Descartes' Challenge: Flexibility and the Architecture of Cognition

1. The Problem

This book concerns the connection between two widespread and fundamental claims about the nature of human cognition. The first, arguably the single most pervasive assumption in all of cognitive science, is that the human mind is a *mechanism* of some sort: a physical device composed of functionally specifiable subsystems. On this view, a central project of psychology is to characterize the mechanisms from which the mind is composed: to provide an account of our *cognitive architecture*. The second claim concerns the open-endedness or *flexibility* of human behavior and cognition. To a first approximation:

> *The Flexibility Thesis:* Human beings are capable of performing a huge –perhaps even unbounded— range of tasks in a *context appropriate* fashion.

Taken in isolation these claims both seem very plausible; yet it has proven extraordinarily hard to reconcile the two. If both are correct, then presumably some conglomeration of mechanisms should be capable of manifesting the sort of open-endedness characteristic of human cognition. But the task of addressing this *Problem of Flexibility* –of describing a cognitive architecture that accommodates our cognitive plasticity— has met with little success and remains among the most pressing general problems confronting cognitive science today. The central aims of this book are to understand the nature of this problem and to assess the prospects of providing a satisfactory solution. In doing so, I defend a conception of our mental architecture which assigns a central role to general-purpose, classical computational mechanisms for both theoretical and practical reasoning (or planning). This general conception of 'central' cognition played a pivotal role in the development of much early cognitive science, and for this reason I refer to it as the *Standard Model*. Yet over the past few decades it has been the subject of intense criticism and is now widely regarded by philosophers and cognitive scientists as seriously flawed. In the current intellectual climate, then, my suggestion that the Standard Model gets things right is in fact quite a radical position to adopt and, moreover, one that is desperately in need of defense.

The Problem of Flexibility is not, of course, a new one. On the contrary, Descartes provided one of its most insightful formulations almost four centuries ago and in doing so sought to show that cognitive flexibility poses an insurmountable challenge for mechanistic theories of human cognition. Descartes' Challenge can be framed in terms of a tension between the Flexibility Thesis and the following claim about the limits of machines:

> *Mechanistic Inflexibility Thesis* Mechanisms are, by their very nature, specialized and inflexible: each individual mechanism is only capable of satisfactorily performing a very small range of tasks in a context appropriate fashion.

The problem is that, if true, this appears to leave little room for mechanistic theories of mind. On the one hand, the Mechanistic Inflexibility Thesis seems to demand that we posit a specific device —an "organ", as Descartes would say— for each small set of tasks that we can address. But on the other hand, the Flexibility Thesis renders such an approach untenable. For if our cognitive repertoire really is as open-ended as it appears to be, then there will simply be too many tasks to permit an appropriate correspondence between tasks and mechanisms. How, then, can we hope to explain the flexibility of behavior and cognition?

2. Options

Descartes' own response to the Challenge was to maintain that flexibility cannot be explained in mechanistic terms and instead depends on the existence of an immaterial faculty of reason: a "universal instrument" that furnishes us with an unbounded capacity for context appropriate behavior. But such a response is unavailable to the mental mechanist; and (aside from implausibly denying the flexibility of cognition) this leaves only two general options:

1. *The universal machine strategy* Reject the Mechanistic Inflexibility Thesis and maintain instead that mechanisms need not be specialized or inflexible: that there could be a mechanical analogue of Descartes' universal instrument.
2. *The collaborative strategy* Accept both the Flexibility Thesis and the Mechanistic Inflexibility Thesis but maintain that human cognition is a product of the

collective –or better *collaborative*— activity of lots of relatively inflexible devices.

Descartes himself provided little or no serious argument against either strategy, and both have been explored at length by cognitive scientists: the majority assuming some version of the universal machine strategy while some –most notably advocates of massive modularity— adopting the collaborative alternative. Moreover, each strategy has been elaborated in a variety of ways so that there now exist a wide range of different research programs, each claiming to explain –or at least to offer the prospect of explaining— the distinctive flexibility of human cognition. Yet, to date, these programs have met with very limited success, and the challenge that Descartes posed so long ago remains unanswered. The problem is not merely that we have failed to generate any *detailed* mechanistic account of cognitive flexibility –surely an unremarkable state of affairs given the complexity of the phenomenon. Rather the real difficulty with addressing the Problem of Flexibility is that it's far from clear that we even know what *general* approach to pursue or kind of architecture to adopt.

Faced with such an intractable general problem, one eminently sensible research strategy is to go and think about something else. Unfortunately, this is hard to do in good conscience when the trajectory of empirical research becomes distorted by the adoption of implausible and unsupported responses to such general problems. In my view, cognitive science currently finds itself it just such a predicament. In particular, there is an increasing tendency to respond to substantive research problems by faddishly jettisoning older accounts of cognitive architecture in favor of newer ones, even though the arguments are unpersuasive and there is little or no reason to suppose that the newer positions will fair any better. Not that there is anything wrong with trying out new ideas; but it does become a problem when good ones are mislaid in the process.

