Chapter 1

Setting the Scene: Naturalism and the Prospects for Evolutionary Epistemology and Reason

Introduction

The business of philosophy, as I conceive it, is providing a general understanding of the nature of life and of how to live it. We humans go about understanding some facet or domain of life by creating theories about it, testing these out in experience, and reflecting on all this, we slowly elaborate both a conception of the domain in question and a conception of how we know (what we think we do know) about it. I shall describe this process as theorizing a domain or subject matter. Self-reflexively this is the naturalist theory of understanding itself, which I employ.

This is a rather radical naturalist conception of cognition, and of philosophy in particular—just how radical will be seen during the course of its defense in section 1.I below. But first the major aims of this book are briefly introduced below and placed in a larger research context. Subsequently, as remarked, an evolutionary naturalist realist philosophy is outlined and argued and, following on from that, a correlative approach to an evolutionary epistemology (section 1.II). These discussions are necessarily somewhat condensed and programmatic, taking their meaning and force from large bases of research and argument that cannot be explicitly reproduced here. Readers unfamiliar with these may like to read this chapter lightly now and return to it after having read the remainder of the book.

One of the immediate consequences of my naturalism is that philosophy should be deeply engaged with all areas of theoretical and practical understanding. For neither the problems nor the (best conjectural) solutions are given in advance, but only emerge from actual historical engagement. And that brings me to the first and most important theme of this book.

We are today, I believe, at the beginning of an important and fundamental revolution in the conceptual foundations of all the sciences, one with important consequences also for the professions: the shift from linear, reversible, and compositionally reducible mathematical models of dynamics to nonlinear, irreversible, and functionally irreducible complex dynamic systems models, especially for complex adaptive systems (which include all living systems). This claim is not my own; a substantial part of it has been spelled out by Jantsch 1980 and Prigogine and Stengers 1984, among others, complemented by Dyke 1988. As little as a decade ago its character and consequences were discussed only by a handful of the mathematically knowledgeable; today there is a veritable explosion of new literature pursuing these ideas in every field from irreversible thermodynamics and chaos theory in physics through engineering control theory and complex adaptive systems theory in computer science, self-organization and hierarchy theory in biology, and dynamical neural net theory and genetic algorithms in neurophysiology and cognitive psychology, to evolutionary economics and international relations, with increasing cross-application among these. It would be neither practical nor desirable to attempt to review these exciting and increasingly interrelated technical fields here; the interested reader is directed to the asterisked items in the references for an introduction. They provide what the pioneers of the general systems literature of the 1960s and 1970s (Banathy, von Bertalanffy, Margalef, et al.) could not provide: the beginnings of detailed, powerful models of nonlinear self-organizing complex adaptive dynamic systems. (Cf. Margalef 1968 or Odum 1971 with Caplan/Essig 1983, Hannan/Freeman 1989, Peacocke 1983, Weber et al. 1988.) I am a long way from having understood and integrated the flood of systems literature of the past two decades, and with some works that are seminal and popular I nonetheless await the outcome of further examination (see the sequence Brooks/Wiley 1986, Morowitz 1986, Collier 1986, and Hooker 1984 on Prigogine 1980). In this book I am only concerned with trying to develop the right kind of philosophical theory; a detailed exploration of the accompanying science is beyond its scope. (That hard work belongs to a future book and to you, dear reader.)

For all its value and conceptual importance, the ideas deriving from this revolution have as yet scarcely touched philosophy. For the most part, philosophers still model rational agents (explicitly or by tacit presumption) in terms of simple logical structure, and the whole of science likewise. Individual rationality and rational scientific method alike, for example, are standardly seen as reducible to logical inference of some kind; the psychology and sociology (economics and politics) of decision making are at best formally irrelevant, merely part of the "implementation" of the formal program; at worst are the sources of causally interfering factors reducing rationality. But those who think in these terms are effectively locked into the dominant analytic philosophy paradigm—philosophy as logical analysis—inherited from the positivists and logicists early this century and expressed in the artificial intelligence (AI) model of cognition from the 1970s, the simple formal logical programming machine. There are various reasons why this is so, some of them deriving from evidence and the attraction of simplicity, some deriving from presumptions about the nature and status of normative principles, and no doubt some deriving from the attractiveness to a discipline of having clear ownership of a domain. (This latter is actually based on illusion, for today much of the creative work in logical theory comes from mathematics and computer science, not philosophy—further grist to my naturalist mill.)

Whatever the reasons, accepting this dichotomy has had unfortunate consequences. Theoretically it either falsely places psychosocial scientific explanation in a category (causality) wholly distinct from that of philosophical epistemology (logic) with the practical consequence that the two kinds of practitioners often regard themselves as offering competing rather than complementary accounts, or it reinforces the constraint to formal symbol manipulation (logic/language of thought) among cognitive psychologists and philosophers, dividing them from other practitioners. But scientists acting rationally are thereby acting socially, and cognition is far more than narrow formal symbol manipulation. One upshot is that at present many are falsely led from a critique of narrowly formalist logical/AI theory to the conclusion that rational epistemology is bankrupt. The theoretical result has been rampant relativisms, a confusion of context sensitivity with context incommensurability, and a profusion of supposed social replacement models ("conversations," "alliances," "stakeholders," etc.); this is accompanied by the practical danger that what we have of reasoned public life will be replaced by demagoguery and narrowly self-interested force.

