CHAPTER THREE

Is the Human Mind Massively Modular?

Richard Samuels

Introduction: Minds as Mechanisms

Among the most pervasive and fundamental assumptions in cognitive science is that the human mind (or mind-brain) is a *mechanism* of some sort: a physical device composed of functionally specifiable subsystems. On this view, *functional decomposition* – the analysis of the overall system into functionally specifiable parts – becomes a central project for a science of the mind, and the resulting theories of *cognitive architecture* essential to our understanding of human psychology.

None of this is, of course, in any way distinctive of massive modularity (MM). On the contrary, these commitments have very wide acceptance among cognitive scientists and have done so since the inception of the discipline. If MM is of value, then, it is not merely because it advocates a mechanistic conception of minds; it must also embody a distinctive, interesting and plausible hypothesis about our cognitive architecture.

The central aim of this chapter is to argue that, as things stand, there is little reason to suppose advocates of MM have succeeded in articulating such a hypothesis. On some readings, the MM hypothesis is plausible but banal. On other readings, it is radical but wholly lacking in plausibility. And on still further (more moderate but still interesting) interpretations, it remains largely unsupported by the available arguments since there is little reason to suppose that *central systems* – such as those for reasoning and decision making – are modular in character. Contrary to what Peter Carruthers and others maintain, then, the case for MM is not strong. But it would be wrong to conclude, as Jesse Prinz and others have done, that the mind is not modular to any interesting degree (Prinz, chapter 2, IS THE MIND REALLY MODULAR?). This is because it is very plausible that relatively peripheral regions of cognition – especially for low-level perception – are modular in character. I thus advocate a middle way between those who endorse a thoroughgoing massive modularity and those who reject modularity altogether.

Here's how I'll proceed. In section 1, I clarify the main commitments of MM and the attendant notion(s) of a module. In section 2, I sketch some of the main theoretical arguments for MM and highlight their deficiencies. In section 3, I very briefly discuss some problems with the experimental case for MM. In section 4, I outline some reasons for finding MM at least in radical form implausible. Finally in section 5, I argue against those who not merely reject MM but deny minds are modular to any interesting degree.

1 What's at Issue?

To a first approximation, MM is the hypothesis that the human mind is largely or entirely composed from a great many modules. Slightly more precisely, MM can be formulated as the conjunction of three claims:

Composition Thesis: The human mind is largely or entirely composed from *modules*.
Plurality Thesis: The human mind contains a great many modules.
Central Modularity: Modularity is found not merely at the periphery of the mind but also in those *central* regions responsible for reasoning and decision-making.

In what follows I assume advocates of MM are committed to the conjunction of these claims. Even so, each is amenable to a variety of different interpretations that vary considerably in interest and plausibility. More needs to be said if we are to get clearer on what's at issue.

1.1 Composition Thesis

MM is in large measure a claim about the kinds of mechanisms from which our minds are composed – viz. it is largely or even entirely composed from *modules*.[1] As stated, this is vague in at least two respects. First, it leaves unspecified the precise extent to which minds are composed from modules. In particular, this way of formulating the proposal accommodates two different positions, which I call *strong* and *weak* massive modularity. According to strong MM *all* cognitive mechanisms are modules. Such a view would be undermined if we were to discover that any cognitive mechanism was nonmodular in character. By contrast, weak MM maintains only that the human mind – including those parts responsible for central processing – are *largely* modular in structure. In contrast to strong MM, such a view is clearly compatible with the claim that there are some nonmodular mechanisms. So, for example, the proponent of weak MM is able to posit the existence of some nonmodular devices for reasoning and learning.

A second respect in which the above Composition Thesis is vague is that it leaves unspecified what modules *are*. For present purposes, this an important matter since the interest and plausibility of the thesis turns crucially on what one takes modules to be.[2]

Minimal processing modules

At one extreme, modules are just distinct, functionally characterized cognitive mechanisms of the sort that correspond to boxes in a cognitive psychologist's flow diagram

(Fodor, 2002). This *minimal* conception of modules is one some advocates of MM really appear to adopt. So, for example, Tooby and Cosmides (1992) characterize modules as "complex structures that are functionally organized for processing information" (Tooby and Comsides, 1992; see also Pinker, 2005). Even so, it is too weak for an interesting and distinctive Composition Thesis since almost all cognitive scientists agree that minds are composed from such mechanisms. But it does not follow from this alone that a distinctive version of MM cannot be formulated in terms of this anodyne notion (Fodor, 2005; Prinz, chapter 2). In particular, some have suggested that an MM formulated with this notion is interesting not because it claims minds are composed of minimal modules but because it claims there are a *large number* of such mechanisms (Carruthers, chapter 1, THE CASE FOR MASSIVELY MODULAR MODELS OF MIND). I consider this suggestion in section 1.2.

Robust modules

So, the minimal notion of a module won't suffice for an interesting Composition Thesis. But debate in cognitive science frequently assumes some more robust conception of modularity; of which the most well known and most demanding is the one developed in Fodor (1983). On this view, modules are functionally characterizable cognitive mechanisms which are (at least paradigmatically) domain specific, informationally encapsulated, innate, mandatory, fast relative to central processes, shallow, neurally localized, exhibit characteristic breakdown patterns, are inaccessible, and have characteristic ontogenetic timetables (Fodor, 1983).[3]

Though the full-fledged Fodorian notion has been highly influential in many areas of cognitive science (Garfield, 1987), it has not played much role in debate over MM,[4] and for good reason. The thesis that minds are largely or entirely composed of Fodorian modules is obviously implausible. Indeed, some of the entries on Fodor's list – relative speed and shallowness, for example – make little sense when applied to central systems (Carruthers, chapter 1; Sperber, forthcoming). And even where Fodor's properties can be sensibly ascribed – as in the case of innateness – they carry a heavy justificatory burden that few seem much inclined to shoulder (Baron-Cohen, 1995; Sperber, 1996).

In any case, there is a broad consensus that not all the characteristics on Fodor's original list are of equal theoretical import. Rather, domain specificity and informational encapsulation are widely regarded as most central. Both these properties concern the architecturally imposed[5] informational restrictions to which cognitive mechanisms are subject – the range of representations they can access – though the kinds of restriction involved are different.

Domain specificity is a restriction on the representations a cognitive mechanism can take as *input* – that "trigger" it or "turn it on." Roughly, a mechanism is domain specific (as opposed to domain general) to the extent that it can only take as input a highly restricted range of representations.[6] Standard candidates include mechanisms for low-level visual perception, face recognition, and arithmetic.

