

1 Introduction: Massive Modularity and Coming Attractions

Not since the advent of sociobiology has an emerging, interdisciplinary subject attracted as much attention as evolutionary psychology. There has been a flurry of articles and a great deal of lively debate about the status of this exciting new discipline.¹ Since I am a philosophical naturalist, I think it is important for philosophers to learn from, and contribute to, this ongoing discussion.² With that in mind, my goals in this book are of two kinds. First, I want to clarify and evaluate the empirical and conceptual credentials of evolutionary psychology. Second, I want to assess the implications of evolutionary psychology for some issues in epistemology, philosophy of science, and philosophy of mind.

To set the stage for my project, it is first necessary to elucidate the core ideas that constitute evolutionary psychology. This is no simple task, for the following reason. Since this discipline is still in its infancy, there is no consensus among evolutionary psychologists as to the central topics, theories, or methods of the discipline. The best strategy to adopt, given this state of affairs, is to clarify the position of this emerging discipline's best-known advocates, Leda Cosmides and John Tooby. But their "position" is something of a moving target itself in the sense that they say things that appear to be contradictory.³ Despite this appearance, I think that they do share a coherent position. We will need to extract the essence of that position from the sometimes rhetorical prose they employ to extol their fledgling discipline. My strategy will be to clarify and defend this account of evolutionary psychology while, at the same time, assessing the philosophical implications of it. I want to offer, in the process, a new approach to naturalized epistemology. The idea is that we ought to allow epistemology to go modular and view knowledge as a set of natural kinds housed in a massively modular mind. That is, knowledge is not a univocal concept to be clarified by

a priori analysis but an empirically discovered phenomenon, like water, to be elucidated using the results of science and made consistent with other scientific results. My conviction is that if we do this, we can free epistemology from the wheel-spinning scholasticism of conceptual analysis, a strategy that, in my view, is both sterile and moribund. First, however, a brief sketch of evolutionary psychology is needed in order to provide a focus for subsequent philosophical discussion.

1.1 Transtheoretical Consistency

A fundamental methodological assumption of evolutionary psychologists is that a conceptually integrated approach to the behavioral and social sciences is needed. There must be consistency between the results of evolutionary biology, psychology, and disciplines that study culture, such as sociology and anthropology. The idea is not that the logical positivist “unity of science” hypothesis holds in the sense that there would be, in a completed science, theoretical reduction of sociology to psychology, psychology to biology, biology to chemistry, and chemistry to physics.⁴ Rather, the idea is that the results of each discipline must constrain the results of the others, and so conceptual integration and multidisciplinary, multilevel compatibility is needed for transdisciplinary consistency. As Cosmides and Tooby (writing with Barkow) note: “[To] propose a psychological concept that is incompatible with evolutionary biology is as problematic as proposing a chemical reaction that violates the laws of physics” (Barkow, Cosmides, and Tooby 1992, p. 4). Once again, Cosmides and Tooby are not committed to theoretical reduction, since consistency does not imply theoretical reduction. The authors state: “The natural sciences are already mutually consistent: the laws of chemistry are compatible with the laws of physics, even though they are not reducible to them. Similarly, the theory of natural selection cannot, even in principle, be expressed solely in terms of the laws of physics and chemistry, yet it is compatible with those laws” (ibid.). Consistency, itself, is not difficult to achieve in the sense that any two true theories are consistent with each other (though truth can be difficult to determine). But, as Jerry Fodor has recently suggested, what Cosmides and Tooby really want is the stronger notion of mutual explanatory relevance (Fodor 2000, p. 82). For instance, evolutionary biology and cognitive psychology mutually constrain the range of admissible theories that the other can take in.⁵

Unlike Fodor, I do not think that Cosmides and Tooby call evolutionary psychology into being “by methodological fiat” by making this move. Instead, the idea is to employ this methodological strategy and to see what results follow from, for instance, reconfiguring what counts as an interesting experiment from the standpoint of reverse engineering. (See p. 9, this vol., for more on reverse engineering.) In the spirit of free inquiry, it is hard to imagine why we should deny Cosmides and Tooby such freedom even if we acknowledge that their approach is novel and that much of science is not so constrained. Presumably, part of the attraction and the success of their approach can be directly attributed to their willingness to do things differently. Of course, they want to encourage the scientific community to follow suit. Whether we should follow their lead is an important issue, but one that will not be directly dealt with in the following pages. At any rate, I think Fodor is right to point out that what Cosmides and Tooby call “trans-theoretical consistency” is really, at the end of the day, mutual explanatory relevance among the sciences. One might think of this as a postpositivist, nonreductionist, unity of the sciences aspiration.