Instead of leaping for an easy or politically correct relativism, I suggest that we must search for a better theoretical framework providing a richer class of models of action. Specifically, *the working hypothesis for this book is that the currently most fruitful and viable framework is that of dynamic nonlinear self-organizing complex adaptive systems*. The general idea is to reconstruct philosophical ideas, in particular rationality and epistemology, in terms of characteristic adaptive processes in such systems. In this conception, roughly, people are represented as strategic adaptive systems pursuing complex goals, and social processes are understood in terms of the self-organizing dynamics of many-person systems. To this end I propose that we adopt a strategic or decision theoretic conception of cognitive agency whose basic component is the epistemic utility increasing strategic decision in response to a problem posed in a particular decision context. This allows the explicit introduction of problem context to cognitive theory and so an explicit role for social structure, in particular a central role for the institutionalized social structure of science in scientific rationality. And it imports the decision theoretic framework of social context-dependent strategic interaction among rational agents as a basis for a dynamics for science. The complex interactions within these systems include both belief and goal formation and re-formation and structuring and re-structuring of roles/processes. All these processes occur within and between all system levels from subindividual to whole-society and now to whole-species, and they derive from interactions both within and between all system levels.

In this way science can be modeled as a dynamic self-organizing complex adaptive process embedded in wider social and biological complex adaptive systems, and the specifically cognitive is to be recharacterized as (roughly) the information-extracting and organizing aspect of these dynamics, not as either a separate "level" or part of it. Rationality can then be re-theorized in terms of the (self-) organizational properties of such systems; it includes, for example, the capacities to systematically collect and evaluate environmental information of various kinds and to systematically adapt context-specific goals as empirical understanding accumulates. Epistemology can be characterized in the same way, for example, in terms of the satisfaction of various systematically constructed and reconstructed inquiry procedures and the resulting social construction of statements and procedures invariant across various systematic contexts. In this way, I aim to employ our best current scientific models of living processes to provide a principled reconstruction of our conception of science. On the one hand, this conception is to understand science as a specialized case of biological dynamics generally. On the other hand, it is to be one that recognizes the essentially social character of science and provides for a cognitive sociology of science—though without recognizing sociological study as privileged or foundational vis-à-vis other disciplines, and while retaining normative epistemology. *The main aim of this book is to argue an initial case of this kind for the philosophical concepts of reason and epistemology and to show how the reconceptualization might fruitfully be developed.*

Some philosophical work of this kind has already begun. Hooker 1981b placed the concept of reduction within a complex systems framework and distinguished between compositional and functional reduction. That distinction, crucial to discussions of reduction in biology and the social sciences, still does not seem to be very widely understood. (In philosophical biology, for example, intermediate or interfield theories, such as operon models of aspects of genetic regulation, are urged against reduction; they certainly do show the inadequacy of a simple compositional/type reduction, but from a regulatory perspective they simply correspond, for example, to modular functional analysis in electrical circuit theory or cybernetics, all of which are strictly compatible with ontological deflation.) Cummins 1983 provides a thorough analysis of the philosophical concept of functional explanation along these same lines, extending systems functional analysis to cognitive psychology, artificial intelligence, and the relation of both to neurophysiology. (See Hahlweg/Hooker 1989a Part IV for a summary of Cummins's work and its marriage with the theory of systems reduction; this account is briefly sketched in chapter 2, section 2.III.1 below.) Over the past two decades, Paul Churchland has developed naturalized analyses of cognitive concepts (such as that of cognitive agent as epistemic engine), which, although not explicitly set in a systems metaphysics, are particularly appropriate for integrating therein (see Churchland 1979, 1989), and Pat Churchland has similarly produced a major reevaluation of the fundamental concepts appropriate to understanding and integrating a neurophysiological theory of mind (Churchland 1986). And Giere has introduced simple decision theoretic models of scientific decisions and related them to an account of theories themselves as models to yield a much richer conception of scientific method (Giere 1979, 1988), but again without much attention to embedding this account within dynamic systems. However, recently the philosophical specialization known as evolutionary epistemology has proven a fruitful locus for the introduction of functional systems ideas to epistemology (see especially Campbell 1974, 1986, 1987, and generally Hahlweg/Hooker 1989b). Indeed, in the form I shall give it in section 1.II below, its task is precisely to understand science as a dynamic process. So, *a second major aim of this book is to provide a re-working of evolutionary epistemology in the new systems framework.*

These researches are only a few of many that could be mentioned; they are simply examples where I have knowledge derived from personal involvement. My own attempts to begin thinking in these terms were first stimulated by reading Piaget in the late 1960s. But there was then a dearth of good scientific models (it was at the dawn of the discovery of chaos and before the seminal works by Morowitz 1968 and Odum 1971 appeared), and Piaget was under attack from all sides (partly for good reasons, but often, it turns out, for bad reasons; see chapter 5). So I turned instead to an elaboration and defense of the kind of general philosophical naturalist realism that I believed (correctly) would be required as a defensible framework for the rethinking of basic ideas when it became more fruitful to do so (see Hooker 1987, chapters 2–-5, originally published 1974–1976). The limitations of empiricism were clear, and I was preoccupied with generalizing its critique to a matching analysis of Popper, Feyerabend, and others. An extension of that naturalist program, provides *the third major aim for this book: to show how an adequate philosophical conception of naturalizing epistemology and reason can be developed.*