Informational encapsulation is a restriction on the kinds of information a mechanism can use as a *resource* once so activated – paradigmatically, though not essentially, information stored in memory. Slightly more precisely, a cognitive mechanism is encapsulated to the extent it can access, in the course of its computations, less than all of the information available to the organism as a whole (Fodor, 1983). Standard candidates include mechanisms, such as those for low-level visual perception and phonology, which do not draw on the full range of an organism's beliefs and goals.

To be sure, there are many characteristics other than domain specificity and encapsulation that have been ascribed to modules. But if one uses "module" as more than a mere terminological expedient – as more than just a nice way of saying "cognitive mechanism" – yet deny they possess either of these properties, then one could with some justification be accused of changing the subject. In view of this, I will tend when discussing more robust conceptions of modularity to assume that modules must be domain specific and/or encapsulated. This has a number of virtues. First, these properties clearly figure prominently in dispute over MM. Moreover – and in contrast to the minimal module version of the Composition Thesis discussed earlier – the claim that minds are largely or entirely composed of domain specific and/or encapsulated mechanisms is a genuinely interesting one. Not only does it go beyond the banal claim that our minds are comprised from functionally characterizable cognitive mechanisms; but it is also a claim that opponents of MM almost invariably deny. In later sections I will consider the plausibility of this thesis; but first I need to discuss the other two theses associated with MM.

1.2 Plurality Thesis

Advocates of MM sometimes suggest their position is distinctive and interesting in part because it countenances a large number of cognitive mechanisms. In spelling out their position, for example, Tooby and Cosmides maintain: "our cognitive architecture resembles a confederation of hundreds or thousands of functionally dedicated computers (often called modules)." (Tooby and Cosmides, 1995, p. xiv)

This suggests a commitment to what I earlier called the Plurality Thesis: i.e. that the human mind contains a great many modules.

How interesting and distinctive is this thesis? If formulated in terms of a robust notion of modularity, it appears quite radical. After all, there are many who deny that domain specific and/or encapsulated devices have a substantial role to play in our cognitive economy. But what if the Plurality Thesis is formulated using the minimal notion of a module? Would this be an interesting and distinctive claim? Certainly, advocates of MM sometimes appear to suggest that it is. For example Carruthers claims: "those evolutionary psychologists who have defended the claim that the mind consists of a great many modular components (Tooby and Cosmides, 1992; Sperber, 1996; Pinker, 1997) are defending a thesis of considerable interest, even if 'module' just means 'component.'" (Carruthers, chapter 1)

The idea seems to be that while virtually all cognitive scientists think minds are composed from distinct, functionally characterizable devices, many reject the claim that there are *lots* of such devices. According to Carruthers, for example, it is a claim rejected by "those who . . . picture the mind as a big general-purpose computer with a limited number of distinct input and output links to the world" (Carruthers, chapter 1, p. 000). In which case, it may seem that MM is an interesting thesis even when formulated in terms of the minimal notion of a module.

On reflection, however, it's hard to see how this could be right: how a mere plurality of functionally specifiable mechanisms could make for an interesting and distinctive MM. This is because even radical opponents of MM endorse the view that minds contain a great many such components. So, for instance, the picture of the mind as a big general-purpose, "classical" computer – roughly, the sort of general-purpose device that manipulates symbols according to algorithmically specifiable rules – is often (and rightly) characterized as being firmly at odds with MM. Yet big general-purpose computers are not simple entities. On the contrary, they are almost invariably decomposable into a huge number of functionally characterizable submechanisms.[7] So, for example, a standard von Neumann-type architecture decomposes into a calculating unit, a control unit, a fast-to-access memory, a slow-to-access memory, and so on; and each of these decomposes further into smaller functional units, which are themselves decomposable into submechanisms, and so on. As a consequence, a standard von Neumann machine will typically have hundreds or even thousands of subcomponents.[8] Call this a version of massive modularity if you like. But it surely isn't an interesting or distinctive one.

1.3 Central Modularity

So far we have discussed the Composition and Plurality theses and seen that both are interesting on a robust construal of modules, but that neither seems interesting or distinctive on the minimal conception. But there is another thesis that requires our attention, the thesis of Central Modularity, which states that modules are found not merely at the periphery of the mind but also in those *central* regions responsible for reasoning and decision making.

This does not strictly follow from any of the claims discussed so far since one might deny there are any central systems for reasoning and decision making. But this is *not* the view that advocates of MM seek to defend. Indeed, a large part of what distinguishes MM from the earlier, well-known modularity hypothesis defended by Fodor (1983) and others is that the modular structure of the mind is not restricted to input systems (those responsible for perception, including language perception) and output systems (those responsible for producing behavior). Advocates of MM accept the Fodorian thesis that such peripheral systems are modular. But *pace* Fodor, they maintain that the *central* systems responsible for reasoning and decision making are largely or entirely modular as well (Jackendoff, 1992). So, for example, it has been suggested that there are modules for such central processes as social reasoning (Cosmides and Tooby, 1992), biological categorization (Pinker, 1994), and probabilistic inference (Gigerenzer, 1994 and 1996). In what follows, then, I assume MM is committed to some version of Central Modularity.

Again, how interesting a thesis is this? If formulated in terms of the minimal notion, it's hard to see how Central Modularity could be an interesting and distinctive one. After all, even those who endorse paradigmatically nonmodular views of central cognition can readily accept the claim. For example, advocates of the "Big Computer" view of central systems can accept the claim that central cognition is entirely subserved by a great many minimal modules since big computers are themselves composed of a great many such entities. All this is, of course, wholly compatible with there being some suitable modification that makes for an interesting version of Central Modularity. But the point I want to insist on here is that if one's arguments succeed only in supporting this version of the thesis, then they fail to support a distinctive and interesting version of MM.

Things look rather different if a robust conception of modules is adopted. Here, the degree to which one's version of Central Modularity is interesting will depend on (a) the extent to which central cognition is subserved by domain specific and/or encapsulated mechanisms and (b) how many such modules there are. Both these questions could be answered in a variety of different ways. At one extreme, for example, one might adopt the following view:

> *Strong Central Modularity*: All central systems are domain specific and/or encapsulated, and there are a great many of them.

That would be a genuinely radical position since it implies that there are no domain general, informationally unencapsulated central systems. But this Strong Central Modularity is very implausible. For as we will see in later sections, there are no good reasons to accept it, and some reason to think it is false. At the other extreme, one might maintain that:

> *Weak Central Modularity*: There are a number of domain specific and/or encapsulated central systems, but there are also nonmodular – domain general and unencapsulated – central systems as well.

