13

The Distinctively Human Mind: The Many Pillars of Cumulative Culture

PETER CARRUTHERS

This chapter argues that there are multiple adaptations underlying the distinctiveness of the human mind. Careful analysis of the capacities that are involved in the creation, acquisition, and transmission of culture and cultural products suggests that it is very unlikely that these could be underlain by just one, or even a few, novel cognitive systems. On the contrary, there are at least eight such systems, each of which is largely independent of the others.

INTRODUCTION

What makes human minds distinctive? The question is ambiguous, depending on the extent of the implied contrast. It might mean, "distinctive in relation to the minds of non-human animals, especially the other great apes." Or it might mean, "distinctive in relation to the minds of earlier forms of hominin, especially the Neanderthals." My focus in this chapter will be on the former issue. I am by no means indifferent to questions of hominin cognitive evolution, especially as revealed in the archaeological record. But my main interest, here, is not in the timing and sequencing of the adaptations that led from the last common ancestor of humans and chimpanzees to ourselves, but rather in the question of *how many* distinct capacities are needed to explain the difference. I argue that there are at least a handful of such capacities (probably many more), rather than just one or a few.

Such claims are not new, of course. Evolutionary psychologists often claim that there are multiple adaptations—or "modules"—underlying the distinctiveness of the human mind (Barkow, Cosmides, and Toody 1992, Sperber 1996, Pinker 1997). I have elsewhere endorsed such views, and have laid out the evidence that supports them (Carruthers 2006). But in the present chapter I propose to proceed more negatively, and mostly by way of task analysis (while also drawing on a range of other forms of evidence). I shall consider the views of some of those who have hoped to explain human cognition in terms of one or two major new capacities, arguing not only that these accounts are insufficient, given an analysis of the tasks involved, but also that the specific ways in which they are insufficient demonstrate the need for a number of further adaptations.

What exactly is meant, however, by "explaining human cognition"? And if one is to proceed by way of task analysis, which cognitive tasks are to be analyzed? Ultimately, of course, our goal must be to explain *all* of the ways in which humans are cognitively and behaviorally distinctive. And there are a great many of them. My book, *The Architecture of the Mind* (2006) lists 22, ranging from a capacity for science, through dispositions to sing, dance, and listen to music, to a capacity for humor. But even this catalog is very partial, and should really be much longer. (See Brown 1991 for an extensive list of human universals that are nevertheless unique to the human species.)

Not all of these capacities will be cognitively primitive, of course. Some, at least, will be explicable in terms of others, or in terms of combinations and interactions of others, together with various kinds of learning. But taken in their totality they present a formidable (and in my view insurmountable) challenge to anyone who claims that there is just one adaptation distinctive of the human mind (or even a few such adaptations). Yet this is a challenge that such authors rarely, if ever, attempt to meet. Their explanatory focus is almost always a great deal narrower, often on the distinctively human capacity for cumulative culture. But even in this restricted domain multiple adaptations are required, as we will see.

There is no doubt that our capacity for culture is distinctive in its extent, and that it both underlies our unique adaptability (accommodating ourselves successfully to environments from tropical rainforests to arctic tundra) as well as our most impressive technological and scientific achievements. Tomasello (1999), Donald (2001), and Sterelny (2003) all emphasize the importance, in explaining these facts, of *cumulative* culture. This enables the operation of a "ratchet effect" for human behavioral capacities, as well as for our powers to control, transform, and interact with natural environments. The discoveries and successful practices of one generation can be passed on to the next, providing both the context in which the new generation does its learning, as well as a novel platform from which new discoveries and practices can be developed. In the present chapter I shall join these authors in taking the capacity for cumulative culture as my target for explanation. I shall argue that there are a significant number of adaptations underlying it, rather than just one or two.

A note on terminology before I proceed: although evolutionary psychologists use the language of "modularity," and although my earlier book (2006) is written in defense of a massively modular account of the human mind, I shall not employ that term in the present chapter. This is because too many people have become fixated on what Fodor (1983, 2000b) understands by "module" (and on the requirement that modules should be informationally encapsulated, in particular). Talking in terms of modules would require an extended discussion of the appropriate notion of modularity to employ in the context of an evolutionary psychology, and would distract from the main point at issue. (That discussion is undertaken in the first chapter of Carruthers 2006.) Accordingly, I shall talk of "adaptations" and "systems" instead.