By the early 1980s the required systems revolution had recognizably begun (see further below). A paper like Hooker 1987, chapter 7 ("Understanding and Control"), became writable. The wheel has now come full circle, and the conceptual tools now exist to understand Piaget and to integrate his work into the larger philosophical and scientific enterprise (chapter 5 and Hooker 1994d). The common problems of empiricism and Popper have also come into correlative focus as centering on the formalist logical models they employ, and it is now possible to provide a pointed critique of

Popper's evolutionary epistemology as well as his general philosophy and that of Feyerabend in a precisely complementary manner (Hooker 1991a). Indeed, it becomes possible to extract from Popper the beginnings of a very different theory of science from his official one, one suited to integration into the dynamic systems framework (chapter 3). In this way the philosophical rethinking presented in this book is self-reflexively a model of the developmental systems processes described therein. One can only hope that the result here also is widening theoretical adequacy and cognitive autonomy.

An enormous amount of work, mathematical, and scientific, but especially philosophical, still needs to be done on the fundamental regulatory systems concepts, for example, on the notions of functional versus compositional entropy and their relation to Shannon/Weaver information and to semantic information. Indeed, these and other fundamental concepts, such as self-organization, are as yet relatively poorly understood (not surprising in a young research field) and in urgent need of conceptual analysis. So the reader is not going to find a detailed theory of cognitive dynamics in this book. At this time we have neither the theorems nor even maturely developed appropriate conceptions of reason and cognition. Yet we are in a better position than our mentors, who nonetheless succeeded in providing us with key ideas: that cognitive dynamics is the key problem and is driven by knowledge begetting problems that disturb the present situation and so lead to new knowledge, that the connection of internal truth criteria to external success is complex but central for an adequate account of epistemic progress, that progress is to be linked to competence in a wider range of environments, and that it is also linked to self-organization and autonomy since it corresponds to a capacity to preserve internal operations invariant across a wider range of environments. Some or all of these themes are already in Popper and Rescher and especially Piaget (see quotes PQ1, 2, 19, 20, 22, 31 from chapter 5), and they form the backbone of the systems evolutionary epistemology espoused in Hahlweg/Hooker 1989a. They also have a long history in the scientific literature (cf. Margalef 1968, Buss 1987, Smith 1992, and their references). We do them most justice by trying to reunderstand their significance in the light of contemporary advances in knowledge.

Science is a complex system; because of that complexity we cannot reasonably expect detailed but general theorems about cognitive systems dynamics even when the subject has matured, if by that is meant being able to deduce from general principles precise behavior in individual contexts. Until recently established science had developed satisfactory detailed models only for systems with either just one or two components or with large numbers of components interacting randomly so that local detailed fluctuations could be ignored, whereas we are dealing with systems of many components all interacting nonrandomly in complex ways. While we have recently gained some insight into these latter systems, it has basically been insight into their general dynamical features; to obtain detailed models requires specialization to the contextual details case by case. Consider this simpler parallel: No one doubts that quantum mechanics applies to bridges, yet no rational engineer tries to write down Schrodinger's equation for a bridge as a basis for its engineering design. Rather, engineers continue to rely on simpler approximate models while using quantum theory selectively to shape general concepts (e.g., metal fatigue and surface corrosion), to predict relevant kinds of structural features (e.g., the role of impurities in tensile strength), and in this way reshape and place refined limits on the applicability of the simpler models. We should expect to deploy complex adaptive systems models in a similar way. Traditional logic/AI models of science were apparently detailed, but only at the cost of wiping out nearly all relevant contextual and goal structure, reducing all scientists to ciphers for an abstract rational mind. Complex adaptive systems models will be less sweeping, but only because they recognize a greater wealth of detail as relevant.

The research presented in this book is part of a larger philosophical program of re- conceptualization, which can be conveniently summarized in this way, where S-ORS refers to self-organizing regulatory systems:

LOGICAL CATEGORY	MAIN TASK
1. General Philosophical Doctrine	General Theory of the Nature of Reason
2. Specific Philosophical Doctrine	Reason as S-ORS
3. Psychosocial Implementation	Cognitive Psychology of Rational Processes; Theory of Rational Institutional Designs and Processes
4. S-ORS Models	Specific Models of Cognitive, Institutional Processes in S-ORS Terms
5. S-ORS Theory	General Theory of S-ORSs