Such a proposal is not without interest. But it is not especially radical in that it does not stray far from the old-fashioned peripheral modularity advocated by Fodor. Nor is it implausible in the way that Strong Central Modularity is. Nonetheless, as we will see, there is at present little reason to accept it. The general arguments for MM don't support this view; and the empirical case for even this version of Central Modularity is surprisingly weak.

2 The Allure of Massive Modularity

Though MM is an empirical hypothesis, many claim it is plausible in the light of quite general, theoretical arguments about the nature of evolution, cognition and computation. But even the most prominent and plausible of these arguments are unsatisfactory since they fail to discriminate between interesting versions of MM – such as those that imply Strong or Weak Central Modularity – and other widespread views.[9]

2.1 Evolvability

A discussion of evolutionary stability in Simon (1962) is sometimes invoked as an argument for MM (Carston, 1996; Carruthers, chapter 1). According to Simon, for an evolutionary process to reliably assemble complex functional systems – biological systems in particular – the overall system needs to be *semi-decomposable*: hierarchically organized from components with relatively limited connections to each other. Simon illustrates the point with a parable of two watchmakers – Hora and Tempus

– both highly regarded for their fine watches. But while Hora prospered, Tempus became poorer and poorer and finally lost his shop. The reason:

> The watches the men made consisted of about 1000 parts each. Tempus had so constructed his that if he had one partially assembled and had to put it down – to answer the phone, say – it immediately fell to pieces and had to be reassembled from the elements . . . The watches Hora handled were no less complex . . . but he had designed them so that he could put together sub-assemblies of about ten elements each. Ten of these sub-assemblies, again, could be put together into a larger sub-assembly and a system of ten of the latter constituted the whole watch. Hence, when Hora had to put down a partly assembled watch in order to answer the phone, he lost only a small part of his work, and he assembled his watches in only a fraction of the man-hours it took Tempus. (Simon, 1962)

The obvious moral – and the one Simon invites us to accept – is that evolutionary stability requires that complex systems be hierarchically organized from dissociable subsystems; and according to Carruthers and Carston, this militates in favor of MM (Carston, 1996, p. 75).

Response. Though evolutionary stability may initially appear to militate in favor of MM, it is in fact only an argument for the familiar mechanistic thesis that complex machines are hierarchically assembled from (and decomposable into) many subcomponents. But this clearly falls short of the claim that all (or even any) are domain specific or encapsulated. Rather it supports at most the sort of banal Plurality Thesis discussed earlier; one that is wholly compatible with even a Big Computer view of central processes. All it implies is that if there are such complex central systems, they will need to be hierarchically organized into dissociable subsystems – which incidentally, was the view Simon and his main collaborators endorsed all along (Simon, 1962; Newell, 1990).

2.2 Analogy with other biological systems

Throughout the biological world – from cells to cellular assemblies, whole organs and so on – one finds hierarchical organization into semi-decomposable components. We should expect the same to be true of cognition (Chomsky, 1980; Carruthers, chapter 1).

Response. Same problem as the previous argument. Though all this is correct, it is at most an argument for the claim that our minds are semi-decomposable systems – hierarchically organized into dissociable subsystems – a conclusion that is in no way incompatible with even the most radically nonmodular accounts of central systems.

2.3 Task specificity

There are a great many cognitive tasks whose solutions impose quite different demands. So, for example, the demands on vision are distinct from those of speech recognition, of mind-reading, cheater-detection, probabilistic judgment, grammar induction, and so on. Moreover, since it is very hard to believe there could be a single general inferencing mechanism for all of them, for each such task we should postulate the existence of a distinct mechanism, whose internal processes are computationally specialized for processing different sorts of information in the way required to solve the task (Carruthers, chapter 1; Cosmides and Tooby, 1992, 1994).

Response. Two points. First, if the alternatives were MM or a view of minds as comprised of just a single general-purpose cognitive device, then I too would opt for MM. But these are clearly not the only options. On the contrary, one can readily deny that central systems are modular while still insisting there are plenty of modules for perception, motor control, selective attention, and so on. In other words, the issue is, not merely whether some cognitive tasks require specialized modules but whether the sorts of tasks associated with central cognition – paradigmatically, reasoning and decision making – typically require a proliferation of such mechanisms.

Second, it's important to see that the addition of functionally dedicated mechanisms is not the only way of enabling a complex system to address multiple tasks. An alternative is to provide some relatively functionally inspecific mechanism with requisite bodies of information for solving the tasks it confronts. This is a familiar proposal among those who advocate nonmodular accounts of central processes. Indeed, advocates of nonmodular reasoning architectures routinely assume that reasoning devices have access to a *huge* amount of specialized information on a great many topics, much of which will be learned but some of which may be innately specified (Newell, 1990; Anderson, 1990). Moreover, it is one that plausibly explains much of the proliferation of cognitive competences that humans exhibit throughout their lives – e.g. the ability to reason about historical issues as opposed to politics or gene splicing or restaurants. To be sure, it *might* be that each such task requires a distinct mechanism, but such a conclusion does not flow from general argument alone. For all we know, the same is true of the sorts of tasks advocates of MM discuss. It may be that the capacity to perform certain tasks is explained by the existence of specialized mechanisms. But how often this is the case for central cognition is an almost entirely open question that is not adjudicated by the argument from task specificity.

2.4 Bottleneck argument

If central processing were subserved by only a small number of general-purpose mechanisms as opposed to a multiplicity of modules, there would be a kind of tractability problem: a serious *processing bottleneck* that would prohibit the formation of beliefs on different topics in real time (Carruthers, chapter 1).

Response. For this to be an objection to the claim that we possess only a small number of central systems it would need to be that, as a matter of fact, central cognition is *not* subject to such bottlenecks. But it's far from clear this is true. Second, even if we assume for the sake of argument that there are few such bottlenecks, it's utterly unclear MM is the only way to avoid them. What MM permits is the *parallel* operation of multiple mechanisms. But it's not the only way to exact this benefit since there are at least two other (not mutually exclusive) ways that parallelism can increase the speed of processing. First, at the level of cognitive architecture, it may be that the component cognitive parts of a given non-modular mechanism operate at the same time. Second, there may be (and surely is) parallelism at the level of *implementation* so that individual, non-modular cognitive mechanisms have a parallel implementation within the circuits of the brain.

2.5 Computational tractability

It is common to argue for MM on the grounds that the alternatives are computationally intractable. In brief, the argument is as follows: Human cognitive processes are realized by computational mechanisms. But for this to be the case, our cognitive mechanisms would need to be computationally tractable; and this in turn requires that they be informationally encapsulated – that they have access to less than all the information available to the mind as a whole. Hence, the mind is composed of informationally encapsulated cognitive mechanisms.