In this spirit, a central aim of the book is to defend one of the earliest architectural proposals to emerge from cognitive science; one widely associated with the seminal work of Newell and Simon and which, for want of a better name, I call the Standard Model (Newell & Simon, 1972 & 1976). At the heart of this view is the contention that many human behaviors depend crucially on means-ends reasoning (or planning) –a process which, in turn, causally depends on theoretical reasoning for suitably rich and accurate

representations of the world— and that such reasoning processes are underwritten by one or more general-purpose, classical, computational mechanism: in effect, computational analogues of Descartes' universal instrument. In one form or another, this general conception of central cognition dominated much early research in cognitive science; and though it continues to be actively explored, it is now viewed by many as a thoroughly discredited idea. In its stead, philosophers and scientists have advocated a remarkable range of alternative proposals –often heralded as 'paradigm shifts'— which purport to resolve the problems that allegedly undermine the Standard Model. I argue, however, that this rejection is a mistake. First, while the Standard Model raises lots of difficult research issues which newer approaches may help address, its core commitments have *not* been refuted (or even seriously challenged) on either empirical or conceptual grounds. Second, not only are the objections unsatisfactory, but there is a strong case to be made for the claim that the Standard Model captures a mandatory component of any mechanistic account that can handle the undeniable fact of human cognitive flexibility. If I am right, then much of contemporary cognitive science is embarked on a fool's errand: a search for a new paradigm that is neither required nor likely to be found.

3. Aims

This book has three central aims:

a) To clarify the general challenge that cognitive flexibility poses for mechanistic theories of cognition
b) To assess the available responses that have emerged from cognitive science and, in doing so, to argue for the superiority of the Standard Model
c) To address the diagnostic issue of why, despite the availability of this approach, efforts to provide any *detailed* explanation of cognitive flexibility have met with such limited success.

Let me say a bit more about each of these aims.

3.1 Clarifying the challenge: In pursuing the first aim, I provide answers to a range of broadly conceptual issues. First, I clarify the notion of cognitive flexibility and provide a taxonomy of some core aspects of cognitive flexibility that are exhibited by human

beings. Second, I address some core issues about the individuation of cognitive mechanisms which, though central to understanding debates over cognitive architecture, are very frequently ignored. Finally, I clarify a variety of important but vexed notions –including, modularity, innateness and domain-specificity— and explain their relevance to understanding Descartes' challenge.

3.2. Assessing responses to the challenge: Although a broad array of responses to Descartes' Challenge has been proposed in recent decades, the position I defend throughout is that a satisfactory proposal will probably need to conform to the Standard Model in positing general-purpose, classical computational systems for practical and theoretical reasoning. To be sure, this proposal is not without its problems, and I know of no decisive, "knock-down" arguments in support of it. Nonetheless, I maintain that the Standard Model should be adopted on broadly empirical grounds. Specifically, I argue that it constitutes the best available approach to understanding cognitive flexibility and, moreover, one whose core commitments we have no good reason to reject. In contrast, the main (putative) alternatives to the Standard Model neither possess its virtues as an account of flexibility nor --contrary to what many maintain-- address its central problems.

Of course, the real work of the project consists in a detailed and systematic assessment of the arguments and proposals that figure in debate over cognitive architecture, in particular as it relates to discussions of cognitive flexibility. But the general line of argument can be summarized as follows:

1. The Standard Model provides an elegant, general approach to explaining core aspects of our cognitive flexibility; and this confers a certain *prima facie* plausibility on the view.
2. Although there are serious practical problems articulating the Standard Model in detail, the main objections do not undermine its core commitments but rather, impose constraints on how the model should be developed.
3. To the extent that recent proposals –including, massive modularity, connectionism and embodied/situated approaches—are intended to *displace* rather than merely supplement the Standard Model, they perform less well in accounting for cognitive flexibility. Moreover, to the extent that such proposals are supposed

to address problems with the Standard Model, they simply *lose track* of the phenomena that required explanation in the first place.

4. Nevertheless, careful analyses of these alternatives yield some important lessons. In particular, they signal the dangers of *over-extending* the Standard Model and expecting too much of cognition to depend on the activity of general-purpose, classical computational mechanisms.
5. Contrary to what many have claimed, the Standard Model is wholly compatible with our current understanding of the relevant evolutionary, neurobiological and psychological considerations.
6. Conclusion: 1-5 provides good reason to take very seriously indeed a version of the Standard Model that accommodates the morals of recent research.

