# Concepts and the modularity of thought

Daniel A. Weiskopf[†]

ABSTRACT: Having concepts is a distinctive sort of cognitive capacity. One thing that conceptual thought requires is obeying the Generality Constraint: concepts ought to be freely recombinable with other concepts to form novel thoughts, independent of what they are concepts of. Having concepts, then, constrains cognitive architecture in interesting ways. In recent years, spurred on by the rise of evolutionary psychology, massively modular models of the mind have gained prominence. I argue that these architectures are incapable of satisfying the Generality Constraint, and hence incapable of underpinning conceptual thought. I develop this argument with respect to two well-articulated proposals, due to Dan Sperber and Peter Carruthers. Neither proposal gives us a satisfactory explanation of Generality within the confines of a genuinely modular architecture. Massively modular minds may display considerable behavioral and cognitive flexibility, but not humanlike conceptualized thought.

Since the publication of Fodor's *The Modularity of Mind* (1983), discussions of cognitive architecture have often been couched in terms of whether the mind is to any interesting degree modular. While Fodor argued that only input and output systems (perception-action systems plus the language faculty) were modular, in recent years, *massively modular* theories of the mind have gained in popularity, due in large part to the rise of evolutionary psychology (Buss, 2008; Pinker, 1999; Tooby & Cosmides, 1992). According to massive modularists, the mind is modular not just at the periphery, but through and through. In its strongest form, the massive modularity thesis holds that there is *no* single central system where processes like belief fixation, cross-domain reasoning, informational integration, planning, and decision-making take place. Rather, there are only domain-specific systems for processing information about one or another circumscribed sort of information; e.g., systems for reasoning about other minds but not about physical objects, or for deciding who to mate with but not what sort of car to buy.

[†] Department of Philosophy, Georgia State University, P.O. Box 4089, Atlanta, GA 30302; Email: dweiskopf@gsu.edu

My goals in this paper are twofold. First, I aim to sketch a view of concepts that shows how possessing them places constraints on the sort of cognitive architecture a creature can have. Second, I will turn to massively modular accounts of the mind and see whether they can meet these strictures on conceptual thought. One frequent line of argument against massive modularity takes a defensive tack, attempting to show merely that the arguments in favor of such an architecture are not decisive. Here I take a stronger line. I will argue that, in attempting to account for the phenomena centrally associated with fully conceptualized thought, massively modular architectures invariably end up positing non-modular components. Thought in the robust, fully conceptual sense is incompatible with the sorts of representational and architectural assumptions characteristic of massive modularity.

## 1. Concepts, Generality, and isotropy

To a first approximation, *thoughts* are states of a creature that play a causal role in mediating between the creature's perceptions of the world and the actions that it takes. This conception of thoughts as centrally situated internal control states reflects Armstrong's famous claim that a mental state is "a state of the person apt for bringing about a certain sort of behavior" (Armstrong, 1968, 82). However, behavior is normally produced by the interaction of many different sorts of inner states. Both perceptual and motor capacities, for instance, are implemented by complex systems realized in a creature's nervous system. Given how sophisticated we now understand these systems to be, what distinctive role is there for thoughts, as distinguished from states involved in perception and action, to play in the causation of behavior? More narrowly, what is required for a creature to have *conceptualized* thoughts?

Many conditions on having concepts have been proposed (Bernal, 2005; Camp, 2009; Peacocke, 1992; Prinz, 2002). Here I wish to focus on a plausible necessary condition on the possession of fully conceptualized thought of the sort that humans, at least, are able to enjoy. This condition is the Generality Constraint proposed by Gareth Evans. In Evans' words, the Constraint says:

> (Generality) "[I]f a subject can be credited with the thought that *a* is *F*, then he must have the conceptual resources for entertaining the thought that *a* is *G*, for every property of being *G* of which he has a conception." (Evans, 1982, 104)

By 'thought' here, Evans means only conceptualized thought; other information processing systems may not be subject to Generality.

The Generality Constraint embodies the idea that creatures capable of conceptual thought do not have 'gappy' mental lives. To have a thought with the conceptual structure A IS F (where A is an individual concept and F a general concept) requires being able to entertain thoughts with similar structure, but containing any other general concept that the creature possesses.[1] With some possible exceptions (Camp, 2004), human conceptual thought also appears to obey the Constraint. Where we can think that a certain object is F, and we also have the concept G, we can think that that object is G (see Peacocke, 1992, 42-51).

The Generality Constraint is fundamentally a closure principle for the conceptual system. It requires that any concepts falling within the same broad class—e.g., concepts of individuals, properties, relations, events, kinds, etc—must be able to enter into the same range of

[1] Evans, of course, does not think—or at least does not explicitly commit himself to the idea—that the structure of thought is a matter of their being mental representations of a certain kind. Hence his warnings against reading the idea of a language of thought into the description of Generality. Rather, Evans thinks of concepts as being abilities to entertain certain contents, and thoughts as being the 'products' of the joint exercise of these abilities. I will continue to read the Generality Constraint in a way that treats it as a constraint on the sorts of representation-combining capacities a creature must have in order to possess concepts, however.

combinations as all other concepts of that class. So if I can think that Sophie is a cat, I can also think that she is an elephant, a tapir, a Hilbert space, a prime number, a hatrack, and so on. And if I can think this of Sophie, I must also be able to think it of any other individuals for whom I have concepts. Nothing could be a concept unless it was capable of entering into this kind of system of relations, and nothing could count as possessing a conceptual system unless it had a system of representations that were organized in such a way.

One powerful rationale for imposing the Generality Constraint is its essential role in facilitating inference. This link is made explicitly by Susan Hurley (1997, 207):

> If the information that a given object has a certain property is conceptualized, it is not context-bound. It has a structure that enables the person in principle to generalize systematically, to decompose and recombine, to make other applications of the same concepts. She can entertain and act not just on this information in this context, or closely related contexts, but also in principle on the information that the same object has a different property or the information that a different object has the same property. And she can quantify and make inferences that depend on such decompositional structure and context-freedom.

Consider a creature that can think about Fs, Gs, and Hs. If its thoughts obey Generality, and it possesses some basic logical operations, it can think the following sequence: FS ARE GS; GS ARE HS; so FS ARE HS. Transitive categorical inference like this is only possible in general if a creature can entertain the full range of combinations of thoughts involving its predicative mental terms. A creature whose thoughts were restricted in some way—e.g., it could only think about Fs in relation to Gs, but not Hs—would be incapable of drawing indefinitely many potentially useful inferences. And one benefit of having a conceptual system is precisely to enable informational integration in this way.

To anticipate the main argument of this paper, I note that there are interesting connections between the Generality Constraint and the property of cognitive systems called ‘isotropy’ by Fodor (1983). Isotropy claims that “the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything the scientist knows is, in principle, relevant to determining what else he ought to believe” (108). If hypothesis confirmation is typically isotropic—as I assume that it is—then it is necessary to draw inferences from (potentially) any domain of knowledge. And this will involve engaging in chains of reasoning such as the one sketched abstractly above, since at some point we will have to make use of an inferential ‘bridge’ that has a foot in at least two separate domains.