1.2 Universal Human Nature

Cosmides and Tooby are committed to the idea of a universal human nature at the level of evolved psychological mechanisms, or Darwinian modules, not of expressed cultural behaviors. Second, these evolved psychological mechanisms are adaptations that were constructed by natural selection over evolutionary time. Hence, Darwinian modules are innate cognitive structures whose main properties are largely determined by genetic factors. A third idea is that the evolved structure of the human mind is adapted to the way of life of Pleistocene hunter-gatherers, not necessarily to our modern circumstances. The idea is that what we think of as recent human history, that is, the last two thousand years, does not have much to do with the shaping of the human mind. This is because our minds were largely shaped by the last two million years as Pleistocene hunter-gatherers. As Cosmides and Tooby note:

Complex, functionally integrated designs like the vertebrate eye are built up slowly, change by change, subject to the constraint that each new design feature must solve a problem that affects reproduction better than the previous design. The few thousand years since the scattered appearance of agriculture is only a small stretch in evolutionary terms, less than 1% of the two million years our ancestors spent as Pleistocene hunter-gatherers. For this reason, it is unlikely that new complex designs—

ones requiring the coordinated assembly of many novel, functionally integrated features—could evolve in so few generations. (Barkow, Cosmides, and Tooby 1992, p. 5) Even a staunch defender of the punctuated equilibrium view of evolutionary change, such as Stephen J. Gould, would not have quibbled with Cosmides and Tooby on this point.⁶

1.3 Domain Specificity

Darwinian modules are domain specific. According to Cosmides and Tooby, our minds consist mostly in “a constellation of specialized mechanisms that have domain-specific procedures, operate over domain-specific representations, or both” (Cosmides and Tooby 1994, p. 94). What this means is that a particular innate cognitive structure will respond to only a particular kind of representational input. As such, the existence of Darwinian modules seems, in principle, to limit frame and relevance problems. The “frame problem” is the question of how one can reconcile a local notion of computation with what seems to be the holism of rational inference. To wit, abductive inference seems to be able to draw on the entire corpus of one’s prior epistemic commitments. As Fodor points out: “Frame problems and relevance problems are about how deeply, in the course of cognitive processing, a mind should examine its background of epistemic commitments. Modular problem solving doesn’t have to worry about that sort of thing because its searches are constrained architecturally; what is in its data base can count as in the frame, and nothing else counts as relevant” (2000, pp. 63–64). The domain-specificity of Darwinian modules severely limits, ex hypothesi, frame and relevance problems.

1.4 Computational Mechanisms

Darwinian modules are also computational mechanisms. As Cosmides and Tooby put it: “our architecture resembles a confederation of hundreds or thousands of functionally dedicated computers (often called modules)” (Tooby and Cosmides 1995, p. xiii). And again: “The brain must be composed of a large collection of circuits, with different circuits specialized for solving different problems. One can think of each specialized circuit as a minicomputer that is dedicated to solving one problem. Such dedicated minicomputers are sometimes called modules” (Cosmides and Tooby 1997b, p. 81). The idea that underlies this conception is due to Turing. It was Turing who

first introduced the idea that mental processes are computations. Such computational devices are classical computers. Hence, cognitive mental processes are formal operations defined on syntactically structured mental representations that are similar to sentences. A computation is a causal process that is syntactically driven. Cosmides and Tooby are not explicit about computations being classical computers, however, and so it is not clear that they take all modules to be formal operations defined on syntactically structured mental operations that are sentence-like. But, surely, many modules must be such classical computers.