3.3. The Diagnostic Issue: If the Standard Model is broadly correct, then why in practice, has it proven so hard to specify computational systems that generate flexible context appropriate behavior? I conclude by considering this issue. Drawing on arguments from earlier chapters, I maintain that this practical failure is unlikely to indicate any principled problem with the Standard Model. Moreover, I argue, contrary to what Chomsky and others have suggested, that it is unlikely to indicate that such explanations are 'cognitively closed' to us: a mystery that will forever remain beyond our capacity to explain. Rather, I argue that these failures are more likely the consequence of some very severe methodological problems that one faces in trying to understand reasoning processes in a computational manner.

4. Context

4.1 Relation to the author's previous work. The book will draw on previously published material of mine on modularity (1998a, 2000, 2002a, 2005), innateness (1998b, 2002b, 2004) and reasoning (Samuels et al. 1999; Samuels et al. 2002; and Samuels & Stich 2003). But it will also extend and integrate this work to a significant degree by showing, for example, how theories of reasoning and cognitive architecture are relevant to

understanding cognitive flexibility and by assessing non-classical approaches to the study of cognition.

4.2 Relation to research of others. Though cognitive scientists often –and quite reasonably— bracket concerns about flexibility, such issues very frequently come to the fore in large-scale theoretical disputes over cognitive architecture. This is very much reflected in the way that concerns about flexibility have been addressed by prominent researchers across a broad spectrum of disciplines, including philosophy[1], developmental psychology,[2] evolutionary psychology,[3] neuroscience,[4] artificial intelligence[5] and robotics.[6] But despite this multidisciplinary interest, there exists no sustained and systematic book-length treatment of the problems that flexibility poses for contemporary theories of cognitive architecture. The proposed book will fill this gap by mapping the problems and their interconnections, by providing much needed clarification of relevant core concepts and by offering a measured assessment of current proposals in the light of a broad range of theoretical and empirical considerations. The result will, I hope, form the basis for further research and debate within the philosophy of psychology and cognitive science.

4.3. Audience. The proposed book is intended primarily for advanced undergraduates, postgraduate students and scholars in philosophy –especially the philosophy of psychology/cognitive science. Yet the topic –and my approach to addressing it— is a thoroughly interdisciplinary one that bears on concerns in a wide array of other disciplines including psychology, artificial intelligence, robotics and neuroscience. I therefore expect that the book will be of interest to theoretically inclined students and

[1] Carruthers 2003, 2004, 2005; Clark, 1997; Fodor, 1983, and 2000; Horgan and Tienson, 1996; Godfrey-Smith, 1996; Sterelny, 2003
[2] Karmillof-Smith, 1992; Spelke, 2003, Hermer & Spelke, 1996; Hermer-Vazquez *et al*. 1999
[3] Sperber, 1994; Tooby & Cosmides 1992; Pinker, 1998; Gigerenzer et al. 1999, Gigerenzer et al. 2000; Gigernzer & Selten, 2001
[4] Miller and Cohen, 2001
[5] Newell, 1990
[6] Brooks, 1999; Kortenkamp,1998

scholars within the cognitive and social science communities at large. In view of this, it will be written in a way that avoids unnecessary jargon and does not assume any detailed knowledge of the issues under discussion.

Although there are, to my knowledge, no recent books that compete directly with the proposed volume, there are many with which it overlaps both in subject matter and with regard to intended audience. Among the most influential of these are Jerry Fodor's *The Mind Doesn't Work that Way*, Stephen Pinker's *How the Mind Works* and Andy Clark's *Being There*. The present book differs, however, from these in a variety of important respects. Most obviously, in focusing on the Problem of Flexibility, I take as central an issue that receives relatively cursory treatment in these other volumes. Moreover, the position that I defend is different from –and indeed rejected by—these other authors, albeit for reasons which, in my view, are unsatisfactory.

5. Outline of the Book

The book is divided into four parts. In part 1, I spell out the core problems that cognitive flexibility poses for mechanistic accounts of cognition and introduce the Standard Model as a *prima facie* plausible response to these difficulties. In the remaining parts, I develop and defend this proposal in light of the broad array of arguments and alternatives that have emerged from cognitive science in recent decades. In part 2, I discuss the massive modularity hypothesis and its prospects of explaining cognitive flexibility without positing general-purpose reasoning systems of the sort envisaged by the Standard Model. In part 3, I consider a variety of 'non-classical' approaches to the study of cognition and assess their prospects of displacing the Standard Model. Finally, in part 4, I consider the diagnostic issue of why our attempts to understand cognitive flexibility have, to date, met with such limited success.