(My use of the former term will be somewhat loose, however. For a case could be made that all of the systems that I shall discuss evolved together in a complex interacting package as the life-history of humans became increasingly social and cultural. In which case they might really constitute a single, complex adaptation. I shall not pursue this question any further.)

It should be stressed, moreover, that most of the systems that I shall be discussing are unlikely to be either unitary (or "atomic") in structure or localized to just one area of the brain. On the contrary, for the most part they possess a complex internal organization, with the various components located in different brain areas. Moreover, those components might initially have evolved for other purposes, and might still also be employed in the service of other systems and other tasks (Anderson 2008). To take just one example: the "mindreading" or "theory of mind" system to be discussed shortly appears to consist of at least four dedicated components distributed in three different areas of the cortex, together with a number of others that are also utilized for other purposes (Frith and Frith 2003, Saxe and Powell 2006). And in the course of its normal operations the mindreading system will need to recruit and interact with multiple other systems also (Nichols and Stich 2003).

One final preparatory comment before we get on with the main business: someone might wonder about the criteria that I have employed when selecting uniquely human capacities as candidates for discussion, since there are a number of other capacities that have also been discussed in this context. A capacity for "mental time travel"—recalling the past and foreseeing the future—would be one. My answer is that I have selected only those capacities that I believe I can defend successfully as being both uniquely human and necessary for cumulative culture. I do not discuss mental time travel, for example, because there is evidence that non-human animals can both recall specific events and plan for the future (Clayton, Emory, and Dickinson 2004, Sanz, Morgan, and Gulick 2004, Mulcahy and Call 2006, Correia, Dickinson, and Clayton 2007, Raby et al. 2007).

IMITATION, SHARED INTENTIONALITY, AND MINDREADING

Tomasello (1999) emphasizes the role of true imitation in enabling cultural accumulation. (True imitation is opposed to the forms of goal imitation or emulation-imitation that he says one finds among apes and monkeys. It involves identifying the goals of the actor, as well as a capacity to parse an action into its component parts.) Indeed, he claims that it is sufficient. One aspect of his argument is negative: he asserts that there hasn't been enough time, in the 5 to 7 million years since the hominin line separated from the other great apes, for a suite of multiple cognitive adaptations to evolve; and he claims that the genetic disparities between ourselves and our nearest relatives, the chimpanzees and bonobos, are too minor to explain widespread differences in cognition. These points are supposed to motivate the search for a single major adaptation that can make the difference.

I have replied to this argument elsewhere (Carruthers 2006: chap. 3), explaining how recent research demonstrates that the genetic differences between ourselves and chimpanzees are much more extensive than was previously believed (Ebersberger et al. 2002, Anzai et al. 2003, Pan et al. 2005, ENCODE consortium 2007), and pointing out that since there are a great many physiological differences between humans on the one hand and chimps and bonobos on the other—which have plainly been able to evolve in 5 million years—there is no reason why a significant number of new cognitive capacities should not have appeared also.

In a positive vein Tomasello (1999) argues that the evolution of true imitation resulted in a high-fidelity trans-generational cultural copying mechanism. This enabled humans to gradually ratchet up their behavioral capacities generation by generation. Innovations and improvements created by individuals (often by happenstance or trial and error) could spread through the local population by imitation and thereafter be transmitted to children in the same way. This provided a platform from which new discoveries and improvements could be made, which could in turn be copied and stabilized in the population.

It is natural to wonder why, if this account were correct, we should see such long periods of stasis in the archaeological record, culminating in a veritable "cultural explosion" a mere 40,000 years ago (especially in Europe)—by which time *Homo sapiens* had already been in existence for more than 50,000 years. Two points can be made in reply. One is that the cultural explosion may be partly in the eye of the beholder, resulting from the historical concentration of most archaeological research in Europe. When the full range of data from Africa are examined carefully, in contrast, what emerges is a long period of gradual cultural accumulation of skills and abilities (McBrearty and Brooks 2000). The second point is that the cultural ratchet effect is highly sensitive to population density (Cavalli-Sforza and Feldman 1981, Shennan 2000, 2001). In dispersed populations even successful technologies can be lost, as appears to have happened in the case of Tasmania (Henrich 2004). Hence for long periods of human evolution, cultural progress may have been a fragmentary and local affair, and thus perhaps too ephemeral to show up in the archaeological record. On some accounts human population density passed through an especially tight bottleneck about 70,000 years ago, with the total human population perhaps dropping as low as 10,000 individuals (Ambrose 1998, Lahr and Foley 1998).