Response. If one is disinclined to accept a computational account of cognitive processes, this is not an argument one is likely to take seriously. But even if one endorses a computational account of the mind, the argument still doesn't work. There is a long story about what's wrong (see Samuels, 2005). But the short version is that computational tractability does not require informational encapsulation in any standard sense of that expression; and the claim that some *other* kind of "encapsulation" is involved is tantamount to relabeling. As ordinarily construed, a mechanism is encapsulated only if – by virtue of architectural constraints – there is some relatively determinate class of informational states that it cannot access. (For example, the paradigmatic cases of encapsulated devices – low-level perceptual devices – cannot access the agent's beliefs or desires in the course of their computations.) While such an architectural constraint may be one way to engender tractability, it is clearly not the only way. Rather, what's required is that the mechanism be *frugal* in its use of information: that it not engage in exhaustive search but only use a restricted amount of the available information. Moreover, this might be achieved in a variety of ways, most obviously by heuristic and approximation techniques of the sort familiar from computer science and Artificial Intelligence.

So, it would seem that the tractability argument fails: frugality, not encapsulation, is what's required. Carruthers has responded to this, however, by drawing a distinction between two notions of encapsulation: the standard notion, which he calls "narrow-scope encapsulation" and another notion, "wide-scope encapsulation," on which the "operations of a module can't be affected by most of the information in a mind, without there being some determinate subdivision between the information that can affect the system and the information that can't" (Carruthers, chapter 1, p. 000). But this really isn't an interesting notion of encapsulation. For not only is it different from what most theorists mean by "encapsulation," but it's simply what you get by denying exhaustive search; and since virtually no one thinks exhaustive search is characteristic of human cognition, the present kind of "encapsulation" is neither distinctive nor interesting.

3 The Empirical Case for Massive Modularity

So far I have focused on the theoretical arguments for MM and found them wanting. But what of the empirical evidence? Many have claimed MM is plausible in light of the available experimental data (Sperber, 1996; Pinker, 1997; Pinker, 2005). I am unconvinced. Although there is considerable evidence for relatively low-level modular mechanisms – a point I return to in section 5 – the experimental case for *central* modularity is not strong (Samuels, 1998; Fodor, 2000). To the extent things seem otherwise, I suspect it is because the interpretation of data is heavily driven by a prior acceptance of the general theoretical arguments for MM – arguments which as we have seen there is little reason to endorse.[10]

In some cases, the reliance on general theoretical considerations in the absence of convincing data is egregious. (My current favorite example is the "homicide module" hypothesis advocated by Buss and his collaborators; see Buss and Duntley, 2005.) But even in cases where the influence of such considerations is less obvious, the data for Central Modularity are uncompelling. Often the problem is that the putative evidence for specialized modules can be better explained in terms of other, less specific processes. This is well illustrated by arguably the flagship case of a putative reasoning module: a dedicated mechanism for social contract reasoning (Comsides and Tooby, 1992; Gigerenzer and Hug, 1992). Advocates of this hypothesis sometimes represent it as a kind of modularist assault on the "doctrinal 'citadel' of . . . general-purpose processes": human reasoning (Cosmides and Tooby, 1992). But the main experimental support for the hypothesis – evidence of so-called content effects in Wason's Selection Task – can be very plausibly explained in terms of quite general features of language comprehension and, hence, provides no support for a dedicated social contract reasoning mechanism (Sperber et al., 1995; Fodor 2000; Sperber and Girotto, 2003).

Another problem with the experimental case for Central Modularity is that even where the data suggest some kind of specialized cognitive structure for a given domain, it seldom adjudicates clearly between claims about the existence of specialized *mechanisms* and claims about the existence of specialized bodies of *knowledge*, such as a mentally represented theory. The former are modularity theses in the relevant sense while the latter are wholly compatible with a highly nonmodular account of central processing on which different bodies of knowledge are used by a small set of nonmodular mechanisms (Newell, 1990; Gopnik and Meltzoff, 1997; Fodor, 2000). This point is well illustrated by the debate over folk biology. Many have proposed the existence of specialized cognitive structures for folk biology. But while some maintain it is subserved by a dedicated module (Pinker, 1994; Atran, 1998, 2001), others claim merely that we posses a body of information – a theory – that is deployed by relatively inspecific inferential devices (Carey, 1995). The problem is that the main available evidence regarding folk biology – e.g. cross-cultural evidence for universality, developmental evidence of precocity, anthropological evidence for rapid cultural transmission, and so on – fails to adjudicate between these options. What they suggest is that folk biology involves some dedicated – and perhaps innate – cognitive structure. But once the general arguments for MM are rejected, there is little reason to interpret it as favoring a *modular* account of folk biology.

Finally, even where the evidence for previously unrecognized modules is strong, it seldom turns out to be clear evidence for *central* modularity. Consider, for example, the "geometric module" hypothesis advocated by Cheng, Gallistel, and others (Gallistel, 1990; Hermer-Vazquez et al., 1999). Though contentious, the evidence for such a device is quite compelling. But this does not support a modular view of *central* processes since the geometric module is most plausibly construed as part of vision or visuomotor control (Pylyshyn, 1999). Perhaps surprisingly, a similar point applies to what is widely regarded as among the strongest candidates for Central Modularity: the theory of mind module (ToMM) hypothesized by Alan Leslie and his collaborators (Leslie et al., 2004). The ToMM hypothesis is not without problems (see Nicholls and Stich, 2002). But even if we put these to one side, the existence of ToMM would do little to strengthen the case for Central Modularity since in its most recent and plausible incarnations, ToMM is characterized as a relatively low-level device for selective attention; and moreover one which relies heavily on decidedly nonmodular executive systems – most notably the "selection processor" – in order to perform its characteristic function (Leslie et al., 2004). Thus what is widely viewed as a strong candidate for Central Modularity both fails to be a clear example of a central system, and also presupposes the existence of nonmodular systems.

4 Some Reasons to Doubt Massive Modularity

So, the experimental evidence for Central Modularity is rather less strong than one would expect if MM were true. But there are also familiar – and in my view very plausible – claims about human thought which collectively suggest the existence of paradigmatically nonmodular central systems. In particular:

Conceptual integration. We are capable of freely combining concepts across different subject matters or content domains (Carruthers, 2003). So, for example, it's not merely that I can think about colors, about numbers, about shapes, about food, and so on. Rather I can have thoughts that concern all these things – e.g. that I had two roughly round, red steaks for lunch.

Generality of thought capacities. Not only can we freely combine concepts, we can also deploy the resulting representations in our theoretical and practical deliberations – to assess their truth or plausibility, but also to assess their impact on our plans and projects.