So a system’s being isotropic strongly suggests that its representations obey something like Generality. It is difficult to see how a system might be arranged in such a way as to make potentially any piece of information inferentially available for combination with any other piece without something like the Generality Constraint in place. For isotropy requires being able to draw inferential links from one domain to another, and Generality ensures that for any concepts, it will be possible to think a sequence of transitive inferences that starts in one domain and ends up in another. So long as the appropriate inferential processes are in place, Generality is clearly sufficient for isotropy. Moreover, it seems necessary as well, since without it one can’t in general entertain the cross-domain thoughts that form the premises and conclusions of these inferences.

This also provides a response to the objection that Generality should not be understood as a necessary condition on concept possession. Human-like conceptualized thought looks as if it were isotropic. This property is explained by conceptualized thought obeying Generality. Isotropy and Generality mark out a real difference between kinds of cognitive systems. Even if one wants to deny that this distinction is essential to something's having concepts, what matters most for the present argument is the distinction, not precisely what capacity it marks. So even if one adopts a different mark of the conceptual, the arguments given here against massive modularity still have force.

## 2. Modularity and massive modularity

Since Fodor brought the term to prominence, many different notions of modularity have developed. I will not review all of them here, but will focus on distinguishing a few central kinds. First there is the classical notion of a *Fodor module* as spelled out in his (1983). Fodor modules are computational devices that are domain-specific, informationally encapsulated, shallow in their outputs, mandatory in operation, fast, largely innate, associated with a typical developmental trajectory, and neurally localized. It has long been noted that this conception is highly restrictive. In one sense this is a benefit: these properties form a potentially interesting cluster. Fodor thought that input/output systems (plus language) were modular for principled reasons having to do with the kind of task they perform, which in turn justified the theoretical utility of the concept of a module. On the other hand, if there are no, or very few, Fodor modules, the concept may not be of much interest or use. Many theorists have responded to this latter possibility by loosening the notion in various ways.

At the other extreme from Fodor modularity is a notion we can call *minimal modularity* (after Cowie, 2008). A minimal module is just a distinct mental subsystem that carries out some proprietary function. To say that it is a distinct system is to say that its processing can in principle be separately impaired or affected, although in practice this might be difficult to arrange. That is, while it may be embedded in a larger subsystem and may have other subsystems as components, aside from these mereological relationships it is an independent component carrying out an independent task. This also comports with Shallice's (1988, 24) notion of a module in neuropsychology.

While all Fodor modules are minimal modules, minimal modules need have virtually none of the properties of Fodor modules. They are quite unconstrained with respect to information flow, processing speed, relative permeability, etc. What they have in common is that they are specialized computational devices that perform a certain task. The attraction of the notion is that it is hard to deny that the mind does contain minimal modules. Denying this would be to deny that minds possess any sort of stable functional decomposition at all, and this seems implausible. (Perhaps no one has held this view apart from behaviorists and certain kinds of pure associationists; see Fodor, 2000, 56-7 and Wilson, 2008, 274.) To see the latitude the notion of minimal modularity allows, note that if functions may be defined in extremely general terns, then the classical belief box might well count as a minimal module, since it was hypothesized to be a separate cognitive system dedicated to the function of forming and evaluating beliefs. Separate memory stores such as primary and secondary memory, or the various kinds of working memory now posited, would also count as modular on this criterion. So this notion is one that should command widespread support.

What is wanted is something between the overly restrictive and overly permissive notions sketched so far. In recent discussions, two properties have emerged as central to the notion of modularity. These are informational encapsulation and domain specificity. Fodor now (2000, 58) holds informational encapsulation to be the most important property that determines whether a system is modular or not. Encapsulation just means that a system can or does access less than all of the information stored in the mind in the course of carrying out its functions. In particular, it must mean that a system can or does access *vastly* less than all of the information in the mind, since having access to greater quantities of information leads (the argument goes) to computational intractability and ultimately to the Frame Problem. The point of modules as contrasted with more central systems was, in Fodor's (1983) conception at least, that modules are systems for which the Frame Problem does not arise. Informational encapsulation purports to be one way of achieving this aim: if there is less accessible information, problems of search and relevancy do not arise in as pointed a fashion.

Finally, modules are domain specific. For convenience, I will distinguish two forms of domain specificity. The first is *input* domain specificity. On this view, a module's domain is whatever is represented by the inputs that activate it. Each modular subsystem is presumed to be sensitive to only some inputs and not others. So a face-recognition module might accept a subset of visual images as its input, and produce as output labels indicating whether or not they are recognized as particular individuals' faces. This module's domain would be the range of perceptual inputs that trigger it, which presumably does not include every perceptual input the organism as a whole can experience, let alone most of the other representations in its cognitive economy. This view of domains is endorsed by Fodor (2000) and Carruthers (2006a).

The second notion I will call *functional* domain specificity, which is the claim that a module's domain is determined by the kind of computation it is designed to carry out. So a cheater detection module might accept an enormous range of possible inputs, from various sensory modalities, including inputs that come in the form of complex, linguistically presented scenarios. These inputs might be comparatively heterogeneous, but the module's domain need not be, since its function is to detect individuals who cheat in social exchanges. These individuals in these contexts comprise the module's domain, even if they are not part of the set of inputs that turn the module on. The range of objects and situations that the system has the function of dealing with constitute its domain. This notion seems to be endorsed by Barrett & Kurzban (2006, 634), Boyer & Barrett (2005), and Cosmides & Tooby (1994).

Henceforth in talking about modules I will mean mental systems that have a dedicated function, are informationally encapsulated, and are domain-specific (in one of the two senses listed). This notion seems to capture much of the core of the idea as it is used today, and avoids being either overly restrictive (as Fodor modularity is) or overly permissive (as minimal modularity is).

With this notion of a module in mind, what does it mean to talk about *massive* modularity? The thesis of massive modularity is actually best thought of as a cluster of related views, rather than a single unified doctrine (Samuels, 2000; Segal, 1996). What these views are commonly committed to, however, is the claim that the mind is *largely or entirely* composed of modules.[2] What they are also generally committed to is the idea that these mechanisms are informationally encapsulated: they have access to a highly restricted subset of the total information stored in the mind. What they are opposed to is the notion that there exists a domain-general system for reasoning about arbitrary content domains, ranging over a potentially unbounded body of information, using a common set of computational mechanisms.

[2] Even massive modularists typically allow for such non-modular capacities or resources such as attention or a kind of global memory workspace. These are conceived of as being available to the inputs and outputs of all modular systems. But these are typically the only exceptions. Here I focus on massively modular systems that lack any dedicated domain-general subsystems.

On the latter view, there exist 'central systems' in the mind where all of a creature's beliefs and desires are stored, processed, and made available for the control of behavior. On a massively modular architecture, however, there is no such central functional space. Rather, there are a host of more restricted systems that deal with reasoning about, say, social exchanges, mindreading, the physics of inanimate bodies, path finding and geographic location, mate selection, language, numerical and statistical reasoning, etc.