Part I: Cognitive Flexibility and the Mechanization of Rationality

Chapter 1: Descartes' Challenge

In this chapter I set out Descartes' Challenge and survey the main available responses. In section 1, I provide a first pass formulation of the problem in terms of a tension between two theses: the Flexibility Thesis and Mechanistic Inflexibility Thesis. (See above.) In section 2, I discuss Descartes own response to the Challenge; and in section 3 I map out the main mechanistic responses both in terms of how they respond to the Challenge itself and how they deviate from Descartes' own positive proposal. Broadly speaking, these responses fall into three main groups:

> 1) *Deflationary strategies* that maintain there is no serious challenge to address since human cognition really isn't that flexible after all.
>
> 2) *Universal machine strategies* that deny the Mechanistic Inflexibility Thesis and maintain instead that cognitive flexibility can be explained in terms of the operation of a highly flexible, general-purpose machine of some kind.
>
> 3) *Collaborative strategies* which accept both the cognitive flexibility and mechanistic inflexibility theses but argue that this is compatible with a mechanistic conception of minds on which the overall flexibility of cognition results from complex patterns of collaborative interaction between specialized, inflexible devices.

In section 4, I rehearse the main reasons for thinking that some version of mechanism about the mind is true and, hence, that one of the above strategies needs to be adopted in response to Descartes' Challenge. Finally, in section 5, I address a worry about the individuation of cognitive mechanisms that threatens to collapse the distinction between the universal machine and collaborative strategies.

Chapter 2 The Deflationary Strategy: How Flexible Are We?

One possible response to Descartes' Challenge is to deny the Flexibility Thesis and thereby dismiss the Challenge on the grounds that it seeks an explanation where none is required. In section 1, I argue that this *deflationary response* is deeply unsatisfactory both because it is highly implausible to reject the Flexibility Thesis and because such a

rejection –even if endorsed— would fail to get to the heart of problem that Descartes posed for mechanistic theories.

Nonetheless, the deflationary strategy does raise an important challenge that has received surprisingly little attention, namely: to provide a plausible and suitably perspicuous characterization of flexibility in human beings. With this in mind, I spend the latter sections of chapter 2 developing a characterization that draws on both philosophical and empirical considerations (especially from cognitive neuroscience). First, I provide a plausible characterization of the range of tasks in which we can engage –very roughly, though not infinite, we are in a sense capable of indefinite novelty. Second, I characterize the sense in which our behavior tends to be context appropriate. Finally, I further enrich the account by highlighting certain features characteristic of normal human behavior and explaining how they contribute to our flexibility. These include: a relatively high degree of stimulus independence, the typical absence of perseveration, the capacity to sequence tasks in a multitude of different ways and the related ability to switch between unrelated tasks. What emerges is an account of cognitive flexibility that both supports the contention that humans are flexible is some clear and quite intuitive sense and provides a realistic standard against which to assess mechanistic theories which purport to explain our cognitive flexibility.

Chapter 3 The Universal Machine Strategy and the Mechanization of Practical Reason

If Descartes' Challenge cannot be met by denying the Flexibility Thesis, then an obvious alternative –and the one I seek to defend— is to deny the Mechanistic Inflexibility Thesis and thereby endorse the existence of cognitive mechanisms that are neither specialized nor inflexible. In chapter 3, I introduce the version of this universal machine strategy that I propose to defend: the Standard Model. The discussion is divided into three parts. First, I describe (and illustrate with examples from AI and cognitive science) the sort of architecture advocated by this proposal and highlight the central role it assigns to general-purpose, classical computational mechanisms for practical reasoning. Second, I consider some very real virtues of the proposal that have in my view been seriously underestimated by much recent research in cognitive science. In particular, I spell out the

strikingly elegant manner in which the Standard Model proposes to explain the flexibility of human cognition and argue that this confers on it a certain *prima facie* plausibility.

Despite these virtues, however, the Standard Model has been the subject of a huge array of objections and is now widely regarded by philosophers and scientists as a proposal that should be *displaced* in favor of some more plausible alternative. Many of the objections concern the alleged biological –evolutionary or neurobiological— implausibly of the view. (I discuss these objections in later chapters.) But perhaps the most damaging kind of objection –and the one to which I pay most careful attention in the book— are those that focus on the Standard Model's persistent failure to address the notorious *frame problem* in all its myriad forms. I argue that these problems are in large measure a consequence of the need to simultaneously accommodate both the flexibility and the context-appropriateness of cognition but that strong arguments is required in order to show that they are any more problematic for the Standard Model than any other account of cognitive architecture. I conclude by setting out the main agenda for Parts II and III of the book: To assess the case for displacing the Standard Model in favor of some alternative conception of our cognitive architecture, and to develop further the Standard Model in response to these challenges.