1.5 Poverty of the Stimulus Arguments

Finally, some of the motivation for thinking that Darwinian modules exist comes from Chomsky's idea that poverty of the stimulus arguments determine the information a mind must have innately.⁷ One must subtract the information that is in the environment from the information required for a child to attain linguistic mastery. What is left over is what the child's innate knowledge contributes to the language acquisition procedure. For evolutionary psychologists, what is left over that is innate is enormous: a mind that is largely, but not completely, constituted by hundreds or thousands of functionally dedicated computers. Darwinian modules, in sum, are innate, naturally selected, domain-specific, Turing computational mechanisms that often work alongside domain-specific bodies of data or representations. Call this the *massive modularity hypothesis*.

1.6 The Massively Modular Representation and Processor Model of Cognition (MMRP)

Cosmides and Tooby are uncommitted on the issue of connecting domain-specific computational processors with domain-specific bodies of information. As they note:

[In] reading the literature on domain-specific reasoning in children, one could come away with the impression that the study of cognition is nothing more than the study of representations. But representations are, by themselves, inert. Obviously, there must be procedures that operate on representations if the brain is to process information. So the next step for many researchers lies in discovering where the domain specificity lies—in the child's mental representations, in the procedures that operate on these representations, or in both. (Cosmides and Tooby 1994, p. 105)

Since there must be procedures that operate on representations for the brain to process information, and, indeed, if the brain were to avoid being inert, it would seem to follow that the two go together. The way that computational processors and bodies of data might go together in a massively modular mind is that both could be domain specific. As such, it might seem that it is an a priori truth that, if there are innate, domain-specific computational processors or Darwinian modules, then there must be innate, domain-specific bodies of data, or Chomsky modules, that the processors are tied to and operate on. Every domain-specific module, therefore, would be a Darwinian/Chomsky module. It follows straightaway that those, like Spelke, who argue that the infant's object concept is embodied in procedures that are domain specific but amodal in that they operate on both visual and tactile data, would be wrong to think that such data bases are domain general (Spelke 1988, 1990). At best, Spelke might say that the child's object concept is bimodal, though still domain specific. Samuels (1999) and Samuels, Stich, and Tremoulet (1998) would be, likewise, mistaken to think that one can be committed to innate, domain-specific, computational processors but not innate, domain-specific bodies of data.⁸ Despite the intuition that "If there are domain-specific modules at all, then they must be Darwinian/Chomsky Modules (DCM)," Cosmides and Tooby do not take this position. That is, they think there are DCMs but that there may well be other aspects to our cognitive architecture too.

Cosmides and Tooby, in fact, do not commit themselves to the notion that the mind must be entirely modular. To a first approximation, one might think about the possibilities as being fourfold:

	Domain-specific processors (Darwinian modules)	Domain-general processors
Domain-specific bodies of data (Chomsky modules)	A	B
Domain-general bodies of data	C	D

Box A, for instance, represents the conjunction of domain-specific processors, or Darwinian modules, with domain-specific bodies of data, or Chomsky modules.

Cosmides and Tooby suggest that boxes A, B, C, and D are all live options and that these four options are not mutually exclusive (Cosmides and Tooby 1994, p. 104). Here is what they say about these four options: "Any of these possibilities may be correct. Indeed, all may be correct, although for

different domains” (ibid.). Beyond proving that they do not blithely accept A as true (though they certainly think A is more likely to be true than B, C, or D), this proves that they do not think that only A need be true. They leave the door open to several of these options being true, though for different domains. For Cosmides and Tooby, the mind need not be only modular (or only massively modular). The mind, on their view, can also be nonmodular in certain respects.⁹

This claim is important since some authors, such as Fodor and Samuels, take them to task for defending the view that the mind is composed entirely of Darwinian modules.¹⁰ In my view, the correct interpretation of Cosmides and Tooby’s standpoint is that they are committed to the notion that the mind is largely composed of a vast array of Darwinian/Chomsky modules. In addition, some instances of boxes B, C, or D exist. This is fortunate, since I think that there may well be some aspects of mental architecture that are nonmodular because there is a place where, to some degree, as Fodor says, “it all comes together.” If that were not possible, abductive or global inductive inference would not exist. But abductive inference, within reasonable bounds, appears to exist. It would seem to follow that the mind cannot be completely modular. (But see chapter 2 for more on this point.) We can now state the position that Cosmides and Tooby accept. The massively modular representation and processor model of cognition states that:

MMRP Principle: The mind is largely composed of a vast array of Darwinian/Chomsky modules. Caveat: There must be some box B, box C, or box D components to the mind.