On a massively modular architecture, the flow of information and control operates as follows.[3] States of the environment cause activation of a creature's perceptual systems, which produce an initial representation of its surroundings. These representations are then passed to the next level of modular subsystems, which carry out their own proprietary computations on those inputs. For instance, a representation of a moving object might activate a folk physics module, which produces predictions about its likely course. Or it might activate a folk biological module, which attempts to identify it as an animate or inanimate object by the sort of trajectory it follows. Representation of a certain kind of visual array might activate a facial identification or conspecific detection module. A module mandatorily initiates processing of an input when it matches its characteristic triggering conditions. Output from these modules is then passed to the next layer, as well as laterally within layers. The flow of information among modules may be complex, but eventually information about the creature's surroundings reaches a suite of domain-specific desire-production mechanisms and practical reasoning systems, which generate specific desires and motor plans appropriate to the situation—e.g., if an available object is represented as being FOOD, this might activate a desire to eat it in a module dedicated to food-seeking. If this

---

[3] Barrett (2005) lays out in great detail the principles that might guide massively modular design.

food-seeking desire is the strongest one active, then the creature will in fact try to eat the alleged foodstuff.

But systems that contain numerous *central* modules—domain-specific belief/desire processing systems—seem incompatible with the Generality Constraint. According to the Constraint, a genuine concept possessor must be able, for any conceptualized thought it can entertain, to entertain a thought similar in structure, but containing a different predicative or singular concept than the original thought. This means not only that one must be able to apply any predicative concept to any singular concept, but also (given minimal logical apparatus) that one must be able to entertain logically complex predicative concepts that freely combine properties drawn from any domain that a thinker is capable of reasoning about.

So if I can think that this dessert is edible, and I possess the concepts EXPENSIVE and LIKELY TO MAKE HER HAPPY, then I can think (individually) that this dessert is expensive and that this dessert is likely to make her happy; but I can also think that this dessert is edible, expensive, and likely to make her happy. If I cannot form this latter, logically complex thought then either I cannot genuinely entertain the initial thought, or I do not genuinely possess any of the latter concepts, or both. Failure of one's thoughts to be closed under this sort of substitution is an indication that these thoughts are not properly thought of as conceptualized.

But this poses a problem for the massive modularist. How domains are individuated is not a simple question to answer, but *prima facie* this toy case appears to involve concepts from the folk domains of gastronomy, economics, and psychology. If these distinct domains are *solely* processed by separate, encapsulated mechanisms, then thoughts involving EDIBLE are processed by the gastronomy module, those involving EXPENSIVE are processed by the economics module, etc., and there is no further system that operates over complex representations drawn from arbitrary domains. Indeed, there is no system that could even construct such thoughts. If this is the case, then the Generality Constraint is violated. And this in turn entails that the system does not genuinely possess concepts, given that the constraint is a necessary condition on having a conceptual system at all.[4]

So, on the one hand a massively modular mind exclusively processes information about different domains in separate subsystems that cannot generally 'see' each others' contents and processes. But on the other hand, humanlike conceptualized thought necessitates having the capacity to construct thoughts that arbitrarily cross domain boundaries. It looks as if either the massive modularity hypothesis or the Generality Constraint must go. However, several theorists have proposed ways that we can, in effect, keep both. I now turn to assessing these proposals.

## 3. Against central modularity: Sperber and Carruthers

### 3.1. Sperber and the metarepresentation module

Sperber, in a series of papers, has articulated a massively modular view of cognition (Sperber, 1994, 2002, 2005). While his picture is complex, what is important for present purposes is that he is committed to the central modularity thesis: indeed, he holds not just that certain 'core domains' (e.g., living kinds) are modular, but also that individual concepts (e.g., DOG, CAT, GOLDFISH) are best thought of as 'micromodules'. As he puts it: "Concepts are domain-specific (obviously), they have a proprietary data-basis (the encyclopedic information

[4] Why not say that the creature possesses many separate systems of concepts, corresponding to the many distinct domain-specific 'belief fixation' mechanisms? Recall that the Generality Constraint operates as a kind of closure condition on concepthood. So if a creature has any concepts at all, it must have some system where information from all of these distinct subsystems is capable of being brought together.

filed under the concept), and they are autonomous computational devices (they work, I will argue, on representations in which the right concept occurs, just as digestive enzymes work on food in which the right molecule occurs)" (Sperber, 1994, 48).[5]

At the same time, he acknowledges Fodor's point that "[t]he conceptual level is the level at which information from different input modules, each presumably linked to some sensory modality, gets integrated into a modality-independent medium" (Sperber, 1994, 40).[6] To borrow his example, we have no trouble considering the possible relations between quarks and Camembert cheese, which are not plausibly part of the domain of any single module. This potential for the free integration of information is precisely what the Generality Constraint is intended to capture. The question, then, is how concepts can be understood as being modular, while at the same time being integratable across domains.

Sperber proposes that there is a separate modular component of the mind whose existence is independently certified that has the capacity to perform the required integration task. This system is the so-called *metarepresentation module*. As its name suggests, the function of this system is to produce representations of representations. The main role that this capacity serves, according to some (e.g., Perner, 1991, 1994, 1995), is in underpinning our folk psychological competence: in attributing a mental state to another, I must represent her mental state, which in turn requires representing the way that she is representing the world. Understanding the representational states of others is crucial to explaining their behavior, both verbal and non-verbal. If someone races out the door while shouting "I left the stove on!", I may attribute to her

[5] These criteria make it clear that Sperber is working with something like Fodor modularity in mind. But since modularity in my sense is entailed by Fodor modularity, this does not affect the argument.

[6] The claim about modality-independence here is dispensable. The main relevant point is that the conceptual system is the locus where information from multiple different sources can be freely combined. Whether this requires a separate, amodal representational code is a different matter. I won't address questions of representational *format* here.

not just the belief (a representational state) that the stove is (still) on, but also the belief that the house might be on fire, and the desire (another representational state) to keep the house from burning down. These states go some way towards explaining her mad dash towards her car.

Apart from its potential utility in underpinning folk psychology, the metarepresentation module purportedly carries out the kind of conceptual integration characteristic of Generality. It is capable of doing this because *all* of the mind's representations (or at least all of its conceptual representations) are among its possible inputs. So while the folk physics module deals only with representations of physical properties, and the 'table module' deals only with tables, the metarepresentation module has access to all concepts—TABLE, QUARK, CAMEMBERT, MANATEE, MORTGAGE, etc. This is required if it is able to satisfy the folk psychological task description: to attribute to a subject, or to myself, the belief that manatees like Camembert requires representing those concepts, combining them into a single representation, and prefacing it with a belief operator. Using bold type to indicate metarepresentations, the system must be able to construct thoughts such as: BELIEVES(HER, **MANATEES LIKE CAMEMBERT**), which deploys the (first-order) concept of belief plus an individual concept and the (second-order) concept of a particular complex representation.[7] This states that an individual is belief-related to the thought containing the concepts MANATEES LIKE CAMEMBERT. Just as this can be applied to interpreting the cross-domain representational states of others, so too it can be applied to my own representational states. If I possess MANATEE and CAMEMBERT, I can think about a Camembert-liking manatee by letting the metarepresentation module take those concepts as input to produce **MANATEES LIKE CAMEMBERT**.

[7] To be absolutely clear on the notation here: I continue to use small caps for concepts. Bold type indicates metarepresentations; i.e., representations not of properties or individuals, but of the concepts of those properties or individuals. So **MANATEE** indicates the concept that represents the concept of a manatee.