Pat II: The Rejection of Universal Machines: Collaboration and Massive Modularity

Part II is an extended discussion of one of the most recent and intriguing alternatives to Standard Model: the view that our minds –including those parts responsible for reasoning and decision-making—are composed entirely of highly specialized computational mechanisms or 'modules.' One reason to provide an extended treatment of this *massive modularity hypothesis (MM)* is that it's widely believed to provide a satisfactory response to the frame problem within a classical computational framework. Moreover, it does so by pursuing a collaborative response to Descartes' Challenge and denying the existence universal, reasoning systems of the sort envisaged by Standard Model. Indeed, MM constitutes the most plausible version of this view that has been developed so far. I argue, however, that there is no good reason to suppose that MM is required in order to resolve

the frame problem and excellent reason to doubt that (at least in pure form) it can furnish us with an account of cognitive flexibility. I conclude that while it is an open empirical question precisely how modular the human mind is, MM is not plausibly viewed as displacing the Standard Model.

Chapter 4: On the Proper Treatment of (Massive) Modularity

Before addressing issues about the plausibility of MM, we need to get clearer on what the proposal is; and this is not an easy task since the notion of a module is used in such a bewildering array of different ways. In chapter 4, I provide a systematic characterization of modularity that explains the interrelations between the various notions. More precisely, I argue that the main concepts used to characterize modularity –e.g. domain-specificity, encapsulation and innateness— are fruitfully viewed as identifying core respects (or *dimensions* along which) a cognitive mechanism can be more or less flexible and that different conceptions of a module can be thought of as occupying distinct regions in this space of possibilities. I conclude by arguing that the resulting framework provides a useful way of formulating the core similarities between versions of MM and articulating the main respects in which they differ from less modular conceptions of cognitive architecture.

Chapter 5 Modularity, Tractability and the Frame Problem

Perhaps the central argument for MM is that it provides a way of addressing the frame problem within a classical, computational framework. Roughly speaking, such arguments purport to show that only MM can handle the computational feasibility problems associated with the frame problem because only modular mechanisms engage in tractable computation. In this chapter I consider the main versions of this *tractability argument* and argue that they fail to support MM. Nevertheless, I do think that they impose some important constraints on how the Standard Model should be developed. First, I argue that tractability considerations suggest that the Standard Model needs to embrace some form of nativism about domain-specific knowledge. Second, I argue that such considerations mitigate in favor of a version of what is sometimes called *the dual-systems theory*: a view that posits both highly flexible, resource-intensive general-

purpose reasoning mechanisms and more modular inferential devices that provide rapid analyses of the inferential problems that we confront. This chapter draws on material from Samuels (forthcoming).

Chapter 6 Evolutionary Considerations

In addition to the tractability arguments, the most prominent general arguments for MM are ones which purport to show that evolution is more likely to produce specialized modular mechanisms than general-purpose ones of the sort envisaged by Standard Model (Tooby and Cosmides, 1992, 1994; Sperber, 1994; Pinker, 1998). In Chapter 6 I consider the main versions of this argument and show that they are unsatisfactory. Furthermore, drawing on the work of Godfrey-Smith, Sterelny and others, I argue that there may have been strong selectional pressure for flexibility in the evolution of human cognition and that this might have resulted in our possession of general-purpose practical reasoning mechanisms (Godfrey-Smith, 1996; Sterelny, 2003).

Chapter 7: Massive Modularity and the Psychology of Reasoning

Although the general theoretical arguments for MM are unsatisfactory, it might still turn out that MM is vindicated by the experimental evidence. With this in mind, in chapter 7, I consider the current experimental case for modular reasoning mechanisms. No one would claim –not even the most ardent modularist— that the current evidence *precludes* the existence of general-purpose reasoning of the sort posited by the Standard Model. Nonetheless, if the empirical evidence for modular reasoning systems were strong, then the plausibility of such a displacement thesis would be considerably increased. But as a matter of fact, such an analysis of is not plausible. In arguing for this, I consider three 'flagship' proposals where modularity claims are widely endorsed: theory of mind, cheater detection and frequency-based inference. I argue that even in these cases, the empirical evidence for modularity is not strong and that this casts doubt on arguments from the current success of modular approaches to the rejection of non-modular reasoning mechanisms. I conclude by discussing various sorts of evidence from the reasoning literature –especially on individual (e.g. Stanovich, 1999) and cross-cultural differences (e.g. Nisbett et al, 2001; Norenzayan, et al, forthcoming)— and arguing that

they accord well with the sort of dual systems theory sketched towards the end of chapter 5.

