator. Within its capabilities, the animal does a number of things (runs, hides, pretends to be dead) to bring about the goal state. Each of the behaviors is a causal sequence of functional processes triggered by appropriate stimuli. Each sequence is ultimately made of atoms and molecules pushed around causally. So are its external stimuli. The behaviors can be individually framed and causally explained, in a proximate fashion, as motions-triggered-by-stimuli. Functional redescriptions are of course much better but not totally revealing unless we find their common denominator in the form of escape behaviors. Escape is the goal that explains these behaviors in a uniform manner. The ultimate explanation in terms of goals grounds the proximate and local explanations in terms of specific causal/functional processes, and portrays the latter as implementations of goal-directedness.

Another example, much discussed in the teleological literature (e.g., Braithwaite 1953, ch. 10; Nagel 1961, pp. 409f; Somerhoff 1974; Collins 1984), is that of homeostasis, in general, temperature stability, in particular. An increase in the temperature around the body triggers an increase in perspiration

which in turn causes a decrease in the bodily temperature. The proximate causation at work here can be superficially schematized as follows:

(2.4) external heat increase --> increased bodily perspiration -->
decreased
bodily temperature

Two facts are critical for the relation between causation and goal-directedness. One is that the environmental changes in temperature can and very often are so drastic that (on a purely physical and proximate account of causation) they ought to cause equally drastic changes in the bodily temperature, yet they don't. The temperature compensation moves within narrow limits. The other fact is that there are several independent ways of maintaining temperature stability. Orthogonality and plasticity again. The rate of perspiration is one way, the dilation and contraction of blood vessels is another, and so are the secretion of various glands, the adrenaline stimulation in the blood, and the muscular contractions. What we have here is temperature compensation (a) within pre-established limits and (b) by various means. (a) and (b) characterize the range and the variety of internal bodily reactions to outside heat increase, relative to preset internal parameters. So (2.4) should read:

(2.5) external heat increase --> internal compensation mechanisms
activated
==> decrease in bodily temperature to preassigned limits ==> internal
equilibrium maintained

Although it is the causal chain described in (2.5) that brings about the outcome (maintenance of a preset internal temperature), the proximate causation involved in the operation of one particular mechanism (say, perspiration) explains neither the existence nor the nature nor the rationale of the outcome. Other causal chains, involving other inputs and internal mechanisms, could have had the same outcome. This means that the sameness of outcome cannot be retrieved from and explained by the analysis of the interactions and mechanisms that instantiate the causal production of the outcome.

The explanatory demise of (functional) causation in teleological contexts affects every form of implementation along the hierarchy of levels of complexity we care to consider. I am stressing this point because of a widespread propensity to confine the notion of implementing causation to a territory ranging from physical to neural states, which is the territory of hardware, and show that the hardware states par excellence fail, in their diversity of instantiations, to reveal the unity of teleological phenomena. I want to push this failure further up in the functional territory of software. All biofunctions redescribe causal processes and interactions of some sort or another, in the light of their dedicated contributions and effects in a larger system. Since they have a limited role and range, and locally carry out larger biological tasks (metabolic, homeostatic, cognitive), under a narrow and proximate analysis, these biofunctions may also fail to reveal the nature of the tasks they execute.

The explanatory demise of causation thus extends to functionalized causation (whence the 'causal/functional' combination in my terminology). In particular, as I argue in chapter 7, although the symbolic and syntactic (or connectionist) redescriptions of brain causation define functions that drive our

mental mechanisms, the redescriptions fail to reveal the unity and character of the cognitive activities these mechanisms implement (which is mental guidance to goal) to the same extent to which, deeper down, descriptions of electrical synapses or calcium transfers among cells fail to reveal the unity and character of the symbolic (activation) structures being syntactically (vectorially) processed.

The implementations that underdetermine goal-directedness are not only vertical (inside organisms) but also horizontal (among organisms). This is the frame relativity of goal-directedness. The contours of goal-directedness are not fixed forever, either geographically or temporally. Consider a mundane example. I want to go to the market. That is my goal. I know where it is and how to get there, by walking. I am a self-contained and self-sufficient goal-directed system. If I don't know where it is, and have to ask somebody for instructions, or if I have to drive, or if I am driven to the market by somebody who knows where it is, then I am much less self-contained and self-sufficient as a specific goal (market)-directed system. For, in these latter cases, either the knowledge of the goal or the means to satisfy it or both transcend my resources. I cast around for other (eco-social) resources that can help. My goal-directedness is framed (expanded) relative to a context of behavior and what it takes to have it bring about the goal state. All our lives we navigate from one frame of goal-directedness to another. No single frame is absolute. This is true of many species, and has to do with how the genes implement their policies. For many species, the frame of goal-directedness need not always be fixed in advance; for complex species, the frame can change from one context of behavior to another.