Tomasello et al. (2005b) shift the focus of the above account somewhat, since they are more sanguine about the imitative powers of chimpanzees than Tomasello (1999). They cite recent evidence suggesting that chimpanzees are capable of figuring out both the goals and component actions of the activities of others, and that chimpanzees do sometimes engage in true imitation. But they don't do it very often. And chimpanzees show none of the forms of *shared intentionality* that are so distinctive of human infancy and childhood. From an early age, human children will engage in protodeclarative pointing, drawing the attention of a care-giver to interesting objects and events in the environment, seemingly for its own sake. Likewise, they will monitor the patterns of attention and emotional expressions of those around them, whether to check on the safety of a new person or object, or to figure out the meaning of a new word. And they also display a drive towards cooperative activities with others, such as shared pretense and conversational turn-taking, which require them to monitor the goals and intentions of others. Chimpanzees, in contrast, seem only to pay attention to the intentions of others in competitive situations.

According to Tomasello et al. (2005b), then, the single distinctively human cognitive adaptation is actually more motivational than cognitive. It consists in an initial drive to share mental states with others. This sets up a developmental learning sequence that results not only in frequent imitation of the behavior of others (issuing in gradual skill accumulation), but also in full-blown theory of mind a few years later (especially an understanding of false belief). It also results in the acquisition of language. For according to Tomasello et al., language is entirely a cultural product—perhaps the archetypal cultural product, providing the initial "ratchet" for almost everything else that follows. (I shall return to this claim shortly.)

I have no quarrel with this stress on shared intentionality and its facilitatory effect on human imitation abilities, nor with the importance of the latter in enabling cumulative culture. But it is very likely a mistake to think that mature mindreading abilities (including false-belief understanding) are acquired via general learning from an initial starting point of shared intentionality. This is because infants who are a little over a year old (only a few months older than the first emergence of shared intentionality behavior) display an implicit understanding of false belief and its consequences for action in their patterns of expectation and surprise (Onishi and Baillargeon 2005, Southgate, Senju, and Csibra 2007, Surian, Caldi, and Sperber 2007, Song et al. 2008, Scott and Baillargeon 2009), and in their helping behavior (Buttelmann, Carpenter, and Tomasello 2009). The best explanation of young children's failures in explicit false-belief tasks is therefore some variant of what Birch and Bloom (2004) refer to as "the curse of knowledge." It is hard even for adults to set aside their own understanding of a situation and to reason and give answers from the epistemic perspective of another (Keysar, Lin, and Barr 2003). So what probably changes in the first few years of development isn't children's understanding of the mental states of others, as such, but rather an increase in their "executive function" ability to control and suppress responses based on their own view of a situation, together with an increase in working memory capacities.

The above line of criticism has little bearing on the question of the *number* of distinctively human adaptations, however. Whether we think that it is shared intentionality that leads to mature mindreading abilities via general learning (as Tomasello et al. 2005b maintain), or whether we think of the core mindreading system as emerging under maturational constraints over the first year or two of life (as I believe), we still have just a single (albeit complexly structured) distinctively human adaptation for social cognition. Of much more consequence is the question whether natural language is entirely a cultural construct, or whether it, too, has a distinct innate basis. I shall pursue this in the section that follows.

LANGUAGE

Tomasello et al. (2005b) and Donald (2001) argue that language is a cultural product, and that nothing more is needed to explain how children acquire it than general learning abilities combined with capacities for shared intentionality and true imitation. While the former claim is, in some sense, obvious, I shall argue that the later claim is (in part) mistaken, and that humans have a distinct adaptation designed for language learning. I shall begin, however, by outlining the part of the claim that is generally correct.

Bloom (2002) demonstrates in detail how a capacity to read the intentions of others is both necessary and (perhaps) sufficient for the acquisition of the vocabulary of a language. Children learn the meanings of words by figuring out what speakers are intending to refer to, especially by tracking patterns of speakers' attention (in the initial stages at least; later on, contextual cues of other sorts may be sufficient). And one can then imagine a stage in human evolution during which people were able to learn words and to communicate quite successfully with two- or three-word strings, somewhat like the simple sign languages that are created by deaf children who have had no exposure to a language exemplar (Goldin-Meadow and Feldman 1977, Goldin-Meadow and Mylander 1983)—only presumably a good deal more elaborate, since vocabulary would be shared and accumulated culturally over time. But the syntax of language is surely another matter. Here we have every reason to think that there exists a distinct language-specific learning mechanism.