Neither Cosmides and Tooby nor I are committed to the idea of an entirely modular mind. Their position, one might say, is fairly liberal concerning what might be the case. But this fact is not so much a result of a “California” attitude as it is a result of the current state of empirical play. In short, the data simply do not warrant ruling out any possibilities. At the same time, it must be emphasized that Cosmides and Tooby are certainly committed to the view that Darwinian/Chomsky modules dominate mental architecture.

1.7 Evolutionary Psychology and Human Reasoning

A significant amount of Cosmides and Tooby’s work in evolutionary psychology has been devoted to the study of human reasoning. In particular, the content effects on the Wason selection task have been a benchmark for

those seeking an adequate account of human reasoning in the sense that any acceptable account must explain these effects. Studies by Peter Wason (1968) and Wason and Philip Johnson-Laird (1972) demonstrated that reasoning performance on distinct tasks that require the use of a single rule of deductive inference varied as a function of the content plugged into the inference rule. This violates the most fundamental idea of formal logic, namely, that arguments are valid purely as a function of their abstract form regardless of their content. That humans consistently fail to observe the content-neutrality aspect on deductive reasoning tasks came as an enormous surprise. For instance, consider the following Wason selection task (Johnson-Laird 1983, p. 30). An experimenter lays out four cards in front of a subject with the following symbols:

E K 4 7

The subject knows that each card has a number on one side and a letter on the other side. The experimenter now presents the following generalization to the subject:

If a card has a vowel on one side then it has an even number on the other side.

The subject's task is to turn over only those cards that need to be checked to see whether the generalization is true or false. The order of card turning is not at issue. This seemingly simple task turns out to be very hard to solve. Most everyone sees that the card with the vowel needs to be turned over. The generalization is left untouched if this card, once turned over, is even. If the card is odd, then the generalization must be false. Similarly, most subjects realize that the card bearing the consonant need not be touched since the rule says nothing about consonants. Some subjects turn the card bearing the even number; some do not. But the even card need not be turned over since whether there is a vowel or not on the other side, the conditional will be true.

The central problem concerns the card with the odd number. This card must be turned over since, if it contains a vowel, then the generalization is false. But few subjects insist on turning it over even though the reason for checking it is exactly the same as the reason for checking the card with the vowel: the generalization is clearly false where a vowel and an odd number occur on two sides of one card. Wason and Johnson-Laird tried many changes of procedure and materials in an attempt to improve performance

until they found a simple alteration that had a striking effect (Johnson-Laird 1983, p. 31). When subjects were presented with four cards representing journeys, that is, with a destination on one side and a mode of transportation on the other side, the results were much better. So, with the cards

Manchester Sheffield Train Car

and the general rule

Every time I go to Manchester I travel by train

over 60 percent of subjects understood that they should turn over the card with “car” on it. This stands in sharp contrast to the previous example where almost no one understood, in the parallel case, that the odd-numbered card must be turned over. And, in the control group, just over 12 percent of the subjects made the equivalent choice where abstract materials were involved. This suggests that realistic or familiar materials produce much better results than abstract or unfamiliar materials, regardless of the fact that distinct experiments employed generalizations with the same logical form and truth conditions. Many other variations on these experiments were performed that seemed to confirm these results.

But it was Cosmides and Tooby’s groundbreaking work explaining the content effects on the Wason selection task that, in part, set the stage for the more general project of evolutionary psychology. Cosmides and Tooby explained content effects by appeal to the presence or absence of a “social contract” in the selection task, rather than the Wason/Johnson-Laird associationist notion of “familiarity.” Interestingly, the discussion of evolution and computation in Cosmides’s award-winning 1989 article is restricted to a footnote.¹¹ But the idea of an innate, massively modular mind was clearly in the background. The strategy of “reverse engineering” guided the methodology of these studies. “Reverse engineering” refers to an experimental design strategy where one must attempt to determine the adaptive problems that Pleistocene hunter-gatherers faced, and then design experiments that would make perspicuous the functional adaptations that arose in response to those adaptive problems.¹² One needs to understand the design features caused by natural selection in order to clarify such complex functional adaptations. As Cosmides and Tooby note concerning the role of chance in evolution: “Random walks do not systematically build intricate and improbably functional arrangements such as the visual system, the language

faculty, or motor control. The only known explanation for the existence of complex functional design in organic systems is natural selection" (Cosmides and Tooby 1994, p. 86). This point also holds true for the human reasoning capacity. As such, the reverse-engineering methodology guided Cosmides and Tooby's studies of the Wason selection task.