Inferential holism. Given surrounding conditions – especially, background beliefs – the *relevance* of a representation to the theoretical or practical tasks in which one engages can change dramatically.[11] Indeed, it would seem that given appropriate background assumptions, almost any belief can be relevant to the task in which one engages (Copeland, 1993).

Task range. We are capable of performing an *exceedingly wide* – perhaps unbounded – range of cognitive tasks of both a practical and theoretical nature (Newell, 1990).

Between task correlations. There is a huge amount of empirical evidence that performance on many central cognitive tasks *co-varies*. Thus psychometric studies appear to show that performance in any single task (e.g. arithmetic) is a reliable predictor of performance in a wide range of other cognitive tasks (e.g. spatial reasoning, deductive reasoning, spelling, reading, and so on) (Anderson, 1992).

The existence of general mental retardation. There are various organic disorders – most notably Down's syndrome – which produce across the board deficits in the performance of such highly correlated intellectual tasks (Anderson, 1998).

Taken individually, none of these phenomena mandates the existence of nonmodular central systems; and nor collectively do they preclude the existence of modular ones. In other words, the above are clearly compatible with what I called Weak Central

Modularity. Nevertheless, taken together, they do strongly suggest the existence of nonmodular – domain general and unencapsulated – mechanisms for thought and, hence, that Strong Central Modularity is false. This is because the assumption there are such mechanisms goes some way to help explain the above facts.

In contrast, if one rejects nonmodular mechanisms altogether, the prospects of accommodating these phenomena appear bleak; and to date, efforts to do so have been deeply unsatisfactory. Some strongly modular architectures – such as those embodied in Brook's animats – are precisely specified computational proposals but clearly inadequate to explaining the above phenomena (Kirsh, 1991; Bryson, 2000). In contrast, other modular proposals are so vague as to be almost wholly lacking in content. So, for example, it is common for advocates of MM to claim that the above kinds of phenomena can be largely explained by the fact that minds contain *so many* modules: "a network of subsystems that feed each other in criss-crossing but intelligible ways" (Pinker, 2005; see also Pinker, 1994; Buss, 2000). But these are not explanations so much as statements of the problem given a commitment to MM. The challenge is to sketch the *right sort* of plurality and "criss-crossing" between mechanisms – one that could plausibly exhibit the above sorts of phenomena. Finally, even those proposals that avoid vacuity or obvious inadequacy only seem to accommodate the above phenomena to the extent that they smuggle nonmodular mechanisms back in. So, for example, Carruthers has sketched what purports to be an MM account of practical and theoretical reasoning that seeks to accommodate the above sorts of phenomena (Carruthers, 2003, 2004). But the proposal fails since it posits various mechanisms – specifically, a "relevance module" and a practical reasoning system – which are domain general and, despite Carruthers' claims to the contrary, unencapsulated as well.[12] So, the prospects of accommodating the above phenomena without positing nonmodular mechanisms appear bleak; and in view of the lack of argument or evidence for MM, I'm inclined to think the effort of trying to do so is, in any case, wasted.

5 Whither Modularity?

The main burden of this chapter has been to argue that the case for an interesting and distinctive version of MM is not strong because there is little reason to suppose *central* processes are modular. In this section I conclude with some comments on a more radical view: that little or none of cognition – including peripheral systems for low-level perception – are modular in character. In my view, this claim goes too far since there are strong empirical grounds for positing a wide array of non-central modules. Efforts to draw grand conclusions about the irrelevance of modularity thus strike me as, at best, premature and, at worst, a serious distortion of our current, best picture of the cognitive mind.

5.1 Evidence for modularity

As with almost any scientific hypothesis, the empirical case for peripheral modularity is not so strong as to exclude all competitors. But what matters is the relative plausibility of competing hypotheses; and on balance I'm inclined to think the evidence – convergent evidence from neuroanatomy, psychophysics, computational modeling, developmental psychology, and so on – supports the existence of modules for many non-central cognitive processes.

The case is most plausible for low-level (or "early") visual processes – for example, those involved in the perception of color, shape from motion, and depth from disparity. In such cases one finds illusions strongly suggestive of mechanisms encapsulated from belief (Fodor, 1983; Pylyshyn, 1999). Moreover, in contrast to central processes, we possess quite detailed and plausible computational models of these perceptual processes: models which suggest the kinds of computations involved are highly specialized for very specific sorts of input and draw on very restricted kinds of information (Palmer, 1999). So, for example, the computations involved in computing depth from disparity are quite different from those involved in computing lightness; and the sorts of information relevant to one are irrelevant to the other. Finally, and again in contrast to central processes, a fair amount is known about the neural basis of such processes; and again what we know suggests that the modular picture is very plausible (Zeki and Bartels, 1998).

Nor is the case for modularity only plausible for low-level perceptual mechanisms. Though perhaps less compelling, there is still a good case to be made elsewhere, including: face recognition (McKone and Kanwisher, 2005), place-specific visual processing (Epstein et al., 1999), approximate arithmetic (Butterworth, 1999; Feigenson et al., 2004), and grammar-specific processing (van der Lely et al., 2004). In each such case, one finds a pattern of results strongly suggestive of modularity: evidence of dissociations, evidence for idiosyncratic styles of processing or distinctive representational formats, and ERP and imaging studies suggestive of distinctive neural responses. All this is, I suppose, good evidence for modularity. Not beyond dispute, but strong nonetheless.

Why, then, do so many reject the modularity of such processes? Since the empirical literature on such matters is enormous, it will not be possible to address much of what has been said here. But fortunately, Prinz (chapter 2) provides an excellent (and brief) presentation of the case against peripheral modularity: one that draws together many familiar strands of argument, but also exhibits many of the problems found in other critiques. In what follows, then, I focus primarily on Prinz's arguments, though many of my comments apply more broadly than this.

5.2 Arguments against domain specificity

According to Prinz, "alleged modules may have domain-specific components" – e.g. proprietary rules and representations – but they "don't seem to be proprietary throughout." Prinz's argument for this conclusion divides in two. First, he distinguishes between some different senses of "specificity" and "domain" and argues that for a mechanism to be interestingly domain specific it should be *exclusively* dedicated to processing some *relatively inclusive* subject matter. Next, Prinz argues that the main alleged examples of modules – vision, language, mathematics, folk physics etc. – are not domain specific in this sense. Neither part of the argument strikes me as convincing.

Characterizing domain specificity

I see no reason to assume that for a mechanism to be interestingly domain specific, its domain must be inclusive – as broad as language, vision or theory of mind, for example. First, the assumption fits poorly with how the notion of domain specificity in fact gets used in cognitive science. Many of the most prominent advocates of modularity – Fodor and Sperber, for example – explicitly deny that domains must be broad in this way (Fodor, 1983; Sperber, 1994); and many of the most carefully studied examples of domain-specific mechanisms – depth from disparity, for example – are not all that inclusive.