One set of arguments is collected under the title "the poverty of the stimulus" (Laurence and Margolis 2001, Crain, Gualmini, and Pietroski 2005, Segal 2007). Children acquire syntax pretty much effortlessly on the basis of only positive data, and without correction or training. Adult linguists studying the syntax of a new language can elicit data from subjects about what can or—especially—*cannot* be said in that language in order to confirm and disconfirm their hypotheses, but children know, by the age of about 3, more than a linguist will know after a lifetime of theorizing. Moreover, children never make the sorts of mistakes that one would expect them to make if they were inducing the grammar of their language by general-learning methods. Indeed, some principles seem to be acquired unerringly, which would be extremely surprising if they were learning them through the use of some sort of general-purpose system. Moreover, when they do make mistakes, it often turns out that the mistaken grammatical forms that they use would actually be permitted by the grammars of another set of extant languages. Furthermore, there is the evidence that children exposed to extremely impoverished linguistic models (whether pidgin or deaf home sign) will nevertheless acquire a language with all of the syntactic subtlety and complexity possessed by any other natural language (Bickerton 1981, Pinker 1994).

In addition, Valian, Solt, and Stewart (2009) provide a direct test of the general-learning claims made by Tomasello (2000a, 2003). According to the latter, acquisition of syntax proceeds "bottom-up," gradually generalizing from specific examples and formulaic phrases to abstract principles. If there is an innately structured syntax-learning mechanism, in contrast, then abstract syntactic categories should be available from the outset. Valian et al. (2009) set out to test the predictions of these different views in connection with children's understanding of determiners like "a," "the," and "this." On Tomasello's general-learning account, children's understanding should initially be exemplar-based, and generalization to new instances should be slow. If children already possess the abstract category *determiner,* in contrast, then generalization to new cases should be immediate, and the distribution of young children's determiner phrases should closely resemble those of their mothers. While reviewing a range of additional supporting evidence, Valian et al. (2009) were able to show that the predictions of the specialized learning-mechanism account are strongly supported by their data, whereas Tomasello's general-learning approach is disconfirmed.

I think we can assume, then, that our syntactic abilities cannot be explained in terms of our capacity for shared intentionality and/or mindreading combined with general learning. But what about the reverse? Can our distinctive forms of social cognition be explained in terms of our capacity for language? Just such a position has been proposed and defended by Bickerton (1995, 2005), who claims that the evolution of language was the event that explains all other distinctive features of human cognition. Such a position is highly implausible, however, and partly for reasons internal to Bickerton's own theory. He argues persuasively that an intermediate stage in the evolution of a language faculty would have been a capacity for *proto-language,* which would have consisted in communicative uses of strings of words unregulated by grammar, not unlike those that we see in pidgin languages, or in the home sign languages of the deaf. But it is impossible to see how these proto-languages could be created and sustained without sophisticated forms of social intelligence. It takes mindreading to figure out what someone is trying to communicate to you in proto-language, and it takes mindreading to use such a language in successful communication. Indeed, the only plausible accounts that we have of the evolution of language presuppose that some form of mindreading was already in existence (Origgi and Sperber 2000).

It appears, then, that we need to postulate at least two distinctively human cognitive capacities: one for shared intentionality and/or mindreading, and one for natural language syntax. And indeed, just as this account would predict, there exist double dissociations between each of these capacities, both in the course of their development in childhood, and in mature individuals. Let me briefly elaborate.

People with disabilities that place them on the autistic spectrum can be deficient in developing social cognition while being normal in their acquisition and use of syntax (but not in the pragmatics of language—Baron-Cohen 1995). And within the set of people diagnosed with Specific Language Impairment (SLI), there exists a subset whose deficit appears to be localized to the acquisition of syntax only, and yet whose mindreading abilities are normal (van der Lely 2005). Moreover, adults suffering from schizophrenia are often deficient in mindreading tasks while being normal in their use of syntax (Pickup and Frith 2001, Sprong et al. 2007). Whereas, in contrast, adults suffering from aphasia can lose all capacity for syntax while still being normal in their social and mindreading abilities (Varley 1998).