Hooker 1991a focused first on level 1 theory, outlining the argument for the failure of formal reason in philosophy of science, and then turned to sketch a systems alternative at level 2. The latter is taken up in this book, which is focused at levels 2 and 3. The former level 1 project is to be taken up in Hooker 1994c and at length in Brown/Hooker 1994, which also develops a richer general philosophical account of reason with which to replace the formalist version. Other work at levels 3–5 is also underway on biologically and cognitively implementable control and cognitive modeling (Hooker et al. 1992a, b; Hooker 1994e), on cybernetic modeling vis-à-vis S-ORS models (cf. variously Churchland 1986; Churchland 1989; Cunningham 1972; Lloyd 1989; Powers 1973; Thagard 1988 with Jantsch 1981: Jantsch/Waddington 1976: Pines 1988: Stear 1987: Yates 1987) and on foundational concepts of S-ORS, but further discussion will be left for another occasion. Let us be clear, the hardest work is yet to come. When readers finish this book, I would like to have persuaded them that a naturalized regulatory systems approach is the most promising route into the future and that it has the support of much of the best of past theorists as well. But then it is necessary to create and investigate detailed working models of biology, cognition, social dynamics, AI, and engineering control, and so on, so that the promise of the ideas is put to the test. That exacting task, though well under way in the sciences, is largely yet to come for philosophers.

The basic naturalist realist philosophical and metaphilosophical framework within which this book is formulated is set out in the next section below, followed by a brief sketch of the advocated approach to evolutionary epistemology. The former is necessary because naturalist realism has radical consequences for the formulation of the philosophical enterprise itself and readers need to be clear about these before proceeding. (Self-reflexively, it is itself a fallible theory, but one for which the analyses of this book provide heartening support and promise of future fruitful engagement with science. In another context, the views developed here can be used to defend the account against relativism; see Hahlweg/Hooker 1988.) The latter is undertaken because I believe that the currently dominant models of evolutionary epistemology, though they do move traditional epistemology one step in the right direction, are not themselves appropriate for proper integration within a dynamic systems framework and it is necessary to set out my own account for the reader.

Following this, chapter 2 provides a first sketch of a dynamic systems account of science. This is a first attempt at theory construction, an attempt to develop enough concepts for the description of science as a dynamic system to motivate and direct further developing the approach. There is an enormous amount of work yet to be done, both by way of incorporating detailed systems models into the analysis on the one side and by co-opting useful philosophical analyses into the model on the other, and finally by working out more detailed historical applications to test the ideas. But I believe that there is already presented here a sufficiently substantive framework to be able to limn the outlines of a new paradigm for philosophy of science that can be presented as a principled alternative to the standard logicist/AI paradigm. In particular, after laying out some essential preliminaries, concepts of objectivity and scientific progress are developed in terms of the unfolding of regulatory systems organization. It is not, note, an attempt to develop an alternative normative epistemology in which the old logical rules (e.g., of induction) are replaced by some equally simple formal rules connecting the new concepts. Rather, normative epistemology is to be redeveloped within this dynamic account by relating norms to the fundamental capacities of the system processes to extract and systematize information. As noted earlier, epistemic norms can be characterized in terms of the satisfaction of various systematically constructed and reconstructed inquiry procedures together with the resulting social construction of statements and procedures invariant across various systematic contexts. From various angles most of the remainder of the book after chapter 2 is devoted to exploring this approach. What emerges is a conception of science as an abstraction from—an aspect or dimension of, but *not* a part of—a highly complex, highly interactive dynamic system of nested subsystems ranging all the way from internal neural organization to social institutions. As noted an important consequence of this is the capacity to provide a principled cognitive sociology of science, as epistemic system design; this provides the proper theoretical ground for those who wish to join the critics of traditional logicist/AI theories of science and yet do not wish to join the strong program sociological, Marxist, deconstructionist, Feyerabendian, and other relativisms in their flight from normative epistemology.

From this perspective it is satisfying to be able to take three evolutionary epistemologies regarded as friendly to rationalism, those of Piaget, Popper, and Rescher, and show how they provide fundamental insights for naturalized theories of evolutionary epistemology and reason. (Similarly, I argue in chapter 6 that Putnam's attempted rejection of naturalized reason only succeeds in providing desiderata for an adequate regulatory systems naturalization.) In Popper's case a thorough reconstruction from minority fragments is required, but required internally in any case by the failure of his theory on its own terms. Rescher's theory requires serious renovation but not reconstruction. Piaget requires only stripping and polishing to produce a vigorous naturalist regulatory systems framework. It is their amenability to naturalist reconstruction of the regulatory systems kind that accounts for their order of treatment. These three studies stand in their own right, and as well they contribute to developing a regulatory systems philosophy of science.

The point of chapter 3 is to take an important and well known philosophy (not to mention one from which I gained all my own initial conceptions of critical reason), namely Popper's, and use it to demonstrate (1) how metaphilosophical attachment to the logicist/AI model can defeat even the best intentions of an evolutionary epistemology and (2) how within someone as honest as Popper is in facing the functional reality of science one can still find and extract a more appropriate philosophical framework. A similar lesson applies exactly to Rescher, who is nearly unique among advocates of evolutionary epistemology in pressing the original formal analogy at least as far as incorporating one order of methodological dynamics. As with Popper, we may set aside the failures caused by attachment to the logicist/AI model and generalize the insights contained in his analysis to form an essential component of an adequate epistemology in a dynamic systems framework—see chapter 4. Indeed, there is a natural progression from the concluding analysis of Popper's alternative account of the control of decisions to Rescher's account, for the latter begins to furnish the control structure that Popper's rudimentary account left undeveloped.