For philosophers, perhaps the most stunning result of empirical psychology over the past thirty years has been the deconstruction of the notion that there is an underlying Russellian psycho-logic that guides thought and language. By rigorous logical standards, humans appear to be wildly irrational in the sense that there are no truth-preserving, content-neutral, domain-general, logical systems that humans employ in everyday reasoning. We make inductive and deductive errors that are widespread and pervasive. The experimental results supporting these claims are robust and replicable. We commit the conjunction fallacy and are guilty of base-rate neglect and overconfidence.¹³ In chapters 4 and 5, I review some of the empirical data concerning these studies. As earlier noted, Cosmides and Tooby have tried to overturn the interpretations of the associationism-based availability theorists, such as Wason, Kahneman and Tversky, Nisbett and Ross, and others, concerning the content effects on the Wason selection task. Associationists argue that familiarity with the data (or differential experience) explains the fact that subjects will reason in accordance with, for instance, *modus tollens* in some cases but violate *modus tollens* in other cases. The supposition by associationists was that humans do possess a domain-general, content-independent reasoning capacity. Later, pragmatic reasoning theorists, such as Cheng and Holyoak, argued for a similar domain-general reasoning capacity, but they suggested that humans employ "pragmatic reasoning schemas" that were induced through recurrent experience within goal-defined domains (Cheng and Holyoak 1985, 1989; Cheng et al. 1986). The schemas were thought to be content dependent, whereas the inductive cognitive processes that gave rise to the schemas were thought to be content independent. Differential experience explains which schemas were built and why other schemas were not built. But Cosmides tried to show that neither the associationism-based availability theory nor the induction-based pragmatic reasoning theory could explain the content effects on the Wason selection task. She argued that the presence of a social contract embedded in the Wason selection task explains the content effects. If there is a social contract involved in the task, even one that is not familiar to the subjects, then

their results will be enhanced. Cosmides's prediction was confirmed by her experiments. Clearly, these empirical results challenged the truth of the availability theory.

Cheng and Holyoak argued that pragmatic reasoning schemas could explain the content effects on the Wason selection task in the sense that where a schema was ingredient, subjects would reason in accordance with the propositional calculus. Conversely, where no schema was evidenced, subjects would fail to reason in accordance with the propositional calculus. Cosmides noted that all social contract rules involve permission rules (or schemas), but not all permission rules involve a social contract. This is because the social contract statement, "If one is to take the benefit, then one must pay the cost," entails the permission rule, "If one is to take action *A*, then one must satisfy precondition *P*." But the reverse does not hold. All benefits taken are actions taken, but not all actions taken are benefits taken. As Cosmides noted: "A permission rule is also a social contract rule only when the subjects interpret the 'action to be taken' as a rationed benefit, and the 'precondition to be satisfied' as a cost requirement" (Cosmides 1989, p. 237). This makes the domain of permission schemas larger than that of social contract algorithms.

For instance, in one experiment Cosmides tested the following permission schema that was not a social contract: "If a student is to be assigned to Milton High School, then that student must live in the town of Milton." The surrounding story for this non-social contract permission problem gave the rule a social purpose: following the rule will allow the Board of Education to develop the statistics necessary to assign teachers to each school. But notice that no cost-benefit structure is built into this permission schema from the subject's standpoint. According to Cheng and Holyoak, however, permission schemas that are not social contracts should still result in content effects on Wason selection tasks, contra Cosmides's social contract theory. But this turns out not to be the case: there appear to be no content effects in such cases. Cosmides was right. Though this evidence did not constitute conclusive evidence against the pragmatic reasoning schema approach, it certainly was news.