Before leaving the topic of language, we should consider whether a capacity for syntax is the *only* novel adaptation associated with it. There is some reason to think that this is not the case. This is because, as soon as linguistic communication became common (especially involving the transmission of information via testimony), there would have been intense pressure to check linguistic messages for accuracy before believing them (Carruthers 2006: chap. 6). One thing we can be certain about is that communication would often have been used for Machiavellian purposes as well as for altruistic ones. Some aspects of the adaptive problem, here, could no doubt have been taken up by the mindreading faculty, which can try to figure out whether the speaker has any motive to deceive, and so forth. But it is very likely that mechanisms for checking consistency and coherence with one's existing beliefs would also have made their appearance. For the absence of such properties would give some indication of likely falsehood, in communication just as in science. It is also very likely that capacities to engage in, and to evaluate, public forms of argumentation would have evolved for the same reason, as Mercier and Sperber (2011) argue.

CREATIVITY

So far I have argued that there are two or three distinct adaptations underlying the human capacity for cumulative culture. But it should be plain, I think, that cultural accumulation wouldn't progress very far without a capacity for creative thought. Sometimes mere happenstance or trial and error can lead to a new discovery, or a new item of technology. But not often. When humans confront problems, they generate new ideas and hypotheses in an attempt to overcome those problems. Even hunters tracking a wounded animal need to think creatively, inventing and debating explanatory hypotheses that might explain the subtle signs available to them (Liebenberg 1990).

Bickerton (1995) suggests that a capacity for creative thought is a byproduct of language. This is a natural enough idea, for we know that the large vocabulary and recursive grammar of natural language together provide us with an almost unlimited representational resource (even allowing for limitations in working memory). And we know that we can, at will, combine together words in novel ways—thereby entertaining novel thoughts—and that we often do so, e.g., in humor and in metaphor. However, it is one thing to possess a recursive representational system that makes it *possible* for people to formulate thoughts that neither they, nor anyone else, has entertained before, but it is quite another thing to have a disposition to use it thus, and to use it relevantly. And I can see no way that the former could in any sense be sufficient for the latter.

Notice that any animal that is capable of moving its limbs and torso independently of one another has the wherewithal to string together novel sequences of movement. (Indeed, humans often do precisely this, in creative free dance.) But none does so, except in limited domains. Thus many prey animals will engage in "protean" evasion procedures when being chased by a predator. These are genuinely random sequences of leaps, bounds, and turns that are impossible for a predator to anticipate (Driver and Humphries 1988, Miller 1997). Otherwise, creativity of action seems to be mostly absent from the rest of the animal kingdom, with the possible exception of some limited uses by the other great apes (Carruthers 2006: chap. 2).

I have developed an account of creativity that is action-based (2006, 2007), building on a prior capacity for mental rehearsal of action (which is probably possessed by other apes as well, and hence which isn't itself distinctively human). What evolved at some point in the hominin lineage was a disposition to activate and rehearse combinations of action schemata creatively, sometimes at random (within a certain set of contextual constraints), sometimes utilizing weak semantic associations provided by the context or by current stimuli (as when the curved shape of a banana prompts a child to say, "telephone," and to pretend accordingly). And when this disposition is directed towards linguistic forms of action, the result can be novel rehearsed sentences in "inner speech" that lead the subject to entertain new thoughts. Indeed, I suggest that the proper function of childhood pretend play is to practice the process of creative supposition (especially when resulting from rehearsals of novel natural language sentences), building strategies and heuristics for creative activations of action schemata that will be both fruitful and relevant.

(Notice, by the way, that pretend play is another phenotypic character that is unique to the human species. The proposal sketched here ties that character into the foundations of the human capacity for cumulative culture.)

Whether or not this account is correct, it is surely plain that a capacity for creative thinking would not come free with the capacity for language. It would need to be a separate adaptation—unless, that is, it can be explained as resulting from our capacity for behavioral imitation. Perhaps creative thinking is a culturally constructed behavior, which children have to learn by imitating their elders and/or peers? But this suggestion is quite mysterious. It is one thing to imitate a particular action or sequence of actions—for these are determinate and observable. It is another thing to imitate an inner process of some sort for generating novel combinations and sequences of actions. In this situation it is not the actions themselves that are to be copied, but rather the creative cognitive procedure that generates them, which isn't itself open to observation.

FINE MOTOR COORDINATION

Thus far I have argued that there are a number of distinct "pillars" of cumulative culture. But one doesn't have to reflect very hard to realize that there is also another. This is our remarkable capacity for fine-grained, exact control of movement (Lemon et al. 1999). It is this that enables us to throw a projectile accurately and with force at a moving target, to strike one rock against another at the precise point and at the exact angle needed to detach a sharp flake, and to coordinate our breathing together with movements of mouth, tongue, and larynx in such a way as to produce a smooth stream of speech. And it seems plain that without such control, the basics of material culture could never have got started. There would have been no precision-made stone tools; there would have been no multi-part, finely constructed, specialized hunting weapons; there would have been no sewing of clothes; there would have been no weaving of cloth; and there would have been no pottery or basket making.