And that brings us naturally to Piaget, where the Popper-Rescher account may be embedded in a still wider complementary framework. Piaget devoted his life to creating an integrated regulatory systems framework that would provide a unified account of knowledge from biological self-organization to the most esoteric reaches of science and mathematics. We shall be prepared by the preceding chapters to reread Piaget afresh in these terms and to appreciate the systematicity and fruitfulness of his conception of genetic epistemology. It incorporates an appropriate evolutionary epistemology within it as a special case and does so in a way that provides for a unified naturalist biologically based account. And as with Popper, we shall discover in Piaget a much neglected aspect, in this case a philosophically neglected account of the nature and formation of regulative ideals and their norms (now read literally as regulatory systems functions). Piaget's stripped and polished account, enriched by those of Rescher elaborated and the reread Popper, provides the foundations for a new paradigm of epistemology in a dynamic systems framework.

Epistemology is applied rationality (namely, rationality applied to information acquisition, whether sensory, revelatory, rationalist, or whatever). The notion of reason in the dynamic systems framework will have been running as an undercurrent throughout these chapters, so it is only appropriate that the book conclude with a study of the naturalization of reason and a discussion of its retheorizing in the new framework. It turns out that with naturalization appropriately understood, the allegedly strongest arguments against naturalizing reason instead simply provide so many desiderata for doing so. And the analyses of the preceding chapters provide a set of clues to how one might appropriately retheorize the nature of reason in terms of regulatory organization. This brings the initial tasks of philosophical rethinking, and with it the book, to a rounded close. The reader is reminded that, just as we are at the very beginning of a revolution in the foundations of the sciences, so we must be also at the very beginning of a correlative philosophical revolution. This book claims no more than to make a modest contribution to that beginning.

Terminology. Before proceeding it may prove helpful to insert some brief remarks about terminology: *Cognition* is normally understood to refer to the action and faculty of thinking, including perception and conception but excluding all else (e.g., feeling and volition). I don't want to be committed to these divisions; I believe that whole system evaluation (cognitive, moral, and aesthetic) is as essential to the knowing process as to all others. So first I need a term of wider reference encompassing the sum total of ways in which adaptive systems adapt, including marshaling feeling, volition, and evaluation (cognitive, moral, and aesthetic); for this, and despite its sometime usage in a narrowly cognitive sense, I shall use the term *intelligence*. Then I shall understand *cognition* in its broadest descriptive sense to refer to the thinking aspect or dimension of being intelligent, to the action and faculty of thinking, including perception and conception. Intelligence and cognition theories are descriptive, their construction is governed only by the norms for science in general. I shall take cognition to include public thought expressed in written, verbal, and non-linguistic forms (e.g., a laboratory exemplar) and so include science within its scope. Many sciences now contribute to a theory of intelligent systems; *psychology* I take to be the science of the actual intelligent systems historically available to us—living organisms, perhaps later to include robots—

whether or not their behavior is specifically cognitive. Psychology includes the psychology of intelligent behavior, and this in turn includes cognitive psychology.

By contrast with these terms, in its ordinary usage *knowledge* is a "success word," what is known must be true, and it comes weighed down with formal logical associations, the detritus of past attempts to guarantee that what in practice we rely on in science and common sense really is true. Naturalists who use the term have constantly to qualify it to mean just our currently most reasonable, but fallible, assertions—about whatever topic, including knowledge. *Epistemic* action is intelligent action when directed toward knowing and understanding. Its theory, epistemology, includes normative cognitive science (as much cognitive science tacitly is) and its structure as well as its construction is governed by the epistemic norms for knowing and understanding (i.e., by those for science generally). For naturalism epistemology is also a fallible descriptive science (see below).

The basic terms of systems science, such as *system*, *feedback*, *function*, and *stability* are more or less well established and will simply be used herein. Briefly, a system S is stable in parameter P, for example, if the system response to perturbations in the value of P always returns P to its original value and maintains it there. A refined quantitative analysis would introduce the range of perturbations over which stability holds, the limit to within which the system returns to its original value, the time it takes to do so, and the range of transient states it occupies while doing so. But I forego any further analysis here. Other terms, however, remain either vague or ambiguous or controversial, and it is useful to indicate how I use them. By a *regulatory system* I mean a system so equipped as to stabilize those parameters and processes necessary to its continued existence. These constitute respectively its homeostases (e.g., body temperature) and homeorheses (e.g., red blood cell production, immune system antibody adaptation, and cognitive development).[1] Many such systems are highly dynamic in the sense that, rather than maintain a particular form invariant, their interactions are nonlinear and they go on changing in historically idiosyncratic ways. If the system dynamics promotes certain processes (i.e., increases the value of relevant parameters, such as stored environmental information or response strategies) then I shall call the systems *self-organizing*.

Systems are *adaptive* if their interaction with their environment is of a kind that the dynamical sequence of system states is a nontrivial systematic function of the state of the environment. Adaptive processes include both adaptation and adaptability, that is, (respectively) alteration of specific traits so as to increase fitness in their environment and alteration of these alteration processes (chapter 2). The kinds of adaptive processes that chiefly interest us for a study of intelligence are those occurring in self-organizing regulatory systems characterized by homeorheses. In what follows I shall sometimes refer to them as self-organizing complex adaptive systems and often more broadly as self-organizing regulatory systems. Systems of this kind, among others, typically show a temporal development or ontogenesis that is characterized by the selective extraction of ordered energy (negentropy) from their environment to produce simultaneous system expansion and internal differentiation. The growth and development of any living organism is an example. I shall refer to this overall combination of expansion and differentiation as *superfoliation*. A superfoliating adaptive system may show increases in either or both adaptation and adaptability.