Chapter 8: Massive Modularity and the Problem of Flexibility

In Chapter 8, I turn to the issue of whether or not cognitive flexibility can plausibly be explained within a massively modular framework. Advocates of MM, very frequently maintain that this can be done and, moreover, in such a way as to resolve the problems which are widely thought to undermine the Standard Model –i.e. feasibility worries associated with the frame problem (Tooby and Cosmides, 1992; Sperber, 2000; Carruthers 2003). I argue, however, that MM approaches to cognitive flexibility are subject to serious difficulties: Either they end up assuming non-modular cognitive mechanisms of precisely the sort posited by the Standard Model or else they leave it a mystery how flexibility is supposed to be explained. In view of this (and the discussion of the previous three chapters) I conclude that we currently have no good reason to reject the Standard Model in favor of MM.

Part III: Non-Classical Challenges and Alternatives

In Parts I and II, I focused primarily on proposals that cleave to a classical computational account of cognitive processes. In Part III, I consider a range of prominent research programs that make no such a commitment. Although such non-classical approaches have greatly enriched research within cognitive science, it is quite striking just how little headway they have made in addressing the problems which allegedly undermined the Standard Model: in particular, the frame and relevance problems in all their myriad forms. Again, the logic that drives Descartes' Challenge is much in evidence. Non-classical cognitive scientists have had considerable success developing models that exhibit appropriate response profiles in relatively restricted task domains. But what have not been forthcoming are models that are both context appropriate *and* capable of performing a very wide range of tasks. Indeed, to the extent that such a capacity appears to depend on practical reasoning, non-classical cognitive science has yet to provide *any*

serious mechanistic response to Descartes' Challenge. One principal aim of Part III is to highlight this 'bate-and-switch' pattern in the development of cognitive science and to spell out some of the more serious problems that confront non-classical approaches to cognitive flexibility. A second aim is to further elaborate the Standard Model in the light of the various arguments and theories to have emerged from non-classical cognitive science.

Chapter 9 Embodiment and Situatedness

In Chapter 9, I focus on a broad confederation of views that are often grouped under the headings of *embodied* or *situated* approaches to cognition. Although such views vary considerably in detail, they almost invariably share an antagonism to the central claims of the Standard Model (Anderson, 2003). In chapter 9, I first distinguish between moderate and more radical versions of the Embodiment Thesis and show that the more moderate (and plausible) versions are easily reconciled with the Standard Model. Next, I turn my attention to the radical version. If true, this view would require the rejection of the Standard Model since it denies the existence of practical reasoning systems of the sort envisaged by latter view. I show, however, that the main arguments for radical embodiment are unsatisfactory and that there is also no reason to suppose that it can accommodate the flexibility of cognition. In doing so, I consider what is widely regarded as the most influential research program within the embedded/situated tradition --Rodney Brooks' subsumption approach in robotics (Brooks 1999). I argue that while this and related work in robotics provide important insight into the computational requirements on real-time behavior, far from undermining the Standard Model, it in fact *strengthens* it by highlighting the need for general-purpose reasoning mechanisms.

Chapter 10 Connectionism

In Chapter 10, I consider connectionist challenges to the Standard Model. In my view, connectionism has provided cognitive science with both a battery of useful modeling techniques and important insight into a broad array of cognitive phenomena. But on one familiar story, it was supposed to deliver far more than this. Specifically, it was intended to provide a paradigm shift for cognitive science: one that addressed the central

deficiencies of the classical approach and made good on the blend of associationism and empiricism that has dominated much of psychology since the late eighteenth century. In short: it was intended to *displace* as opposed to merely supplement the sort of proposal developed in the earlier parts of this book.

The central aim of chapter 10 is to argue that research over the past two decades does not bear out this view; and in particular, that connectionist research have yet to yield any serious alternative to the Standard Model as a mechanistic approach to cognitive flexibility. Although a detailed survey of the literature is beyond the scope of the current chapter, I highlight a series of general issues and trends that have emerged over the past few decades and relate them to Descartes' Challenge. First, I argue that the actual models to have been produced –as opposed to the rhetoric that sometimes accompanies them— suggest that connectionist views are committed to a modular-cum-nativist conception of cognitive mechanisms. To that extent these approaches would appear committed to a collaborative explanation of cognitive flexibility and subject to exactly the same sorts of problems discussed in connection with MM. Second, I argue that the prospects of pursuing the universal machines strategy within a connectionist framework appear bleak. As many connectionist researchers themselves acknowledge, there has been little success modeling the sorts of 'higher' processes associated with flexible cognition (Elman et a., 1996). Moreover, there are familiar problems that connectionist approaches encounter in trying to develop highly flexible systems –including, limitations on expressive power and the so-called problem of catastrophic unlearning. To be sure, none of this rules out the possibility of *hybrid* architectures that incorporate both classical and connectionist systems (Sloman, 1996); but it does suggest that the proposed *displacement* of the Standard Model by some connectionist alternative will be hard to sustain.