There is no standard way of being goal-directed. Whatever ensures the maintenance and replication of genes will do, and the space of opportunities and

the manner of exploiting them are wide open. Bacteria are the simplest organisms with autonomous goal-directedness. Bacteria also commune, qua cells, in larger organisms, often only temporarily. In this alternation of sociality and individualism, no posture is definitive; it depends on goals and means. There are good reasons for bacteria to live communally, in colonies, more like cells in an organism, and specialize functionally. In many ecologies, small means vulnerable -- being eaten, crushed, not able to ingest or digest the only stuff available because of its size or hardness. Some microbiologists think that most bacteria are social most of the time, and that their social organization, development, metabolism, communication, and movement are genetically regulated (Shapiro 1988). Social goal-directedness may be more adaptive than the individual one. ¹⁸

We know that many species are goal-directed only socially (mole rats, bees, ants); others are socially goal-directed from time to time in specific contexts (hunting, migration). What about us? Is our social goal-directedness temporary, as my earlier market example suggests, or, deeper down, intrinsic and permanent? Are we in fact sophisticated teleofunctional cogs in larger goal-directed ensembles (families, societies, cultures)? When we behave independently, aren't we in fact implementing the work of larger teleounits (like cells in larger bodies, or insects in colonies)? Doesn't our own development and learning enable larger social teleounits to operate in ways that preserve and replicate their own structures? Don't we mature intellectually to preserve and replicate the cultural heritage of the tribe, society, species? And, like social insects, don't we divide our competence and labor along functional specialities

¹⁸ "By becoming multicellular, an organism can preserve all the advantages offered by cells for efficient metabolism and proper gene distribution and at the same time become very large" (Bonner 1988, p. 62).

such as professional cognizers (the scientists, physicians), professional food growers (the farmers), professional regulators (the politicians, judges) or professional memorizers (story tellers, historians)?

No matter how we answer these questions and draw the boundaries of goal-directedness, it seems clear that the answers would not be intelligible in causal or functional terms. Wherever evolution has rewarded a flexible and context-sensitive framing of goal-directedness, with constantly changing boundaries, no causal explanation in terms of natural selection would capture how we go about implementing the social forms of goal-directedness. Natural selection ultimately explain the social forms themselves, in terms of their adaptive virtues, but cannot explain proximately the ways in which the social forms are managed by the individual goal-directed organisms.

*7. THE MERITS AND HANDICAPS OF TELEOLOGY

It is time to wrap up the case for goal-directedness by confirming the naturalist credentials of our teleology and exposing some of the misconceptions that obscure its truth and explanatory relevance. Our conception of goal-directedness meets the adequacy conditions on mind naturalization set in the first paragraph of section 3, chapter 1. I divide the conditions into metaphysical, having to do with the natural ingredients of goal-directedness; scientific, having to do with how goal-directedness fits into the scientific picture; and logical, having to do with thinking straight about these matters.

The metaphysical condition requires that teleology be about facts, properties, and processes which (i) have material composition, (ii) are caused by other material items in the lawful flow of causation, and (iii) have a natural

history. The metaphysical condition, as an expression of materialism, prejudices neither the type of material items nor their form of organization nor the causal form of their internal processes and behaviors nor finally the particulars of their evolutionary history. These matters are for science to fathom.

All the variables of our account pass the metaphysical test: the genetic program, itself encoded in chemical structures, is responsible for a form of systemic causation that embodies the teleological parameters; the causation flows in the right direction from material structures and processes, and results in other such structures and processes; the genes have a respectable (though still mysterious) material pedigree in the ur properties of mineral formations; through their physical impact on organisms, the evolutionary mechanisms regulate the symbiosis between genes and their ecologies. And the whole spiel, blind and opportunistic, has no goals of its own, nor is it ruled by any design intent. Materialism, causation, and natural history are respected by our teleology.

The logical condition stipulates that a teleological account be coherent and noncircular. Its key notions and propositions should not appeal to those of goal or goal-directedness. Likewise, the notions of goal-directedness, needed to ground and explain guidance to goal, should not presuppose the latter. And the notions of information tasks and cognitive programs should not in turn presuppose what they must ground and explain (semantics, representation, cogitation). I am happy to report that both the general logical condition and its specific implications have been scrupulously obeyed.

Our teleology does not beg any historical question since it does not assume that either the (possibly mineral) pedigree of the genetic program, or its various local mutations and variations under environmental pressures, or the evolutionary selection of the genetic expressions, are themselves goal-directed

and hence in need of a teleological story of their own (no circularity and no regress); and our teleology does not beg any metaphysical question since it does not assume that the chemical components of the genetic code and the various executive and regulative functions are themselves goal-directed and in need of a teleological account.