Later, Gigerenzer and Hug (1992) duplicated Cosmides's results, but they offered a friendly amendment to those results. They agreed that the associationism-based familiarity and the pragmatic reasoning schema hypotheses were false (or likely false). But they demonstrated that, where

social contracts that involved a cheater detection algorithm were involved, content effects took place. In contrast, where only a social contract was involved and there was no cheater detection algorithm, no content effects were evidenced. More recently, Cosmides and Tooby have argued that humans may be much better intuitive inductive reasoners than they are intuitive deductive reasoners. A series of studies by Cosmides and Tooby involving base-rate neglect has shown that if subjects are presented problems as involving relative frequencies, rather than single-case probabilities, the results are dramatically better. Likewise, other researchers have obtained parallel results concerning relative frequencies as applied to base-rate neglect, the conjunction fallacy, and the overconfidence bias. This evidence paints a much more positive picture of human inductive inferential capabilities than that presented by associationists and pragmatic reasoning theorists in the 1970s and 1980s, or even Cosmides and Tooby in the 1980s. However, the status of deductive reasoning remains fairly bleak by strict logical standards.

The empirical story that I have just outlined is essentially that of Cosmides and Tooby. But, as I make clear in chapter 5, there are several alternatives to their analysis of the content effects on the Wason selection task. Some of these studies are examples of evolutionary psychology, some are not. Whether or not one agrees with the methodology and results of evolutionary psychology, this much is uncontroversial: the attempt to link evolutionary theory with empirical psychology, by appeal to the methodological strategy of reverse engineering, has created a booming research industry. It is high time, then, to ask what the philosophical implications of evolutionary psychology might be.

1.8 The Philosophical Implications of MMRP

With this brief survey of Cosmides and Tooby's position concerning that part of evolutionary psychology that deals with human reasoning in hand, it is now time to map out, in rough outline, what I take to be the philosophical implications of these results. I want also to note that evolutionary psychology is a very broad research program with many topics and many viewpoints on these topics. For instance, evolutionary psychologists have written about mating strategies, mate assessment and choice, and other topics. And, Gigerenzer and Hug, Cummins, Manktelow and Over, and others

have provided unique perspectives concerning, and friendly amendments to, Cosmides and Tooby's work on evolutionary psychology. In this text, I focus primarily on Cosmides and Tooby's position concerning human reasoning because that is the literature that is the most germane to the philosophical issues I want to discuss. It should also be noted that there are other accounts of modularity that differ in dramatic ways from that proposed by Cosmides and Tooby. The prime example, of course, is Jerry Fodor's account in *Modularity of Mind*. Fodor posits a nonmodular central reasoning processor, assisted by peripheral input, for example, perception, and output systems, for example, action, that are modular.

But the central rival to an account, such as Cosmides and Tooby's, that sees the mind as involving computations is connectionism. To the extent that Cosmides and Tooby see mental processes as formal operations defined on syntactically structured mental representations, their view will fly in the face of connectionism. As William Bechtel notes: "The connectionist view of computation is quite different. It focuses on causal processes by which units excite and inhibit each other and does not provide either for stored symbols or rules that govern their manipulations" (1991, p. 2). Connectionists, such as Andy Clark, deny that there is any innate representational base: "The point, however, is that the initial weights (assuming a random starting point) are not usefully seen as a set of representational elements (ask yourself what such weights represent?) and, a fortiori, the subsequent learning of the network is not usefully understood as constrained by the representational limitations of an initial 'language'" (1993, p. 36). At the same time, some connectionists do not buy into a tabula rasa model of knowledge acquisition. As Clark says: "Such a model would be implausible on well-documented empirical grounds. . . . The precise way in which knowledge about (e.g.) physics, faces, and language may be built in remains an open question, but one obvious option is for evolution to preset some or all of the weights so as to embody some initial domain knowledge" (ibid., p. 37). Subsequently, learning takes over and develops such innate knowledge. Clark suggests that individuals might acquire modular knowledge through such neural networks. On the other hand, Elman et al. do argue for a tabula rasa theory by suggesting that there are no innate representations and no innate knowledge (Elman et al. 1996, pp. 359–366). However, Elman et al. also argue that there are innate mechanisms for learning and information processing. Suffice it to say that Cosmides and Tooby's account of

the mind is far from being the only game in town. That said, I will briefly summarize coming attractions.