Second, contrary to what Prinz appears to suggest, the admittance of fine-grained domains need not trivialize claims about domain specificity. There is nothing trivial, for example, in claiming we possess mechanisms exclusively dedicated for edge detection, scene geometry, or approximate arithmetic. Of course, if one insisted on treating just *any* entity for which some subject matter could be specified as a *bona fide* domain-specific mechanism, then triviality threatens. Therein lies the route to mechanisms for, say, reasoning about camels in London on Tuesdays, or planning wheat-free picnics in 2005. But this is not a problem with a fine-grained notion of domain specificity, or indeed with the notion of domain specificity at all. Rather, it's a problem that arises when one adopts a silly attitude towards the individuation of cognitive mechanisms. The main problem with a distinct "mechanism" for reasoning exclusively about camels in London on Tuesdays is not that its domain is too fine-grained, but that *there is no such mechanism*. To be sure, there are genuine issues about what the criteria for individuating cognitive mechanisms ought to be. But the present discussion highlights that silliness ensues if one individuates by domain alone. My suggestion: Don't do it.

Evidence against domain specificity?

Even laying aside the above concern, Prinz's argument still fails since the main kind of evidence he cites is largely *irrelevant* to the domain specificity of cognitive mechanisms. Prinz objects to many alleged cases of domain specificity on the grounds that the processes involved in one kind of task share cognitive/neural resources with processes of other sorts. So, for example: language processing shares resources with nonlinguistic pattern recognition; mind-reading shares resources with language processing and also uses working memory; arithmetic shares resources with language and spatial processing, and so on. On the basis of such considerations, Prinz concludes that there are unlikely to be domain-specific mechanisms of the sort that modularists posit.

What are we to make of this argument? The logic is rather murky. In particular, it is very unclear how one gets from claims about resource sharing to conclusions about the absence of domain-specific mechanisms. That said, the general idea appears to be that since domain-specific mechanisms are *exclusively dedicated* to processing a given domain of information, there cannot be domain-specific devices for processes that depend on resources recruited for many *different* domains. On reflection, however, this argument turns on a shift between two different senses in which a domain specific mechanism might be exclusively dedicated:

The mechanism only processes information in a given domain, D;
The mechanism only processes – and, moreover, is the *only* mechanism that processes – information in D.

The former is one to which advocates of domain-specific modules are plausibly committed; but it is not one on which resource sharing counts as evidence against domain specificity. Thus the existence of mechanisms which are exclusively dedicated in this sense is perfectly compatible with the very familiar position – one endorsed by many modularists – on which cognitive processes depend on both domain-specific *and* domain-general mechanisms. For example: Leslie and his collaborators claim mindreading depends on both a domain-specific theory of mind module and domain-general executive systems (Leslie, 2003). In contrast, the second notion of exclusive dedication does license inferences from resource sharing to a failure of domain specificity and, hence, precludes precisely such cases. But the problem for Prinz is that advocates of domain specific mechanisms are clearly not committed to mechanisms which are exclusively dedicated in this way.

5.3 Evidence against informational encapsulation

Let's turn to the case against encapsulated mechanisms. Prinz's strategy is to try to undermine what is widely regarded as the most plausible case for encapsulation: perceptual illusions which modularists claim result from the operation of low-level perceptual mechanisms. But *pace* Fodor, Pylyshyn, and others, Prinz maintains illusions are not good evidence for encapsulation because there is a competing hypothesis – what he calls the "trumping hypothesis" – which explains the existence of illusions and, moreover, does a better job of accommodating evidence of top-down effects on perception. Again, I am unconvinced.

First, the trumping hypothesis itself is hard to make sense of. The rough idea is simple enough: though belief is trumped by perception when the two are in conflict, it can influence perception when no conflict exists. But how exactly is this to occur? On the most natural interpretation, what's required is some kind of consistency check between belief (e.g. that the lines in the Müller-Lyer illusion are of equal length) and a representation produced by some perceptual process (e.g. that the lines are of different length). But if the trumping hypothesis were correct, such a checking process would *presuppose* the existence of encapsulated perceptual mechanisms. After all, for a consistency check to occur at all, there must be a *perceptual* representation – i.e. the output of some perceptual device – that can be checked against belief. And since, according to the trumping hypothesis, beliefs only influence perceptual processing when no conflict exists, it cannot be that beliefs are implicated in producing the output of *this* perceptual device.

In any case, the data cited by Prinz do not merit the rejection of encapsulated low-level perceptual mechanisms. Following a long tradition, Prinz argues that top-down effects on perception are incompatible with encapsulation. But the argument turns on the assumption that a "truly encapsulated system would be insulated from *any* external influence"; and this is simply false. On the contrary, advocates of encapsulation agree with their opponents that there are top-down influences on perception.

What's at issue is the character and extent of this influence. Specifically, what advocates of encapsulated early perception are most concerned to reject is a picture – widely associated with Bruner's New Look psychology – on which early perceptual mechanisms have something approximating an unlimited access to one's beliefs and desires in the course of their online processing (Fodor, 1983, p. 60; Pylyshyn, 1999). But this rejection is wholly compatible with many sorts of external cognitive influence on perception, including:

- Shifts in loci of focal attention brought about by clues, instructions, or preferences about where to look;
- Top-down processing *within* a perceptual modality;
- Cross-talk between perceptual systems;
- Diachronic or developmental effects in which one's beliefs and goals influence the development of perceptual systems – e.g. via training effects;
- Beliefs and desires influencing late perceptual processes, such as perceptual categorization.[13]

As far as I can tell, all the putative objections to encapsulation Prinz cites fall into one or other of these categories. For example, it is very plausible that the influence of verbal cues and decisions on our experience of ambiguous figures consists in the production of shifts in the locus of focal attention (Peterson and Gibson, 1991; Pylyshyn, 1999). Similarly, the role of expectations – e.g. in producing non-veridical experiences – is plausibly viewed as an influence on late perceptual processes. And so on. In view of this, I see no reason here to deny the modularity of early perception.