But *prima facie*, this account seems puzzling. Recall that domain specificity can be understood in terms of inputs or functions. The metarepresentation module is supposed to take all of a creature's representations as its possible inputs. So if inputs determine domains, then its domain must be everything the creature can represent.[8] And this would mean that the metarepresentation module is in fact a *domain-general* system, not a domain-specific one. Obviously this is in tension with treating it as a modular system at all, since part of the point of modularity is that modules deal only with a restricted range of things the creature can think about.

It isn't clear whether this is how Sperber is understanding domain specificity since he does not explicitly discuss how domains are determined. Given this argument, however, the input conception seems unavailable to him. So I will suppose instead that domains are determined functionally; that is, the metarepresentation module takes all concepts as its inputs but its domain is something more narrow, e.g., the domain of *representations of* those input concepts. This comports with the idea that the purpose of the module is to create and manipulate such representations for a range of cognitive purposes.

But while this evades the initial domain specificity worry, there are further problems that arise when we consider the content of the representations the metarepresentation module creates, and the processes it carries out over them. Take the process worry first. Sperber says that the module

> may know something about semantic relationships among representations; it may have some ability to evaluate the validity of an inference, the evidential value of some

[8] Sperber distinguishes actual from proper domains, but both of these might be determined by the inputs to the system—the actual domain being the inputs that now turn the system on, the proper domain being those that did turn it on during its evolutionary history.

> information, the relative plausibility of two contradictory beliefs, and so forth. It may also evaluate a belief, not on the basis of its content, but on the basis of the reliability of its source. (Sperber, 1994, 61)

These are the traditional functions assigned to a domain-general reasoning mechanism, however: weighing the epistemic properties of beliefs without relying on domain-specific information, considering the purely formal features of inferences, etc. One of the key arguments for massive modularity is that wholly general reasoning rules are likely to be both inefficient and prey to various sorts of computational intractability. But on Sperber's proposal, the metarepresentation module carries out many of the same functions as domain general central systems were classically proposed to do. Indeed, it may even have fairly wide-ranging access to information, if it is able to access the sources of the beliefs it deals with. This module, then, might just amount to a kind of 'central system lite'.

We will return to this problem in discussing Carruthers' view. Turn now to the content worry. Consider the way we might interpret non-mental representations such as pictures and words. When I look at a child's drawing of a bear, I need to construct a mental representation of this pictorial representation, and then pass it to a metarepresentation module that 'interprets' the pictorial representation by constructing a representation that matches the picture in its content. Perner (1995) proposes something like this: on his view, the goal of metarepresentation is to construct propositions of the form **R** SEMANTICALLY ENCODES P, where **R** represents some representation R (mental or not), and P is an encoding of R using the attributor's own conceptual resources.

There is some evidence that Sperber adopts this view. He comments (Sperber, 2000, 134) that the ability to assess the validity of arguments, the weight of evidence for a proposition, etc., might appear to be domain-general, since these arguments may concern any subject matter. However, he says, the device does not respond to the content properties of these arguments; rather, it responds only to their *formal* properties. A consistency checker need not know anything about the content of the sentences it is checking—it need only be able to detect the formal indicators of contradiction. So rather than being domain-general, a metarepresentation module actually deals just with the domain of formal or logical properties of representations. This roughly conforms to the line of argument sketched above.

But if this is the case, then the content of the representations that the metarepresentation module manipulates does not seem to be the *right* content to explain the Generality of our conceptual system. For consider: the intuitive content of the thought CAMEMBERT IS MADE OF QUARKS is the proposition *Camembert is made of quarks*, which is composed of the property of being Camembert, etc.[9] It is a proposition that is about cheese and the particles that make it up. But the content of the thought **CAMEMBERT IS MADE OF QUARKS** is the proposition *CAMEMBERT IS MADE OF QUARKS*, which is composed of the property of being the concept of Camembert, etc. This proposition is not about cheese, but about the concept of cheese. More specifically, for Sperber, it is about the formal properties of the concept of cheese, since that is what the module accesses. And these are, obviously, two quite different things.

So what the metarepresentation module can do is to represent any concept in its input range, which is all of the concepts available to a thinker. And it can then take these

---

[9] On notation: I am using italics here to indicate contents. So *cat* is the property of being a cat, which is represented by the concept CAT.

metarepresentations and carry out various processes of logical inference, weighing evidence, assessing relative confirmation, and so on; any sort of processing that involves responding to their formal properties. The thoughts it constructs are not thoughts about the world, but about other thoughts. If this module is supposed to explain cross-domain integration, we are faced with the obvious problem that when we entertain thoughts that span modular domains, we seem to be thinking about things in the world (cheese and particles), not about our concepts and representations of those things.

So Sperber's proposal faces serious difficulties. As we have seen, he needs to adopt the functional notion of domain specificity. But the domain specific module he posits has two problems. First, the processes it involves are uncannily similar to the ones that belong to traditional central systems. Other massive modularists have resisted positing modules capable of carrying out wholly general inferences across domains, but Sperber's module seems to do just this. Second, the content of the module introduces an implausible element of semantic ascent into our everyday cognition. When I wonder whether Camembert is made of quarks, it does not *seem* as if I am asking the second-order question whether CAMEMBERT IS MADE OF QUARKS is true.

None of this is to deny that we can, and perhaps often do, engage in metarepresentation, nor is it to deny that there might be a modular faculty for doing so. It is to deny, however, that such a faculty can explain our evident capacity to entertain domain-crossing thoughts in the way that the Generality Constraint requires.

### 3.2. Carruthers and role of language as content-integrator

Carruthers (2002, 2003, 2006a, 2006b) develops a sweeping picture of a massively modular cognitive architecture with enormous attention to empirical detail. Assessing the plausibility of the architecture as a whole is beyond the scope of my discussion here. Rather, I will focus on the particular mechanism that he proposes to integrate the outputs of the various modular subsystems. As with Sperber, my argument will be that the proposed mechanism is insufficient to the task.[10]

Briefly, the relevant details of Carruthers' architecture are as follows (Carruthers, 2006a, 211-276). We have a host of distinct, domain-specific central modules that have both dedicated input channels (from perception and from other central modules) and access to a common global memory workspace. Suppose two of these modules produce two distinct representations, R1 and R2. How can these thoughts from different domains be combined? Carruthers appeals to the fact that the language faculty is positioned to interface with all conceptual modules. Roughly, we can *talk* about anything we can *think* about. First, R1 and R2 are passed to the language production subsystem, which outputs a third representation, this time of a natural language sentence S that encodes at least some of the content from both R1 and R2.[11] This it can do because natural languages contain indefinitely many slots into which words carrying arbitrary content can be inserted: e.g., another adjective or clause can always be inserted, producing a slightly longer and more complex sentence. The production system then passes S to the perceptual systems, which produce the experience of 'hearing' (having an auditory image of) S. This perceptual state is then fed to the language comprehension system, which produces a final representation R3 that

[10] Machery (2008) makes substantially the same argument against Carruthers as the one presented here. Though he does not connect it with constraints on concept possession, he does contend that Carruthers' sort of modularity cannot explain true content flexibility.