The scientific credentials of the teleological stance are quite another matter. Even with good metaphysical and logical credentials, teleology cannot easily find a scientific home these days. For, given the way the condominium of science is currently set up, and the restrictions on who can live there and how and where, teleology faces a serious housing discrimination. This discrimination gives elbow room to biopsychological Newtonianism, and reinforces many of its prejudices. This is why we have to sweat a bit, not only with argument but also with some historical reflections on, and complaints about, the current enterprise of science, to get teleology the scientific respect it deserves. Time for deconstruction.

The question we want to ask is why teleology is invisible to science. The question is not about nature and telos, for nature is organized and does operate in a goal-directed fashion at the biological level. Organisms are teleosystems. It is an undeniable truth that the biofunctions executed by molecules and other causal/functional pushers have dedicated effects, cooperate, persist, and desist, and do so in an enormous variety of types of biosystems, in an enormous variety of ecologies, subject to an enormous variety of selective pressures. The teleological truth shines through all these diverse manifestations. Explicating this truth, however, is no easy matter. Goal-directedness can be objectively invisible. Evolution implements the goal-directedness of life opportunistically, in small, improvised, disjointed, and proximate steps (by way of mechanisms and programs), without displaying

pervasive and homogeneous properties and regularities. So we cannot speak of a uniform and well structured domain of goal-directedness, with its own properties and laws. Nor, as a result, can we expect a distinct and homogeneous type of teleological explanation, as a distinct form of scientific explanation. The problem, then, is how to translate the teleological truth and display its unity in current scientific discourse. It is a problem not about the facts but about their understanding.

Most philosophers and scientists either do not see the point of teleology at all, or else, when they see it, are uncomfortable about it.¹⁹ J.B.S. Haldane is reported to have said that teleology is like a mistress to a biologist: he cannot live without her but is unwilling to be seen with her in public. His analogy aside, Haldane was right. The facts about means-ends structures and processes are there, solid, visible, stubborn. It is their early conceptualizations that messed things up. The traditional logos of goal-directedness, from Aristotle's to Lamarck's to Bergson's to some current views, construes goal-directedness in a wrong and often ridiculous way, in terms of such notions as the goal of evolution (or of the genes), final causes, nonexistent futures, preformationist blueprints, backward causation, élan vital, dualism, emergence, optimality, rationality, progress, and so on. It is such notions that have given teleology a deservedly bad reputation. Yet it is a fallacy to infer from this conceptual

¹⁹ With a bit of terminological prestidigitation, I could have avoided (as many theorists do) any teleological talk. Instead of goals and goal-directed processes, I could have talked of structure and function, and still get all my conclusions. After all, looked at from below, implementationally, a goal is nothing but the effect of the causal execution of a number of functions by specialized structures. But that is precisely the angle I urge not taking: from below, the more global properties that characterize goal-directedness are not visible, yet those properties are essential to, and explanatory of, life and cognition. The reduction of teleology is a loss of explanation.

nonsense to the nonexistence of goal-directedness in the world.

I take Haldane's insight to signify that whereas teleology is the right stance for the biologist to take, there is no clear way to translate and legitimize its wisdom and explanatory power in the current conceptual framework of science. Goal-directedness is too global a phenomenon to be visible to, and digested by, contemporary science in general, biology and cognitive science, in particular. The means-ends structures and processes operating in an environment range over very diverse properties and regularities across extended and complex patterns of causal and functional chains. To explain fully the development of an organ, for example, one has to proceed from consideration about environment and evolution to genetics to cell biology to developmental biology, thus crossing several well established and often tightly compartmentalized disciplines. If the organ is cognitive, psychology must be brought into the picture. The disciplinary differences would prevent such an overview.²⁰ One can do molecular genetics without worrying much or even having to know much about cell behavior or organ development. And one can work in any of these areas of biology without having to know much if anything about psychology.

Yet the fact is that not all of nature's joints fit neatly into the disciplinary confines of current science. Nature is not that obliging. This is particularly true (is bound to be true) of those of its joints, such as the biological ones, which

²⁰ Think, for example, of vision or language. Their architecture has genetic roots, a developmental stem, of the embryonic sort, and the baroque variety of branches representing the mature specialized programs. Each level (roots, stems, branches) contains different forms of organization, of different complexities, governed by different functional regularities. We would not adequately and ultimately understand the biology of the visual and linguistic programs without crossing the various disciplinary domains. Yet we have no current scientific paradigm with which to attain and regiment such an interlevel and crossdisciplinary understanding, some valiant efforts notwithstanding.

not only build on and exploit a variety of lower-level and widely distributed physical and chemical properties and regularities, but also systematically engage the environment in an unlimited variety of ways. There are biological properties which are not visible if we examine only their implementing structures, and do not track their overall functional interlocking across several forms of organization and levels of complexity, within and outside the organism.