Conclusion

We started by clarifying the sort of view that advocates of MM seek to defend. We then saw that the main theoretical arguments for views of this sort fail to provide reason to prefer them over other competing proposals, such as those on which much of cognition depends on nonmodular mechanisms with access to bodies of specialized information. Next, we saw that the available experimental case for MM is not strong because there is little evidence for the existence of modular central systems. Moreover, we saw that there is some reason to reject a strong MM which claims that all central systems are domain specific and/or encapsulated. Finally, we saw that while MM is not well supported by the arguments and evidence, it would be wrong to maintain that minds are not modular to any interesting degree since there are good reasons to suppose that more peripheral regions of the mind – especially for low-level perception – are modular in character. Where does this leave us? If our assessment of the evidence is correct, then the most plausible position to adopt is one that takes a middle way between those, such as Carruthers, who endorse a thoroughgoing massive modularity and those, such as Prinz, who reject modularity altogether. The situation is, in other words, much as Fodor advocated over two decades ago (Fodor, 1983).

Acknowledgments

I would like to thank Guy Longworth, David Papineau, Gabe Segal, Rob Stainton, and Mark Textor for helpful comments on earlier drafts of this chapter. I would also like to thank Brian Scholl for helpful discussion of the material in section 5.

Notes

1 There is a familiar notion of modularity, sometimes called Chomskian modules, on which modules are not mechanisms but systems of mental representation – bodies of mentally represented knowledge or information – such as a grammar or a theory (Segal, 1996; Samuels 2000; Fodor 2000). Paradigmatically, such structures are truth-evaluable in that it makes sense to ask of the representations if they are true or false. Moreover, they are often assumed to be innate and/or subject to informational constraints (e.g. inaccessible to consciousness). Although Chomskian modules are an important sort of cognitive structure, they are not the ones most relevant to the sort of position advocated by massive modularists. This is because advocates of MM appear to assume that modules are a species of cognitive *mechanism*. I say this for two main reasons. First, it is very plausible on the basis of their explicit comments about what modules are. Dan Sperber, for example, characterizes modules as a species of "autonomous mental mechanism" (Sperber, 2002; Sperber and Hirschfeld, 2004), while Cosmides and Tooby characterize them as "functionally dedicated computers"; all of which indicates a conception of modules as cognitive mechanisms. Moreover, if modules are construed as Chomskian modules, it's hard to make sense of what advocates of MM say about the relationship between their position and other views. So, for example, MM is routinely (and correctly) presented as being in opposition to views on which central cognition depends on the activity of one (or a few) general-purpose computers (Cosmides and Tooby, 1992; Carruthers, chapter 1). But if modules are construed as systems of mental representations, it's unclear why any opposition should exist. After all, the claim that central cognition depends on, say, a single universal Turing machine is wholly compatible with the existence of a great many Chomskian modules – bodies of information deployed by the mechanism. This is, however, surely *not* the sort of view advocates of MM seek to defend. Indeed Carruthers (chapter 1) says as much when explicitly contrasting this sort of "informational modularity" with the "computational modularity" he seeks to defend.

2 The following discussion is by no means exhaustive. For more detailed discussions of different notions of modularity see Segal, 1996, and Samuels, 2000.

3 This list is sometimes construed as a *definition* of modularity (Karmiloff-Smith, 1992). But Fodor rejects this in favor of the idea that modularity admits of degree and that cognitive mechanisms are modular to the extent that they possess all or most of the features on this list to some interested degree. (Fodor, 1983, p. 37; Coltheart, 1999).

4 Incidentally, not even Fodor adopts it in his recent discussions of MM (Fodor, 2000, 2005).

5 To claim that a property of a cognitive mechanism is *architecturally* imposed, minimally implies the following. First, they are relatively enduring characteristics of the device. Second, they are not mere products of *performance* factors, such as fatigue or lapses in attention. Finally, they are supposed to be *cognitively impenetrable* (Pylyshyn, 1984). To a first approximation: they are not properties of the mechanism that can be changed as a result of alterations in the beliefs, goals and other representational states of the organism.

6 It should go without saying – though I'll say it anyway – that the notion of domain specificity *admits of degree* and that researchers who use the notion are interested in whether we possess mechanisms that are domain specific to some *interesting* extent. The same points also apply to the notion of informational encapsulation.

7 Indeed this is more-or-less guaranteed by the widespread assumption that the functional decomposition of a "large" system will typically have many levels of aggregation (Simon, 1962). I return to this point in section 2.1.

8 A similar point applies to the sort of radical connectionism on which the mind is characterized as one huge undifferentiated neural network. This is often – and rightly – seen as the antithesis of MM (Pinker, 1997); and yet it is committed to a vast plurality of mechanisms. After all, each node in a neural network is a mechanism; and on any version of the connectionist story, there will a great many such nodes.

9 There are other less plausible arguments for MM, which due to space limitations I will not consider here. For further discussion of other arguments for MM see Tooby and Cosmides, 1992; Sperber, 1994; and Samuels, 2000.

10 See Prinz (chapter 2, IS THE MIND REALLY MODULAR?) for complementary criticisms of the experimental case for MM.

11 The same would also appear to be true of many of a representation's epistemic properties – e.g. simplicity and conservativism.

12 They are "wide-scope encapsulated." But as already noted, this is not encapsulation as ordinarily construed.

13 Indeed, given that encapsulation admits of degree, an input system might be encapsulated to some interesting degree and yet still have access to some beliefs and goals in the course of its operations.

References

Anderson, J. R. (1990). *The Adaptive Character of Thought.* Hillsdale, NJ: Erlbaum.

Anderson, M. (1992). *Intelligence and Development: A Cognitive Theory.* Oxford: Blackwell.

— (1998). Mental retardation, general intelligence and modularity. *Learning and Individual Differences,* 10/3, 159–78.

Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavioral and Brain Sciences*, 21, 547–609.

— (2001). The case for modularity: Sin or salvation? *Evolution and Cognition*, 7, 46–55.

Baron-Cohen, S. (1995). *Mindblindness: An Essay on Autism and Theory of Mind.* Cambridge, MA: MIT Press.

Bryson, J. J. (2000). Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, 12/2, 165–90.

Buss, D. (2000). *Evolutionary Psychology.* New York: Allyn & Bacon.

Buss, D. and Duntley, J. (2005). On the plausibility of adaptations for homicide. In P. Carruthers, S. Laurence, and S. Stich (eds.), *The Innate Mind: Structure and Contents.* New York: Oxford University Press.

Butterworth, B. (1999). *The Mathematical Brain.* London: Macmillan.

Carey, S. (1995). On the origins of causal understanding. In S. Sperber, D. Premack, and A. Premack (eds.), *Causal Cognition.* Oxford: Clarendon Press.

Carruthers, P. (2003). On Fodor's problem. *Mind and Language*, 18, 502–23.

— (2004). Practical reasoning in a modular mind. *Mind and Language*, 19, 259–78.

Carston, R. (1996). The architecture of the mind: modularity and modularization. In D. Green (ed.), *Cognitive Science.* Oxford: Blackwell.