Chapter 11: The Standard Model Meets Neuroscience

In chapter 11, I consider the connection between recent neuroscientific research and the topics addressed in previous chapters of the book. Broadly speaking, my discussion can be divided into two parts. First, I consider two challenges from neuroscience, which if correct, would appear to threaten the Standard Model. According to the first challenge, the development of cognitive neuroscience calls into the question the claim that we

possess any general-purpose mechanisms for planning (Burgess et al.). According to the second, not only do we lack general-purpose planning mechanisms, but the paradigmatic instances of decision-making –viz. *conscious* decision-making— are not even causally relevant to the production of behavior (Wenger, 2002; 2003). I argue that these challenges fail and, indeed, turn on serious confusions. I conclude the chapter by suggesting that the burgeoning program of cognitive neuroscience –and, in particular, recent research on the prefrontal cortex— comport quite well with the positive views developed in the previous chapters.

Conclusion: Of Problems and Mysteries

A central burden of the proposed book is to argue that, when appropriately elaborated, the Standard Model provides a plausible, general approach to explaining human cognitive flexibility, one which we have no good reason to reject and that comports well with a broad range of theoretical and empirical considerations drawn from philosophy, cognitive psychology, robotics, neuroscience and evolutionary theory. But this poses something of a puzzle: If the Standard Model is such a good idea, then why has so little headway been made in providing detailed, plausible computational accounts of human reasoning and decision-making? In this chapter I conclude the book by considering four possible responses to this question:

i. *The No-Theory View*: Though the Standard Model is our best current account, it is still fundamentally in error. The failure of the Standard Model thus reflects a fundamental limit of contemporary cognitive science: that we have no adequate theory of central processes (Fodor, 2000).

ii. *The Mysterian View*: The correct response to Descartes' Challenge is 'cognitively closed' to us: a 'mystery' that we humans are incapable of penetrating by virtue of our intrinsic mental limitations (Chomsky, 1987; McGinn, 1993).

iii. *Fundamental Insight View*: Though the Standard Model is right as far as it goes, there are fundamental insights about the nature of mental computation that are yet to be satisfactorily elaborated and

incorporated in the study of human cognition. On one recent view, for example, a satisfactory theory of mind must adopt a quantum theory of computation as opposed to the 'classical' view typically assumed by conventional computer science.

iv. *Methodological View*: Failures to formulate detailed and empirically plausible computational models of human reasoning result from a series of deep methodological problems that plague the study of central processes.

These claims are neither exhaustive nor mutually exclusive; and they all remain genuine empirical possibilities. Nevertheless, I argue that the methodological view is both sufficient to explain the lack of progress in elaborating the Standard Model and more plausible than the alternatives.

References

Anderson, M. L. (2003) Embodied cognition: A field guide. *Artificial Intelligence* 149(1): 91-130.

Brooks, R.A. (1999) *Cambrian Intelligence.* MIT Press, Cambridge, MA, 1999.

Burgess, P.W., Simons, J.S., Coates, L.M., Channon, S. (2005) The search for specific planning processes. In R. Morris (Ed.) *Cognitive Psychology of Planning*. Oxford University Press, Oxford.

Carruthers, P (2004) Practical reasoning in a modular mind. *Mind and Language*, 19

Carruthers, P (2003) On Fodor's Problem. *Mind and Language*, 18.

Carruthers, P (2005) Distinctively human thinking: modular precursors and components. In P.Carruthers, S.Laurence and S.Stich (eds.), *The Innate Mind: structure and content*, Oxford University Press.

Chomsky, N (1987) *Language and Problems of Knowledge. The Managua Lectures*. Cambridge, MA and London: The MIT Press,

Clark, A. (1997) *Being There: Putting Brain, Body And World Together Again* MIT Press.

Elman, J. L., et al. (1996) *Rethinking Innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Fodor, J. A. (1983) *The modularity of mind: An essay on faculty psychology.* MIT Press.

Fodor, J. A. (2000) *The mind doesn't work that way; the scope and limits of computational psychology*. MIT Press.

Godfrey-Smith, P (1996) *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.

Karmiloff-Smith, A. (1992) *Beyond Modularity: A Developmental Perspective on Cognitive Science.* Cambridge, MA/London: MIT Press.

Kortenkamp, D, R.P. Bonasso and R. Murphy (eds.) (1998) *Artificial Intelligence and Mobile Robotics.* AAAI Press/ MIT Press.

Hameroff, S. (1998) Quantum computation in brain microtubules?.*Philosophical Transactions Royal Society London* (A) 356:1869-1896 (1998).

Hermer-Vasquez, L., Spelke, E. S. & Katsnelson, A. S. (1999) Sources of flexibility in human cognition: Dual-task studies of space and language. *Cognitive Psychology, 39*, 3-36.