Finally, a note on the concepts of regulatory order and system level. These are distinct. A *system level* is a relatively functionally isolatable component of total system function grounded in causal processes relatively causally buffered or screened from causal influences entering from elsewhere, for example, kidney or immune system function vis-à-vis the statistical behavior of their individual cells (see Collier 1988b). Typically, levels can be associated with a particular system scale (e.g., cellular or phenotypic), and then this gives them a quasi-spatial sense, but this is not always the case; for example, the immune system runs across the cellular and organ scales. By contrast *regulatory order* is concerned with the sequence of conditionalization among regulatory outputs. If an output of regulator B is causally determined by its input and an output of regulator A, but not vice versa, then A is a higher order regulator than B. Thus a homeostatic regulative output occurring is conditional on the prior operation of the corresponding homeorhetic regulative output occurring. It is possible to causally instantiate regulatory order so that level structure materially reflects order structure; this is, for example, how a business office is often arranged, with the chain of command literally reflected in the height above ground of offices. But it is not necessary to do this, complex regulatory order may be modeled, for example, by a tuned neural net that is all at the one level, and an effective secretary may function across many levels of an institution. (Cf. the discussion in Hooker 1994e.) The latter may prove to be the commoner situation, and all kinds of intermediate designs are possible. I shall try to respect the distinction between orders and levels, which is constantly blurred by the casual use of *level* to refer indiscriminately to regulatory orders proper, metalinguistic relations, degrees of generality, and so on. But where others use it ubiquitously, I retain this usage for it, trying to confine it to contexts where it does little harm. (This is the policy pursued with Popper in chapter 3.) I retain *level* also for referring to levels of physical quantities and the like, just as I retain the term *order* in "thermodynamic order" and "mathematical order"; these should cause no confusion in context.

1.I. Evolutionary Naturalist Realism: A Philosophical Outline

The guiding spirit of all my philosophical work has been naturalism. Roughly, this means regarding all aspects of life as part of a single natural world, and this includes philosophy itself. This idea is developed in Hooker 1987. It will be discussed only briefly below. Naturalism, pursued systematically, has rather radical consequences (at least by contrast with much contemporary philosophizing), for example, philosophy is to be treated as fallible theorizing, like science; formal logic is abandoned as the paradigm of rationality, and evolutionary epistemology may be treated as only a first crude approximation to a more system(at)ic theory of knowledge. The purpose of this part is to set out these and other consequences of naturalism systematically, for they form the framework for all that follows.

But first a brief illustration of the general orientation of the position. A standard challenge to realism is that the world itself cannot be inspected independently of human conceptually interpreted sense perception, so how can those notions that lie at the heart of realism, correspondence truth, laws of nature, scientific progress, especially in method, and so on be determined or even coherently specified? Surely we have to confine ourselves to our perceptual experience and whatever can be defined from within its resources. The initial premise is granted, but before we leap to the proffered conclusion, let us pause to reflect that science itself understands us very differently. Current scientific theory presents us as one evolved species among many inhabiting a world we did not make and only very imperfectly understand, finding our way about through the use of highly fallible sensory and motor capacities orchestrated by equally highly fallible theories we construct and are constantly forced to reconstruct as our experience extends across ever wider environmental conditions. From this point of view nothing is certain; cognition is as problematic as the world. So those who counsel that we are to assume instead that our sensory experience is our foundation face an embarrassing dilemma: Either science is hopelessly wrong, or their position is internally inconsistent (science constructed on sensory foundations contradicts the existence of those foundations). By contrast, naturalist realism presents a self-consistent account, scientific and philosophic, and this speaks for it. Further, it leads to a proper analysis of the senses as dynamic adaptive systems, which shows how and why the notion of foundational sensory data fails (cf. Hooker 1978; see the similar response to arguments for relativism in Hahlweg/Hooker 1988). Self-reflexively, this framework is itself a fallible theory, but one for which scientific and philosophical analyses of this kind provide heartening support and promise of future fruitful engagement with science. In this way naturalist realism presents a coherent intellectual account of us as fallible, (partially) self-organizing adaptive systems. Now to develop the general principles of this position in a little more detail. The overall argument is summarized in Diagram 1.1.

From a naturalist viewpoint, human knowledge itself is a natural phenomenon, a complex of individual and species capacities with particular origins and related distinctive characters, to be studied as any other natural phenomenon—with the added complication that such theories must be reflexively consistent. It follows that epistemology (indeed, philosophy at large) should form a coherent unity with science, a single self-consistent conception of us and our cosmos.