In chapter 2, I take up foundational worries that Jerry Fodor has recently voiced against the massive modularity project. In a book entitled *The Mind Doesn't Work That Way*, Fodor argues that the "New Synthesis," that is, the evolutionary psychology of Cosmides and Tooby, Pinker, and Plotkin, is deeply flawed. Fodor thinks that what he calls the "input problem," an a priori problem, stops the New Synthesis in its tracks. He also offers several other objections to the New Synthesis. In particular, he argues that evolutionary psychology cannot accommodate global, abductive inference within the context of local, computational processors. I argue that he misinterprets the position of Cosmides and Tooby and so commits the straw man fallacy repeatedly. Moreover, I argue that the massive modularity hypothesis can be extended to explain what I call "nonexplicit, evolutionary abduction." Evaluating Fodor's critique is important because it allows us to address concerns that many cognitive scientists may have about evolutionary psychology and because it will allow a more detailed understanding of Cosmides and Tooby's understanding of their own project. In particular, and despite repeated attempts on their part to deny it (both in print and at conferences), there is a rumor going around that Cosmides and Tooby are committed to the notion of a completely modular mind. I hope to dispel this completely erroneous idea. Evolutionary psychology may ultimately provide a false picture of the mind, but a demonstration of its failings can occur only if we begin with a charitable and accurate understanding of evolutionary psychology's methodology, central claims, and the evidence for those claims.

In chapter 3, I canvas accounts of misrepresentation for a solution to the disjunction problem. The crude causal theory of representation has it that tokenings of "D" are reliably caused by D. This makes it the case that the condition governing what it means for D to be represented by "D" is identical to the condition for such a token being true. As such, it is not possible to get falsity into the picture. One might think that D-caused "D" tokenings are true and E-caused "D" tokenings are false. But this will not work. "D"s are reliably caused by the disjunctive property of being (D or E). As such, E-caused "D" tokenings are true because they are reliably caused by (D or E), and we have no theory of misrepresentation. That is the disjunction problem. I argue that it is possible to solve the disjunction problem using the

resources of the massive modularity hypothesis of Cosmides and Tooby, combined with an etiological, reliable proper function account of the computational processes that constitute such modules. The result is an account of misrepresentation or error that is fully applicable to that part of the mind that is modular.

Roughly put, if the gap between the proper domain (within which the module was selected for) and the actual domain (in which we now exist) is large enough, then the proper function of a module will fail. Misrepresentation occurs when the proper domain/actual domain gap is significant and results in the malfunctioning of modules. Moreover, false types of beliefs were endemic to our forebears if they resulted in the misidentification of significant adaptive problems. Misrepresentation is part of the process that leads to the malfunctioning of modules. False beliefs about adaptive problems trigger inappropriate responses by modules. Misrepresentation triggers malfunction, and malfunction results in biological error. Conversely, true types of beliefs were important for identifying adaptive problems for our forebears. The account does not inflate the representational abilities of modules, while it succeeds in biologizing Fodor's asymmetric dependence account of misrepresentation. I dispense with the metaphysically otiose notion of a representation-consumer while preserving the core advantages of Millikan's proper function account. The account shows that meaning matters because truth and falsity matter.

In chapter 4, I argue that accurate indexical representations have been crucial for the survival and reproduction of *Homo sapiens*. Specifically, I argue that reliable processes have been selected for because of their indirect, but close, connection to true belief during the Pleistocene hunter-gatherer period of our ancestral history. True beliefs are not heritable; reliable processes are heritable. Those reliable processes connected with reasoning take the form of Darwinian algorithms: a plethora of specialized, domain-specific inference rules designed to solve specific, recurrent, adaptive problems in social exchange contexts. Humans, I argue, reason not logically, but adaptively.

In chapter 5, I first note that reliabilist and externalist conceptions of epistemic justification and knowledge face criticism from two groups: analytic epistemologists and philosophers of science. The first group criticizes naturalized epistemologists for "changing the subject" and so failing to address the long-standing issues of traditional epistemology and other important

issues. The second group, ironically, criticizes naturalized epistemologists for “failing to change the subject” insofar as their work is connected to traditional, analytic epistemology, an allegedly outdated and utopian form of inquiry. Philosophers of science think that traditional epistemology is utopian because it is an attempt to respond to the unreasonably high standards that the skeptic wishes to impose on any adequate account of knowledge. According to philosophers of science, the secret is to reject those unreasonably high skeptical standards.