[11] It's important that what is produced first here is a representation of a sentence that integrates the relevant content, not the content-integrated thought itself.

decodes the message that sentence carries.[12] This last representation, having been produced from the representation of a sentence that carries content from several distinct domains, will itself integrate content from several different domains. And R3 is now available to be globally broadcast to any conceptual modules that can read its input, so they may continue to process it, draw further inferences, and elaborate on its content.

To take an example, suppose I witness a cat wearing a diamond-studded collar. I might then think: $\text{THAT}_1$ IS A CAT (picking out the animal in question); $\text{THAT}_2$ IS A DIAMOND COLLAR (picking out the collar); and $\text{THAT}_1$ IS WEARING $\text{THAT}_2$ (indicating the relation the one perceived thing bears to the other). The first thought might be produced by a folk biological module, the second (let's imagine) by a mineralogy module; the third is just derived from the spatial layout of the scene. To combine these into a single thought, they are passed to the production system, which constructs a representation of the sentence "That cat is wearing a diamond collar". Finally, the comprehension system, in turn, decodes this to produce the thought THAT CAT IS WEARING A DIAMOND COLLAR, a content-integrated representation that couldn't have been produced by any single conceptual module.

Carruthers summarizes his view thus:

> It is, moreover, our capacity for creative assembly of speech action schemata that explains why human thought is fully subject to Evans' (1982) 'Generality Constraint'….

[12] This complex 'looping' arrangement, whereby sentences must be internally produced and decoded, is necessary because Carruthers assumes that there is no mechanism apart from language that can produce representations that cross modular domains in the way that the outputs of the comprehension system do. But to make use of the comprehension system, we need to feed it its proper inputs, which can only come from the production system. One extremely interesting question about this architecture involves whether it makes cross-domain thoughts take too long to compose. Each such thought, on his view, would entail the use of linguistic production, auditory imagery, and comprehension. These all take measurable amounts of time. Can we formulate cross-domain thoughts in less time than our best measures of these systems' lower bounds suggest that they take? If so, this would be a piece of evidence against the architecture.

> Because I am capable of assembling and recombining the speech actions that would produce each of the component words and phrases, I can frame any sentence allowed by the grammar of my language. And I can do this without first having to entertain the thought thereby expressed. (Carruthers, 2006a, 311-312)

This would seem to have solved the problem posed by Generality. However, further problems now arise. For this cross-domain representation must still be passed to a host of distinct modular systems for processing. This is what happens to the outputs of the comprehension module: they are globally broadcast to any modules that can 'read' their constituent concepts. What sorts of operations are these systems likely to be able to carry out on these inputs, however?

Take an abstract example: the thought that A IS F AND G (where F and G belong to different domains). An F-module will be able to draw inferences from this representation, as will a G-module. But these inferences will be independent not just in processing terms (since these are distinct modules, they are distinct computational devices), but also in *informational* terms. An F-module typically draws only on its database of F-related information, and so does a G-module. Modules also frequently rely on specialized, domain-specific processing rules. So such a system will be *unable* to draw any inferences that require integrating information from across these domains.[13]

However, it is plausible that these sorts of inferences are extremely common. To find examples, we need only look at cases of so-called 'emergent' features in conceptual combination. To borrow an example from Kunda, Miller, & Claire (1990), thinking of someone

[13] Carruthers notes that we should expect there to be modules for separate and independently varying domains whenever it is important to process information about those domains simultaneously (Carruthers, 2006a, p. 27). The present point is that many putatively independent domains can be combined in ways that produce emergent sorts of interaction.

as a Harvard-educated carpenter often also results in thinking of them as non-materialistic. They suggest that this is the product of a bit of causal-explanatory reasoning that we engage into account for the co-occurrence of two oddly conjoined properties. Reasoning about carpenters alone, and about Harvard graduates alone, won’t obviously produce this emergent feature. It’s the incongruity between the features that spurs the search for further properties that explain their co-occurrence. And two autonomous inference engines operating over these concepts won’t ‘notice’ any such incongruity.

Examples of this sort can be multiplied more or less indefinitely. To borrow an example from Fodor (1981), GRANDMOTHERS MOST OF WHOSE GRANDCHILDREN ARE DENTISTS. Prinz (2002) suggests that we might attribute having very good teeth to such women, on the basis of the plausible (if defeasible) non-demonstrative inference that their grandchildren would be both able and motivated to take care of them. Here we need to draw on folk causal knowledge about dentistry and family relations, which doesn’t plausibly belong to any single module. Similarly for ARCTIC SNAKE, which is (emergently) WHITE, a fact we can infer using information about Arctic climate and ecologically sensible properties for such an animal to have.

Many more such examples come from the interpretation of nominal compounds (Bauer, 1979; Levi, 1978; Weiskopf, 2007). For instance, the noun-noun combination ‘zebra blanket’ might mean a blanket that is used to keep zebras warm, or a blanket that has a zebra-striped pattern on it. But what modular system could derive these sorts of interpretation? A system that can draw inferences from ZEBRA might be good at making various folk biological inferences, and one that can operate over BLANKET might be good at comprehending the uses of various artifacts we make to keep warm. But what system can figure out that zebras get cold at night, and that blankets can keep them warm, and hence conclude that a zebra blanket might be for that? Or note that blankets are often made in decorative patterns, and that zebras display just such a highly distinctive pattern, and hence conclude that a blanket might be made with just that pattern? These inferences involve not just the kind of transitive categorical inference that is supported by full Generality, but also various kinds of non-demonstrative reasoning about properties drawn from many domains. Just elaborating on one or the other of these concepts in isolation won't solve the problem of how to interpret the complex expression itself.

What these cases show, I think, is that Carruthers gives us a solution to the Generality problem—he does show how fully General conceptual structures might be built up—but at the price of making them incapable of entering into a full range of useful inferential relations.[14] For the point of constructing cross-domain thoughts is not just to do so for its own sake, but to do so to facilitate new and useful inferential connections, either through simple transitive inference or through more creative varieties of causal reasoning. But simply passing a cross-domain thought back to a host of specialized computational systems won't deliver this, since these subsystems lack the ability to reason about properties outside of their proprietary domain. In effect, this would be a kind of Generality severed from isotropy: the materials for cross domain inference are in place, but not the mechanisms. If this were the case, we would certainly purchase Generality at too high a price.

Carruthers may have a built-in rejoinder to this objection in the weakened notion of modularity that he adopts. Fodor and many other theorists assume that modules must be informationally encapsulated in a strong sense: they can access *only* their inputs and their

[14] Something similar seems to occur with Sperber's solution as well. Because he holds that modules *are* encapsulated, the metarepresentational module cannot access information within other conceptual modules. Hence, even if it *could* construct a full range of cross-domain thoughts, it could not draw any inferences from them that depended on the database of beliefs linked to those thoughts' constituents. So it isn't clear what sort of cognitive role these thoughts would have. Some version of the problem Carruthers faces would likely occur in Sperber's framework as well.

proprietary database, and cannot ‘see’ any other representations in the mind. On Carruthers’ view, however, modules are encapsulated only in the sense that no single module can see *all* of the information in the mind at once. But it is compatible with this position that individual modules may access subsets of this information that are stored and processed by *other* modules (Carruthers, 2006a, 57-59).