Chomsky, N. (1980). *Rules and Representations*. New York: Columbia University Press.
Coltheart, M. (1999). Modularity and cognition. *Trends in Cognitive Sciences*, 3/3, 115–20.
Copeland, J. (1993). *Artificial Intelligence: A Philosophical Introduction*. Oxford: Blackwell.
Cosmides, L. (1989). The logic of social exchange: has natural selection shaped how humans reason? Studies with Wason Selection Task. *Cognition*, 31, 187–276.
Cosmides, L. and Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, and J. Tooby (eds.), *The Adapted Mind*. New York: Oxford University Press.
—— (1994). Origins of domain specificity: the evolution of functional organization. In L. Hirschfeld and S. Gelman (eds.), *Mapping the Mind*. Cambridge: Cambridge University Press.
Epstein, R., Harris, A., Stanley, D., and Kanwisher, N. (1999). The parahippocampal place area: recognition, navigation, or encoding? *Neuron*, 23, 115–25.
Feigenson, L., Dehaene, S., and Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences,* 8/7, 307–14.
Fodor, J. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
—— (2000). *The Mind doesn't Work that Way*. Cambridge, MA: MIT Press.
—— (2005). Reply to Steven Pinker: "So how does the mind work?" *Mind and Language*, 20/1, 25–32.
Gallistel, R. (1990). *The Organization of Learning*. Cambridge, MA: MIT Press.
—— (2000). The replacement of general-purpose learning models with adaptively specialized learning modules. In M. Gazzaniga (ed.), *The New Cognitive Neurosciences* (2nd edn.). Cambridge, MA: MIT Press.
Garfield, J., ed. (1987). *Modularity in Knowledge Representation and Natural-Language Understanding*. Cambridge, MA: MIT Press.
Gigerenzer, G. (1994). Why the distinction between single-event probabilities and frequencies is important for psychology (and vice versa). In G. Wright and P. Ayton (eds.), *Subjective Probability*. New York: Wiley.
Gigerenzer, G. and Hug, K. (1992). Reasoning about social contracts: Cheating and perspective change. *Cognition*, 43, 127–71.
Gopnik, A. and Meltzoff, A. (1997). *Words, Thoughts and Theories*. Cambridge, MA: MIT Press.
Hermer-Vazquez, L., Spelke, E., and Katsnelson, A. (1999). Sources of flexibility in human cognition: Dual-task studies of space and language. *Cognitive Psychology* 39, 3–36.
Jackendoff, R. (1992). Is there a faculty of social cognition? In R. Jackendoff, *Languages of the Mind*. Cambridge, MA: MIT Press.
Karmiloff-Smith, A. (1992). *Beyond Modularity*. Cambridge, MA: MIT Press.
Kirsh, D. (1991). Today the earwig, tomorrow man. *Artificial Intelligence,* 47, 161–84.
Leslie, A. M., Friedman, O., and German, T. P. (2004). Core mechanisms in "theory of mind." *Trends in Cognitive Sciences,* 8, 528–33.
McKone, E. and Kanwisher, N. (2005). Does the human brain process objects of expertise like faces? A review of the evidence. In S. Dehaene, J. R. Duhamel, M. Hauser, and G. Rizzolatti (eds.), *From Monkey Brain to Human Brain*. Cambridge, MA: MIT Press.
Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
Nicholls, S. and Stich, S. (2002). *Mindreading*. New York: Oxford University Press.
Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.
Peterson, M. A. and Gibson, B. S. (1991). Directing spatial attention within an object: altering the functional equivalence of shape descriptions. *Journal of Experimental Psychology: Human Perception and Performance* 17, 170–82.
Pinker, S. (1994). *The Language Instinct*. New York: William Morrow.
—— (1997). *How the Mind Works*. New York: Norton.

— (2005). So how does the mind work? *Mind and Language*, 20/1, 1–24.
Pylyshyn, Z. (1984). *Computation and Cognition*. Cambridge, MA: MIT Press.
— (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–423.
Samuels, R. (1998). Evolutionary psychology and the massive modularity hypothesis. *British Journal for the Philosophy of Science*, 49, 575–602.
— (2000). Massively modular minds: Evolutionary psychology and cognitive architecture. In P. Carruthers and A. Chamberlain (eds.), *Evolution and the Human Mind: Modularity, Language and Meta-Cognition*. Cambridge: Cambridge University Press.
— (2005). The complexity of cognition: tractability arguments for massive modularity. In P. Carruthers, S. Laurence, and S. Stich (eds.), *The Innate Mind: Structure and Contents*. New York: Oxford University Press.
Segal, G. (1996). The modularity of theory of mind. In P. Carruthers and P. Smith (eds.), *Theories of Theory of Mind*. Cambridge: Cambridge University Press.
Simon, H. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106, 467–82.
Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In Hirschfeld, L. and Gelman, S. (eds.), *Mapping the Mind*. Cambridge: Cambridge University Press.
— (1996). *Explaining Culture: A Naturalistic Approach*. Oxford: Blackwell.
— (2002). In defense of massive modularity. In I. Dupoux (ed.), *Language, Brain and Cognitive Development*. Cambridge, MA: MIT Press.
— (2005). Massive modularity and the first principle of relevance. In P. Carruthers, S. Laurence, and S. Stich (eds.), *The Innate Mind: Structure and Contents*. Oxford: Oxford University Press.
Sperber, D. and Girotto, V. (2003). Does the selection task detect cheater detection? In K. Sterelny and J. Fitness (eds.), *From Mating to Mentality: Evaluating Evolutionary Psychology*. New York: Psychology Press.
Sperber, D. and Hirschfeld, L. (2004). The cognitive foundations of cultural stability and diversity. *Trends in Cognitive Science*, 8/1, 40–6.
Sperber, D., Cara, F., and Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, 52, 3–39.
Tooby, J. and Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides, and J. Tooby (eds.), *The Adapted Mind*. New York: Oxford University Press.
Van der Lely, H. K. J., Rosen, S., and Adlard, A. (2004). Grammatical language impairment and the specificity of cognitive domains: Relations between auditory and language abilities. *Cognition*, 94, 167–83.
Varley, R. (2002). Science without grammar: Scientific reasoning in severe agrammatic aphasia. In P. Carruthers, S. Stich, and M. Siegal (eds.), *The Cognitive Basis of Science*. Cambridge: Cambridge University Press.
Varley R. A., Siegal, M., and Want, S. (2001). Severe impairment in grammar does not preclude theory of mind. *Neurocase*, 7, 489–93.
Zeki, S. and Bartels, A. (1998). The autonomy of the visual systems and the modularity of conscious vision, *Philosophical Transactions of the Royal Society London*, 353, 1911–14.