Clearly, both groups of critics cannot be right. I mount a case in favor of the naturalization project by first drawing a distinction between meliorative and nonmeliorative senses of justification. I use this distinction to argue that both groups of critics have missed the point of the naturalization project but for different reasons. Later, I review some of the literature on human rationality from empirical psychology and suggest how a naturalized epistemology, which takes seriously these empirical results and is informed by the meliorative–nonmeliorative distinction, might be developed. Along the way, I try to reconnect analytic epistemology with philosophy of science by showing how naturalization projects in these two areas are related. It seems to me that the time is long since due that philosophers recognize that the epistemology of the standard cognizer is continuous with the epistemology of the scientist. Quine, in my view, was right: science is just sophisticated common sense. The two roads to analytic philosophy, symbolized by Moorean informal analysis, on the one hand, and Russellian formal analysis, on the other hand, need finally to be joined once again. The way to do this, I think, is to show how a naturalized epistemology can inform, and be informed by, a naturalized philosophy of science. The two solitudes must merge.

Chapter 6 represents a departure from the previous two chapters in that its discussion centers on the massive modularity literature directly and not on rationality theory. It begins with a discussion of the objection by non-naturalists that science has no place in matters epistemic. Nonnaturalists, such as Richard Feldman and Richard Fumerton, have argued that epistemic issues are normative or evaluative in nature. Since this is the case, it must be irrelevant to employ the factive or descriptive resources of science to reconstruct normative, epistemic notions such as justification and knowledge. I argue that nonnaturalists are simply mistaken about this since knowledge is

actually a set of natural kinds that should be studied empirically, not a conceptual kind to be studied by a priori appeal to reason. I offer the “Knowledge Is a Natural Kind” argument (NKA) to support my view. In making this move, I not only depart from a central tenet of analytic epistemology, I depart from the form of analysis of justification and knowledge adopted in chapters 4 and 5. I would argue, however, that even on the terrain made possible by conceptual analysis, my views are superior to internalist conceptions of knowledge and other externalist conceptions of knowledge. But the virtue of seeing knowledge as a set of natural kinds is that the brief for a fully naturalized epistemology is now, for the first time in the history of epistemology, finally on offer. In my view, this is a novel and important methodological alteration that clears the ground for my modular account of knowledge.

In sum, my account of knowledge depends on the twin ideas that knowledge is a set of natural kinds (and not a conceptual kind) and that such knowledge is housed in a vast array of proper subsets of MMRP modules. Empirical knowledge is the result of Darwinian modules transducing sensory inputs, whereas a priori knowledge is the result of Darwinian/Chomsky modules. Such modular knowledge is instrumentally important since, once properly tethered to our desires, it results in the satisfaction of our biological needs and other goals we might have. The result is that knowledge is composed of a set of natural kinds housed in distinct modules of the mind.

1.9 Conclusion

Steve Stich argued in favor of epistemic relativism in his classic, *The Fragmentation of Reason* (1990). What I hope to show is that there is good reason to think that the more recent literature on rationality and massive modularity present us with grounds for a different picture of reason and representation. The new picture is, I believe, decidedly more optimistic than the one Stich presented. My modular conception of knowledge is a species of foundationalism, though one that conceives of knowledge not as a conceptual kind but as a set of natural kinds. As such, my account of knowledge and justification is incompatible with epistemic relativism even as it argues for the fragmentation of knowledge. Another way that I depart from Stich’s relativism is that I think there are evolutionary grounds for thinking that

truth did, and does, contribute to our fitness as a species. Truth and error are two sides of the representation and rationality coins that I try to reconstruct in the pages to follow, and truth connects these two key notions. In chapter 2, I look at Jerry Fodor's recent book, *The Mind Doesn't Work That Way* (2000), and his interesting criticisms of Cosmides and Tooby's New Synthesis. I will argue that Fodor simply misrepresents their position. These are the coming attractions, and now it is time for our show.