Carruthers repeatedly appeals to this operation of ‘querying’, whereby one module may consult another for advice or information about some input that it is currently processing. He says, for instance, that the mindreading module “can access some of the contents generated by those other systems, no matter whether they be concerned with mates, financial institutions, falling stones, or whatever” (Carruthers, 2006a, 248). And again, the mindreading system “may need to work in conjunction with other elements of cognition in providing us with a solution to a problem, querying other systems for information” (Carruthers, 2006a, 12). Modules can outsource certain questions to other modules. Perhaps something like that occurs in the cases discussed above.

To see whether this strategy will work, consider Carruthers’ (2008) replies to Machery’s (2008) criticisms, which are very similar to the ones presented here. He claims that combining information from across domains results in its being more salient and available. So when I think (for example) the thought that the cheater is my brother, the presence of the concept BROTHER may trigger various emotions and norms related to family and filial loyalty. This might counteract the negative norms attached to the concept CHEATER. Rather than invoking a merely punitive attitude towards this instance of norm violation, then, I might be inclined to take a more gentle or nuanced position in light of my personal relationship to the cheater. Information pertaining to my brother, my family, my obligations to them, etc., is salient in a way that it would not have been otherwise. Combining concepts brings competing norms and attitudes into play, and these influence my deliberation.

Another of Carruthers' replies makes essentially the same point. Metaphors, he notes, have a characteristic cross-domain structure. A child may think or pretend that a banana is a telephone. This in turn can lead to a range of telephone-related behaviors directed at the banana (e.g., pretend calls to grandma, etc.). This behavior would not have been activated if it were not for the child's having predicated a concept from one domain of a concept from another domain. So the charge that cross-domain thoughts are causally inert in a massively modular mind cannot be sustained. Thoughts and actions may be produced by cross-domain thoughts that would not have been produced by thoughts containing either of their constituent concepts in isolation.[15]

These cases do establish that *some* cross-domain thoughts may have distinctive effects; but these are not the effects that are most distinctive of human conceptual combination. To see this, distinguish among *independent* and *interactive* effects. An independent effect is an effect that is produced by the presence of a particular concept no matter whether it is embedded in a larger, containing thought or not. Independent effects are those that a concept has no matter the cognitive context in which it is entertained. So BROTHER in the cheater example produces filial thoughts and emotions more or less independently of the other concepts in that thought. Thinking BROTHER would trigger the same associations whether it was entertained with CHEATER or not. Hence it is independent.

The same holds for the metaphor case. TELEPHONE activates a host of telephone-related behaviors independently of BANANA. True, in the present context the child applies these triggered

---

[15] There is one exception to this: the pair of thoughts $\text{THAT}_1$ IS A BANANA and $\text{THAT}_1$ IS A TELEPHONE together might have the same effect as the thought THE BANANA IS A TELEPHONE in virtue of their having a shared indexical. I will ignore this wrinkle here, however.

behaviors to the banana. But she could equally well do so even if, holding the banana, she thought (or pretended): THIS IS A TELEPHONE. These behaviors would be applied to the target of the demonstrative no matter if it were conceptualized as a banana or not. This indicates that the conceptual combination is doing no work here—the effect of telephone is independent in my sense.

Concepts act as gateways to stored information about their categories, and this information can brought to bear by their presence. Once activated, these bodies of information determine future thoughts and acts. But as the examples of conceptual combination cited earlier show, not all effects are independent in this way. Many cases of conceptual combination require some novel causal or other non-demonstrative reasoning, often of a complex and hard to specify sort. Where the effect a concept has differs depending on the cognitive context in which it is being entertained, I will say its effects are *interactive*.[16] In the Harvard carpenter example, what one knows about Harvard and what one knows about carpenters play a role in what one thinks about Harvard carpenters. But this role is not determined just by the content of those two bodies of information. Neither is it a simple application of one body of knowledge or set of action schema to another object, as in the BANANA/TELEPHONE case. Some extra information and reasoning is required to determine that Harvard carpenters are non-materialistic. One needs to reason about the relative pay and social status of Harvard grads and carpenters, and decide what they might say about the personality of someone who is both. The effect of combining these

[16] I don't mean to suggest that there is a hard and fast distinction between independent and interactive effects. Perhaps every concept's effects are modulated to some degree by the context. But the examples Carruthers gives of how conceptual combinations may affect cognition are clearly closer to the independent side of the spectrum. Emergent features in concept combination are canonical cases of interactive effects. And insofar as Carruthers' modular systems can produce at best independent effects, they are missing an enormous amount of what happens in conceptual combination.

concepts, and many others like them, is interactive: it depends on a much wider body of knowledge than that which is activated by either concept taken alone.[17]

What sort of effects follow from entertaining a cross-domain thought depend entirely on the sorts of processes that are available to operate over the information that they provide. The problem with interactive effects, from the point of view of the massive modularist, is that any such processes will have to be modular, hence domain-specific. And it isn't at all clear that there are such processes, for any set of concepts that might be combined.

Indeed, there may *be* no single system that has the capability to draw the required inferences. In the case of the Harvard carpenter, or the grandmother of a dentist, we rely on what appears to be a general capacity to draw causal inferences from a wide range of background information. These inferences, insofar as we understand the mechanisms behind them, don't appear to rely on any particular domain-specific principles. For example, psychological models of compounding have appealed to general processes like similarity comparison (interpreting BOOK MAGAZINE involves comparing highly similar concepts), property mapping from the modifier to the head (BOX CLOCK → clock that is square-shaped), and thematic role filling (FLOOR PIE → pie that is located on the floor).[18] It is unclear where, in a massively modular architecture, these seemingly domain-general processes should be assigned a place. Positing a 'causal reasoning module' or something similar just seems to reintroduce central system processes by another name—a problem that Sperber's account also faces. So there will still be indefinitely many conceptual combinations that cannot be successfully elaborated by modular minds.

---

[17] This is why such information is often called 'emergent': when these concepts are combined, the result is not something that is encoded in the data base of either concept taken independently. Rather, it is derived from those data bases in combination with many other pieces of information about the world.
[18] See Wisniewski (1996, 1998).

## 4. Objections and replies

### 4.1. Alternative forms of massive modularity

I now turn to some objections to this line of argument. One objection is that the target view is a straw man. Massive modularity theorists are a heterogeneous bunch, and they are not in general committed to the claim that there is *no* central system that can perform the domain-general integration function that I have been highlighting here.

This is true enough. One could consistently believe that some central functions are modular in a certain sense and also think that there exist central systems that integrate all conceptual information available to a creature. Central modules might be devices responsible for carrying out certain rapid and automatic processes of belief fixation concerning their domains, so long as the *products* of those devices were made available to a Generality-conforming inferential system. Perhaps this is what many massive modularists have had in mind, though I would note that it is clearly not the view of the Sperber and Carruthers, who adopt the more radical view that the mind contains *no* non-modular central systems.

The only caveat that needs to be added concerns the intermediate stages of processing that such devices carry out. If central modules depend on a proprietary database of information that *cannot* be accessed by any other system in the mind, then that database cannot, on pain of failing to obey Generality, contain any concepts that can't be part of either the device's input or output representations. The intermediate representational transformations, that is, cannot use any concepts that are in principle unavailable to other processes outside the module itself, since such 'concepts' could not be attached to thoughts other than those processed within the module itself.