From this, realism follows. For science presents a well-confirmed conception of an independent external world and of ourselves as a relatively recently evolved species in it, learning through causal interactions with it. This is the basic metaphysics for realism. It is of course possible to consistently doubt and dispute any particular part of the current scientific image, and that image is clearly only partial and then only approximately accurate. Naturalism requires a critical realism, not a naive one. Nonetheless, one who doubts only on reasonable grounds should acknowledge that at the present time there is no serious reason to question the general lineaments of the scientific image.

According to the scientific image, cognition involves the active construction of some kind of map or model in the head, beginning with elementary sensory-motor coordinations and elaborating on these. However the correct account of meaning goes, precisely, the basic referents of these constructions are situations in our external world causally mediated to us. This is the basic metaphysics for the correspondence theory of truth. To put the matter argumentatively: Either one adopts a coherence or a correspondence theory of truth, but the content of science rationally accepted under coherence truth yields correspondence truth metaphysics, so correspondence truth.

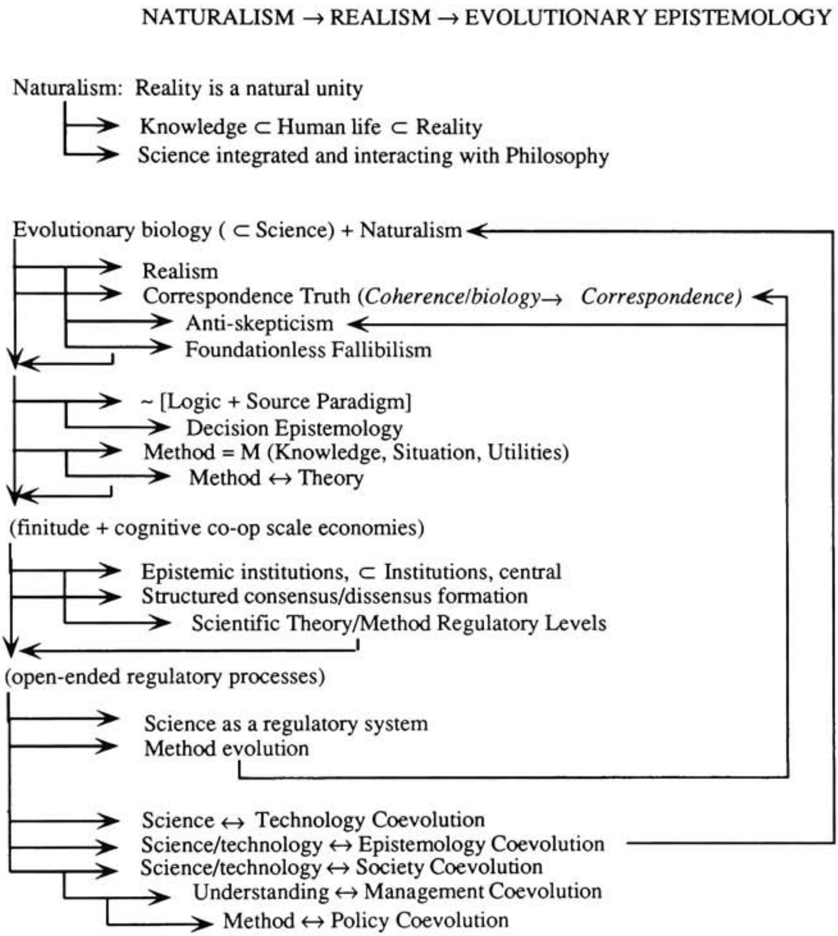

Diagram 1.1 Evolutionary Naturalist Realism
Schematic flow of argument from naturalism to its philosophical consequences. Descending arrows carry the argument forward; ascending arrows represent internal coherence and mutual reinforcement.

Another argument for correspondence truth comes from the open-endedness of the evolving cognitive process. The whole evolutionary process is an open-ended ascent of regulatory control (see below). As an extension of this process, cognition is, and should be, open-ended also because we are born in evolutionary ignorance; at our species birth and at each of our individual births, we arrive lacking explicit knowledge not only of the world, but also of every aspect of what knowing itself is. And as fast as our knowledge has expanded, our knowledge of the extent of our ignorance has grown even faster. But all coherence and pragmatist definitions of truth close off open-endedness. They do this because the truth criteria themselves, whatever they are, cannot in turn be examined for adequacy, cannot be opened up to learning. Coherence and pragmatist definitions of truth are attractive precisely because they offer epistemically accessible truth criteria. But it is just this feature that cuts off their open-endedness. Suppose, for example, one entertains the notion that for all propositions p, p is true $=_{df}$ p is entailed by science S arrived at through idealized use of methodology M. (It will have to be an idealized use in order to cover those parts/aspects of the universe that we would otherwise miss because they were too big, too small, too fast, too slow, too early, too late, etc.) But M is our construct, we can ask whether it is an adequate tool to achieve truth, for creatures like us, in our world. And this is something about which we can learn, and typically have learned, as we gain historical experience with methods and simultaneously improve our knowledge of ourselves and our world. So not only is our cognitive system open-ended, but it ought reasonably to be kept open-ended for creatures like us. But if M is accepted as definitive of truth, it cannot be critically assessed in this manner.[2] The argument generalizes: All definitions of truth employing epistemically accessible truth criteria are inadequate.