Hermer, L. & Spelke, E. S. (1996) Modularity and development: The case of spatial reorientation. *Cognition, 61*, 195-232.

Horgan, T. and Tienson, J (1996) *Connectionism and the Philosophy of Psychology*, Bradford/M.I.T.

Pinker, S. (1998) *How the Mind Works* W. W. Norton & Co Inc.

Gigerenzer, G. & Selten, R. (Eds.) (2001) *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Gigerenzer, G. (2000) *Adaptive thinking: Rationality in the real world*. New York: Oxford University Press.

Gigerenzer, G., Todd, P. M. & the ABC Research Group (1999) *Simple heuristics that make us smart*. New York: Oxford University Press.

McGinn, C (1993) *Problems in Philosophy: the Limits of Inquiry*, Blackwell.

Miller, E.K. and Cohen, J.D. (2001) An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24:167-202

Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001) Culture and systems of thought: Holistic vs. analytic cognition. *Psychological Review*, 108, 291-310.

Norenzayan, A., Smith, E.E., Kim, B. J. & Nisbett, R. E. (In press) Cultural preferences for formal versus intuitive reasoning. *Cognitive Science.*

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, Mass.: Harvard University Press.

Newell, A., and H. A. Simon. 1972. *Human Problem Solving*. Englewood Cliffs, N.J.: Prentice-Hall.

Newell, A., and H. A. Simon. 1976. Computer science as empirical inquiry: Symbols and search. *Commun. Assoc. Comput. Machinery* 19:111-26.

Samuels R. (1998a) Evolutionary Psychology and The Massive Modularity Hypothesis *British Journal for the Philosophy of Science*, 49, 575-602.

Samuels R. (1998b) What Brains Won't Tell Us About the Mind: A Critique of the Neurobiological Argument against Representational Nativism. *Mind and Language*, 4, 548-570.

Samuels R., S. Stich and P. Tremoulet (1999) Rethinking Rationality: From Bleak Implications to Darwinian Modules. In E. Lepore and Z. Pylyshyn (eds.) *What is Cognitive Science?* Oxford: Blackwell.

Samuels, R. (2000) Massively Modular Minds: Evolutionary Psychology and Cognitive Architecture. In P. Carruthers (ed.) *Evolution and the Human Mind*. Cambridge University Press.

Richard Samuels, Stephen Stich & Luc Faucher (2003) Reasoning and Rationality. In I. Niiniluoto, M. Sintonen, & J. Wolenski (eds.) *Handbook of Epistemology* pp. 1-50. Dordrecht: Kluwer.

Samuels R. (2002a) The spatial reorientation data do not support the thesis that language is the medium of cross-modular thought. *Behavioral and Brain Sciences*, 25:6.

Samuels R. (2002b) Nativism in Cognitive Science. *Mind and Language*, 17, 3, 233-265.

Samuels, R and Stich S. (2003) Rationality and Psychology. In A. Mele (ed.) *Rationality: Oxford Readings in Philosophy Series*. New York: Oxford University Press.

Samuels, R (2004) Innateness in cognitive science. *Trends in Cognitive Sciences* Vol.8 No.3.

Samuels, R. (2005) The Complexity of Cognition: Tractability Arguments for Massive Modularity. In Carruthers, Laurence & Stich (eds.) *The Innate Mind: Structure and Content*. OUP.

Sloman, S. A. (1996) The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.

Spelke, E.S. (2003) What makes humans smart? In D. Gentner and S. Goldin-Meadow (Eds.), *Advances in the Investigation of Language and Thought*. Cambridge, MA: MIT Press.

Sperber, D. (1994) The modularity of thought and the epidemiology of representations. In L. Hirschfeld and S. Gelman (Eds.) *Mapping the Mind*, 39-67. New York: Cambridge University Press.

Sperber, D (2002) In Defense of massive modularity (in Dupoux, E. *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*, pp.47-57. Cambridge, Mass. MIT Press.

Sperber, D (2005) Modularity and relevance: How can a massively modular mind be flexible and context-sensitive? In Carruthers, Laurence & Stich (eds.) *The Innate Mind: Structure and Content*. OUP.

Stanovich, K (1999) *Who is Rational?* Laurence Erlbaum Associates.

Sterelny, K (2003) *Thought in a Hostile World* Blackwell.

Tooby, J. & Cosmides, L. 1992 The psychological foundations of culture. In *The adapted mind: Evolutionary psychology and the generation of culture* (ed. J. Barkow, L. Cosmides, & J. Tooby), pp. 19-136. NY: Oxford University Press.

Wegner, D. M. (2002) *The Illusion of Conscious Will* Cambridge, MA: The MIT Press.

Wegner, D. M. (2003) The mind's best trick: How we experience conscious will. *Trends in Cognitive Science, 7,* 65-69.