But so long as the conceptual vocabulary of these modules is potentially publicly available—even if the database itself is not—then the Generality Constraint can be satisfied.

So nothing I have said here conflicts with 'moderate' massive modularity of some varieties. The argument is only that genuine concept possession is incompatible with the strongest variety of modularity, as exemplified by Sperber and Carruthers.

4.2. Causal versus metaphysical Generality

Carruthers has charged that the Generality Constraint should be interpreted in such a way as to block unsavory consequences such as the ones educed here. He suggests that we distinguish between causal and metaphysical readings of the Constraint. Speaking of honeybee cognition, he says:

> from the fact that a bee lacks any mechanisms that could ever cause it to go into a state with the content [nectar is 200 meters north of the brood chamber] it doesn't follow that the bee lacks compositionally structured thoughts. The Generality Constraint, if it is to be defensible, should be seen as a *metaphysical* constraint on genuine compositional structure, not as a constraint on the causal processes that serve to generate thoughts. (Carruthers, 2006a, 81)

Similarly, from the fact that a massively modular mind cannot create thoughts that bring together concepts from arbitrary domains, it shouldn't follow that it doesn't satisfy the metaphysical interpretation of Generality; at least, not if its rules of combination plus its vocabulary of concepts permit, in principle, the formation of such thoughts.

I think we need to adopt a causal, not metaphysical, interpretation of Generality. Consider first a creature that can produce structured thoughts that describe the kinds of objects and situations that it has actually encountered. So it can think THE CAT IS FURRY, THE MAT IS BLACK, and THE CAT IS BLACK; but given that it never encounters a furry mat, it will never think THE MAT IS FURRY. But if it were to encounter such a mat, it *would* think that thought. We might want to complicate things a bit by adding that it has to have the right sort of attention and motivational structures to formulate such a thought, but it is causally available to the creature to do this. Now consider a creature for whom FURRY-thoughts are *architecturally* restricted to the domain of living things. It can apply that predicate to cats and coyotes, but not to inanimate entities. But it still possesses a set of rules of combination such that THE MAT IS FURRY is a permissible formation under those rules; it's simply causally *impossible* for the creature to produce that thought, since the architecture itself doesn't allow those representations to be combined in that ways.

Both of these creatures might, given the course of their experience, fail to be causally capable of thinking that THE MAT IS FURRY. But in the case of the former creature, that thought is still causally *accessible* given the kind of architecture it possesses. It just needs to be situated in the right environment, stocked with the right sort of motivation to pay attention to the appropriate things and properties, etc. For the second creature, no amount of such good fortune will enable it to entertain the thought. For both creatures in a certain situation it might be causally inaccessible. But for the latter creature it is inaccessible for *structural* reasons. Nothing it could do, short of rewiring its architecture, could make the thought thinkable for the creature. The thought is *necessarily* unthinkable for the creature as it stands.

I think that we should say that the kind of 'availability' that the latter creature possesses is not sufficient to satisfy Generality. While the constraint strictly speaking only says that a creature needs to have "the conceptual resources for entertaining" a thought that recombines constituents from other thoughts, it is hard to see how a creature that is necessarily incapable of putting two concepts together has the resources to entertain the thought containing both of them. The rules of combination allegedly permit that thought to be tokened by the creature, but its architecture prohibits it from being tokened. The relevant question in assessing whether Generality is satisfied is: is it possible for the creature to think such-and-such kind of thought?

Metaphysical Generality allows that this might be *impossible*, and hence it is too weak an interpretation of the Constraint. In the case of the former kind of creature, even if it never in its life encounters circumstances that would cause it to token THE MAT IS FURRY, it is still possible for it to do so. This might be so even for thoughts that the creature would find bizarre, e.g., THE NUMBER TWO IS FURRY. Arguably we are able to entertain such thoughts, since otherwise it is hard to see how we could be in a position to find what they describe bizarre. This is further evidence that our kind of conceptual system is one that obeys full Generality.

Indeed, a stronger case can be pressed against the metaphysical interpretation. Complex thoughts are composed according to a certain set of combinatorial rules. But how do we determine what combinatorial rules a creature's thoughts obey? The only evidence that we have for making such attributions is the range of states the creature can be caused to enter. In general, the attribution of a certain combinatorial rule to a creature should be constrained by what states are actually or counterfactually available to it. If we have two candidate sets of rules, one of which is entirely General, and another of which is restricted in some way, in choosing between these attributions we should adopt the one that most closely fits the causal profile of the thoughts the creature can entertain. Otherwise we risk overattribution: assigning a creature a set of compositional rules that exceed its true capacities.

I suggest that precisely this is happening in the bumblebee case Carruthers discusses. If bees, for architectural reasons, cannot enter certain states, it isn't true that their thoughts satisfy metaphysical Generality, since the set of combinatorial rules that they possess isn't in fact General. Of course, all creatures have some limits on the thoughts they can entertain: processing, memory, and resource limits are notable examples. So no creature's causal profile will perfectly match full Generality. There is always some element of idealization in attributing compositional (and inferential) rules to cognitive systems. But where these factors do not come into play, there may still be reasons for denying that a creature meets even the conditions of metaphysical Generality. Failing to be able to entertain even simple recombinations of concepts that cross domains would be evidence that the rules the creature possesses are restricted in their scope. So no appeal to metaphysical Generality will work in such cases.

The upshot is that it is only where a system satisfies causal Generality that we can say it possesses a set of rules that would entitle us to say that its states satisfy metaphysical Generality. So weakening the Generality constraint as Carruthers proposes won't help to satisfy it.

## 5. Conclusions

Interpreted in its strongest form, notion that human minds might be massively modular entails that the mind consists only of specialized computational mechanisms, without any general purpose system capable of integrating information across every modular domain. I've argued that (1) the Generality Constraint requires that such integration be possible, and (2) that it is plausible that massively modular systems make it impossible. Hence massively modular systems cannot possess anything like human conceptual abilities.[19]

That isn't, however, to say that they mightn't be capable of impressive feats of cognition and behavioral control. Susan Hurley, for example, has argued that nonhuman animals may occupy 'islands of practical rationality', such that they can apply general inferential rules in one domain but not another (Hurley, 2003, 238-239). She offers the example of a primate that can recognize the transitivity of dominance relations among conspecifics, but not the transitivity of quantities of fruit available in local trees. Even creatures with this limited ability to reason similarly across domains—and hence without conceptual capacities—might be extremely skilled at navigating their environment (Bermúdez, 2003).