Correspondence truth, by contrast, is theorized as faithfulness to a reality that transcends, because it is "external" to, our cognitive or belief systems. Precisely for this reason, correspondence truth transcends the epistemic, so it can act as an ideal for an open-ended cognitive process.[3] One can always insist on epistemically accessible truth conditions, say by arguing that any others are meaningless (cf. empiricists) or confused (cf. idealists); these arguments aren't convincing, displaying instead just the kind of anthropocentrism from which naturalism is the best protection.

For if anything exists not logically or conceptually bound to cognition, then it is an open theoretical issue what exactly exists and how it is best known. Hence for naturalism all is fallible theory, from the lowliest factual claim to the highest metaphilosophical principle. This position is consistently supported by the scientific image. According to the evolutionary viewpoint, our entire cognitive capacity evolved from ignorance, a lack of knowledge not just of our latest theoretical concepts but even of the nature of conceiving itself. Our senses provide epistemically limited, partially biased, and moderately unreliable information channels, and we start in ignorance not only of our world, but also of the nature and reliability of our own perceptual access to it. The point generalizes to every aspect of cognition. These circumstances underwrite systematic fallibilism, applying at all orders of substantive conjecture, from perception through science to philosophy and metaphilosophy. Indeed, they underwrite a systematic foundationless fallibilism. There are no foundations because there are no guaranteed information channels that could act in this capacity. Not perception, certainly, and not reason either, since however a priori it may seem to the individual, it is a posteriori for the species (and that even over historical times).[4]

As remarked, our historical experience supports this self-conception. It has taken us two millennia to construct our present notions of geometry, for example, and these constructions have at every stage of their development arisen through interaction with our stumbling practical attempts to understand our world, from the mensuration of fields in ancient Egypt to the recent difficulties of reconciling classical mechanics and electromagnetism.

Nowhere is this experience of fallible historical development more vivid than in the development of our understanding of reason itself. There have been explosions in both logic and decision theory this century, resulting in a variety of distinct systems of both; contrast classical logic, for example, to intuitionist, n-valued, relevance, super-valuation, fuzzy, and quantum logics. These latter logical systems have all burst upon the scene this century, spoiling the air of necessity created by the preceding 2,500 years of logical theorizing in the Aristotelian mode. Their arrival followed the coupling of logic to algebra and geometry and the proliferation of mathematics consequent upon the discovery of non-Euclidean geometries and other nonstandard mathematical systems. A parallel story could be told for decision theory. The explosion of research in decision theory and logic has also resulted in the uncovering of numerous "paradoxes" (Arrow, Newcombe, Lowenheim-Skolem, etc.) and of startling theorems, such as the radical distinctness of quantum logic. There is increasingly intimate interaction of both these theories not only with mathematics, but also with economics, engineering, physics, and other scientific theories (cf. Hooker 1987). Again, conceptions of scientific method have changed substantially, especially across the last three centuries (Blake et al. 1960; Oldroyd 1989). And there is now intimate interaction between theory, technology, and method (cf. Hooker 1987 and below). All of this provides vivid examples of our learning historically about reason.

What then is a naturalist to say about the general nature and status of norms? Normative claims break up into two components, one functional and one about status. The functional component divides up the subject matter (inference, methodology, etc.) into the normatively permitted and the normatively excluded. Without this function, one does not have a normative claim at all. The status component claims that the normative principle should be obeyed for certain relevant reasons, for example, that it is a truth of reason. It has been characteristic of Western philosophy to try to defend normative claims by giving them some such privileged status. But from the point of view of naturalist realism, all claims have the status of fallible conjectures, there is no "higher" status with which to privilege normative claims. That leaves the question of status open. The status of philosophy as theory, together with our historical experience of their fallibility, supports the conjecture that there is no more to a claim for normative status than a claim of theoretical adequacy for describing some goal-directed activity. In this case, only their function remains as something distinctively normative about a claim.

This normative function is fulfilled by any adequate theory, for a theory is used to critically assess both methods and data within its domain. The primary methods of science are theoretical methods; that is, they are informed by our theories of the nature of the domain under investigation (one does count stable particles but not turbulent vortices, since theory tells us that the quantity of the former but not the latter is conserved, etc.), the nature of the instruments used to investigate that domain, the causal or statistical character of the relationships involved, and so on.[5] Theories have the role of partitioning methods and data into the acceptable and unacceptable. The more adequate the theory—however adequacy is assessed (explanatory power, empirical accuracy, etc.)—the more forceful is its critical, partitioning function. Thus for a naturalist claim to have a functional normative role for some domain is for it to be part of a fallible theory concerning that domain, and the rationally acceptable normative force it is taken to have is proportional to its adequacy.

In the old terminology, there is therefore a symmetry between normative and descriptive role and force. The more descriptively adequate a theory, the more it carries normative force, and the descriptive presuppositions of prescriptions that regularly lead to successful action thereby accrue respect. Conversely, the more serious the failure of descriptive (respectively prescriptive) adequacy, the more serious the weakening of normative (respectively descriptive) force. We simply don't take seriously methodological prescriptions based on seriously inaccurate or ad hoc conceptions of our situation.

Now we can return to the question of status. Although naturalism rules out any privileged status for normative theories, it does not rule out all differences between description and prescription. Naturalism should retain this traditional general skeleton of nor-