Embryonic development, for one, is beyond the disciplinary grasp of molecular biology because many embryological properties range not only over intergenetic interactions but also over complex subsystems, such as cell formations, organ structures, and even organisms. Embryonic development is a topic the molecular biologist is trained not to think and worry about. Yet embryonic development, sandwiched between genetic instructions, at one end, and organs, at the other end, instantiates an important segment of the teleological descriptions. How could a molecular biologist decide whether such descriptions are legitimate? And why should he be an authority on these matters, when he was trained not to see them? Often specialized scientists do not care even about the truths of neighboring disciplines. Talk to a geneticist about organ development, and he will soon begin to yawn; talk to an organ development theorist about behavior, and, guess what, yawning again. Put all their yawns together, and you have yawned not only goal-directedness but many other trans-domain properties out of existence. With all due but irreverent respect, if current science does not talk and think about it, and does not even know how to think about it, it doesn't mean that it doesn't exist. ²¹

²¹ Here is Robert Weinberg (1985, pp. 55-56), in an otherwise upbeat review of successes in molecular biology: "It is still far from clear that attempts to reduce complex systems to small and simple components, pushed to an extreme, can provide adequate insights for coming to grips with the great problem biologists confront today: describing the overall functioning of a complex organism. Can the biology of a mammal be understood as simply

Given the current organization of science, we have no other choice but regard teleology as a meta- and mega-theoretical stance that reflects the joint but implicit wisdom of a coalition of scientific theories. The teleological stance is megatheoretical since it extends across disciplinary boundaries, from genetics to cell biology to developmental, behavioral, and evolutionary biology, and then on to neuroscience and psychology. The teleological stance is metatheoretical because its concepts and axioms do not belong to the principled vocabulary of any one single scientific discipline. The teleologist is left to engage in a sort of metascientific hermeneutics in order to locate, deconstruct, and assemble teleological truths out of the intradisciplinary fragments and proximate models found in the domain-specialized literature. The result is not a neat picture but then the goal-directedness of life is not neat either.

This diagnosis of teleology affects our project in several ways. Teleology was said to make both an ultimate and proximate difference to our understanding of the cognitive mind. In the area of ultimate difference, the mega/meta stance of teleology is reasonable and fruitful. About the goal-directedness of life, as a means/ends strategy evolved by the genes, we do not yet have an established science (witness the recent debates about artificial life); a teleological account offers as good an insight into the organizational and functional properties of living systems as anything on the current scientific

the sum of a large number of systems, each controlled by a different, well-defined gene? Probably not. A more realistic assessment would be that the interactions of complex networks of genes, gene products and specialized cells underlie many aspects of organismic function. Each gene in an organism has evolved not in isolation but in the context of other genes with which it has interacted continuously over a long period of evolutionary development. Most molecular biologists would concede that they do not yet possess the conceptual tools for understanding entire complex biological systems or processes having multiple interacting components."

market. The scientific resolution might emerge from a future science of self-organizing and self-replicating systems that enter into dynamic arrangements with their environments. I expect this resolution to be teleology-friendly.

The same is true of our teleoevolutionary taxonomy of forms of guidance to goal. We do not yet have a comparative evolutionary biology of forms of cognition, evolutionary psychology is still in its infancy, and the top-down style of analysis has barely begun to make its impact in these domains. Our teleoevolutionary profiles of forms of guidance are meant to contribute to this emerging enterprise. How a teleological account makes a proximate explanatory difference in understanding mental cognition is a matter for the reader to decide after seeing the argument of chapter 7, and its implications in chapters 9 and 10.

One last defensive word about teleology before we put it to work. I want to stress that the explanatory value of the (genetically based) teleology used here is in the same methodological boat with the special sciences from chemistry up, and particularly the functional sciences. The reason we have special sciences, in addition to physics, is that there are emerging patterns of complexity and nonbasic or regional laws (of molecules, wheather patterns, organisms, societies, and minds) which we simply miss, and cannot conceptualize properly, if we look at everything from below, with the theoretical eyes of elementary particle physics. This is a familiar point. I think it also applies to teleology precisely because it applies to biology and psychology. The latter two comply with causal explanation but deviate from physics by being generally functional (biology) or functional with respect to information (psychology). Goal-directedness redescribes complex types of functional patterns. That is all. So there is no point in trying to replace teleological talk with a mere physicalist talk about particles and causes, because that would hurt biology and psychology

first. The real Trojan horse in the Troy of physicalism is the phenomenon of function. Once there are functions, there will be some patterns of them (explainable under, but not eliminated by, natural selection) that exhibit goal-directedness. As noted a few pages ago, I could have talked about such patterns nonteleologically yet functionally, and still get all my conclusions, but what's the point of not calling a spade a spade? Political correctness in science? For those still scared by teleology, this advice: any time you see 'goal' and cognates, remember it's just a pattern of functions satisfying the teleological (oops!) parameters. A quick mental translation, and you are with us.