One might wonder what we should say if it turns out that the best models of human cognition turn out to be massively modular, however. Isn't it possible that we should discover empirically that we don't obey the Generality Constraint? Of course, we shouldn't foreclose on this possibility *a priori*. But I'd suggest that we have ample empirical evidence that we do in fact obey the Constraint. The evidence consists of our abilities to recombine and reason about information across numerous domains, including many domains that are culturally constructed (politics, finance, sports, high-energy physics). If this facility turns out to be merely apparent, of course, then it might be that we need to revise whether the Generality Constraint is really necessary for conceptual thought. This possibility, however, seems fairly remote. More likely is that models of the mind that impose strong modularity on the conceptual system will need to be

[19] For some other doubts about modularity in general, and central modularity in particular, see Prinz (2006) and Sloman, Lombrozo, & Malt (2007).

supplemented or radically revised to account for the flexibility with which we draw novel inferences across the domains provided by whatever evolved modules we possess.*

REFERENCES

ARMSTRONG, D. M. 1968, *A Materialist Theory of the Mind*, London: Routledge
BARRETT, H. C. 2005, "Enzymatic computation and cognitive modularity", Mind and Language 20, pp. 259-287
BARRETT, H. C., & KURZBAN, R. 2006, "Modularity in cognition: Framing the debate", Psychological Review 113, pp. 628-647
BAUER, L. 1979, "On the need for pragmatics in the study of nominal compounding", Journal of Pragmatics 3, pp. 45-50
BERMÚDEZ, J. L. 2003, *Thinking without Words*, Oxford: Oxford University Press
BERNAL, S. 2005, "Object lessons: Spelke principles and psychological explanation", Philosophical Psychology 18, pp. 289-312
BOYER, P., & BARRETT, H. C. 2005, "Domain specificity and intuitive ontology", in: D. M. Buss, ed., *Handbook of Evolutionary Psychology*, New York: Wiley, pp. 96-188
BUSS, D. M. 2008, *Evolutionary Psychology: The New Science of the Mind* (3rd ed.), Boston: Allyn and Bacon
CAMP, E. 2004, "The generality constraint, nonsense, and categorical restrictions", Philosophical Quarterly 54, pp. 209-231
CAMP, E. 2009, "Putting Thoughts to Work: Concepts, Systematicity, and Stimulus-Independence", Philosophy and Phenomenological Research 78, pp. 275-311
CARRUTHERS, P. 2002, "The cognitive functions of language", Behavioral and Brain Sciences 25, pp. 657-674
CARRUTHERS, P. 2003, "Moderately massive modularity", in: A. O'Hear, ed., *Mind and Persons*, Cambridge: Cambridge University Press, pp. 69-91
CARRUTHERS, P. 2006a, *The Architecture of the Mind*, Oxford: Oxford University Press
CARRUTHERS, P. 2006b, "The case for massively modular models of mind", in: R. Stainton, ed., *Contemporary Debates in Cognitive Science*, New York: Blackwell, pp. 3-21
COSMIDES, L., & TOOBY, J. 1994, "Origins of domain specificity: The evolution of functional organization", in: L. A. Hirschfeld & S. A. Gelman, eds., *Mapping the Mind*, Cambridge: Cambridge University Press, pp. 85-116
COWIE, F. 2008, "Us, them, and it: Modules, genes, environments, and evolution", Mind and Language 23, pp. 284-292
EVANS, G. 1982, *The Varieties of Reference*, Oxford: Oxford University Press
FODOR, J. 1981, "The present status of the innateness controversy", in: J. Fodor, *Representations*, Cambridge, MA: MIT Press, pp. 257-316
FODOR, J. 1983, *The Modularity of Mind*, Cambridge, MA: MIT Press

* Thanks to the referees for this journal for their extremely helpful comments on an earlier version of this paper, and to Elisabetta Lalumera for the invitation to participate in this special issue, and her efforts in putting it together.

FODOR, J. 2000, *The Mind Doesn't Work That Way*, Cambridge, MA: MIT Press
HURLEY, S. 1997, "Nonconceptual self-consciousness and agency: Perspective and access", Communication and Cognition 30, pp. 207-248
HURLEY, S. 2003, "Animal action in the space of reasons", Mind and Language 18, pp. 231-256
KUNDA, Z., MILLER, D., & CLAIRE, T. 1990, "Combining social concepts: The role of causal reasoning", Cognitive Science 14, pp. 551-577
LEVI, J. L. 1978, *The Syntax and Semantics of Complex Nominals*, New York: Academic Press
MACHERY, E. 2008, "Massive modularity and the flexibility of human cognition", Mind and Language 23, pp. 263-272
PEACOCKE, C. 1992, *A Study of Concepts*, Cambridge: MIT Press
PERNER, J. 1991, *Understanding the Representational Mind*, Cambridge: MIT Press
PERNER, J. 1994, "The theory of mind deficit in autism: Rethinking the metarepresentation theory", in: S. Baron-Cohen, H. Tager-Flusberg & D. J. Cohen, eds., *Understanding Other Minds*, Oxford: Oxford University Press, pp. 112-137
PERNER, J. 1995, "The many faces of belief: Reflections on Fodor's and the child's theory of mind", Cognition 57, pp. 241-269
PINKER, S. 1999, *How the Mind Works*, New York: W. W. Norton
PRINZ, J. 2002, *Furnishing the Mind*, Cambridge, MA: MIT Press
PRINZ, J. 2006, "Is the mind really modular?", in: R. Stainton, ed., *Contemporary Debates in Cognitive Science*, New York: Blackwell, pp. 22-36
SAMUELS, R. 2000, "Massive modular minds: evolutionary psychology and cognitive architecture", in: P. Carruthers & A. Chamberlain, eds., *The Innate Mind*, Cambridge: Cambridge University Press, pp. 13-46
SEGAL, G. 1996, "The modularity of theory of mind", in: P. Carruthers & P. Smith, eds., *Theories of Theories of Mind*, Cambridge: Cambridge University Press, pp. 141-158
SHALLICE, T. 1988, *From Neuropsychology to Mental Structure*, Cambridge: Cambridge University Press
SLOMAN, S., LOMBROZO, T., & MALT, B. C. 2007, "Ontological commitments and domain-specific categorization", in: M. J. Roberts, ed., *Integrating the Mind*, Hove, UK: Psychology Press, pp. 105-129
SPERBER, D. 1994, "The modularity of thought and the epidemiology of representations", in: L. A. Hirschfeld & S. A. Gelman, eds., *Mapping the Mind*, Cambridge: Cambridge University Press, pp. 39-67
SPERBER, D. 2000, "Metarepresentations in an evolutionary perspective", in: D. Sperber, ed., *Metarepresentations*, Oxford: Oxford University Press, pp. 117-137
SPERBER, D. 2002, "In defense of massive modularity", in E. Dupoux, ed., *Language, Brain, and Cognitive Development*, Cambridge: MIT Press, pp. 47-57
SPERBER, D. 2005, "Modularity and relevance: How can a massively modular mind be flexible and context-sensitive?", in: P. Carruthers, S. Laurence & S. Stich, eds., *The Innate Mind*, Oxford: Oxford University Press, pp. 53-68
TOOBY, J., & COSMIDES, L. 1992, "The psychological foundations of culture", in: J. Barkow, L. Cosmides & J. Tooby, eds., *The Adapted Mind*, Oxford: Oxford University Press, pp. 19-136
WEISKOPF, D. A. 2007, "Compound nominals, context, and compositionality", Synthese 156, pp. 161-204
WISNIEWSKI, E. J. 1996, "Construal and similarity in conceptual combination", Journal of Memory and Language 35, pp. 434-453

WISNIEWSKI, E. J. 1998, “Relations versus properties in conceptual combination”, Journal of Memory and Language 38, pp. 177-202