Of course, many animals have remarkable physical abilities; and many are capable of fine degrees of physical control of a specialized sort. Think of an elephant picking up a peanut or twisting off a bunch of bananas with its trunk; or think of a hummingbird that maintains its position in front of a nectar-bearing flower while being buffeted by a gusty breeze. But no other animal is capable of anything like the *range* of precise physical skills that we possess. This differential is, no doubt, at least partially explained by the vastly increased size of motor and pre-motor cortex in humans, and also by the unusually extensive projections from the human neocortex to the motor neurons in the medulla and spinal chord (Striedter 2005). In addition, no other animal is similarly motivated to practice and rehearse physical skills in the way that humans do, seemingly taking their own skill-acquisition as an intrinsic goal (Donald 2001).

It is obvious, moreover, that our distinctive motor-control capacities can neither explain, nor be explained by, any of the other distinctively human capacities identified thus far. Although fine-grained motor control may be a necessary condition for smooth and fluent speech, it plainly is, by itself, no guarantor of syntactic abilities. But then neither, obviously, can our possession of language explain our capacity to thread a needle. The same is true for social cognition and for creativity of thought. Motor control is not sufficient for either mindreading or creativity. Nor can mindreading or creativity provide any kind of explanation of our distinctive physical abilities.

PHYSICS

Yet another pillar of cumulative culture is provided by our naïve physics system. Although many animals are adept at tracking causal contingencies among events, it appears that we may be the only species capable of understanding the underlying forces and mechanisms involved (Povinelli 2000). Certainly we are unique in the extent and depth of that understanding. And without it, even the very first steps on the road to cumulative culture—the making of stone handaxes—would not have been possible. For as Mithen explains (1996), we know that the manufacture of such tools would only be possible for people who were capable of appreciating the fracture dynamics of rocks of various types, and of thinking about the physical forces necessary to produce a desired effect. Consistently with these points, attempts to teach chimpanzees to make even the very simplest of Oldowan stone tools have generally been a failure (Toth et al. 1993). It is obvious, likewise, that only a species capable of sophisticated thinking about physical principles would be able to manufacture the sorts of specialized multi-part weapons that are utilized by hunter-gatherers.

It is highly implausible that our naïve physics should be explicable in terms of any of the other distinctively human abilities identified so far, or in terms of any combinations thereof. It is plain that neither syntax, nor creativity, nor fine motor coordination could confer on those that possess them an understanding of physical forces. Nor does it appear plausible that mindreading abilities would be of any help. And even if one believes, as do Gopnik and Schulz (2004), that children's theories of the properties and forces underlying both physical phenomena and human behavior are achieved through the use of powerful Bayesian reasoning mechanisms utilizing directed causal graphs, the distinct starting theories for each of the two reasoning processes would have to be innate.

Moreover, we know that physical understanding emerges early in human childhood. Infants as young as 3 months are surprised if a moving object that disappears behind a screen is later revealed in a position that it could only have reached by passing through a barrier; and they are surprised if an unsupported object doesn't fall. Similarly, by 3 months of age infants expect an object to move when it is struck by another; and by 6 months of age they understand that the amount of displacement of the stationary object is a function of the size of the moving one; and so on and so forth (Baillargeon, Kotovsky, and Needham 1995, Spelke, Phillips, and Woodward 1995). We also know that the development of naïve physics is independent of the development of social cognition, since autistic individuals are normal with respect to the former while being deficient in the latter (Baron-Cohen 1995).

INFERENCE TO THE BEST EXPLANATION

Humans are obviously unique in their capacity for science, and science plainly has a huge impact on culture and cultural products. But—equally obviously—science is not a single, unitary capacity. On the contrary, scientific thinking is the very archetype of a "whole person" activity, involving interactions among virtually all components of the mind. For this reason some have despaired of our ability to understand (in the sense of provide a warranted cognitive-scientific theory of) our own scientific abilities (Pinker 1997, Fodor 2000b, but see Carruthers, Stich, and Siegal 2002 for contrary views). Be that as it may, it is possible to identify, and perhaps understand, some of the components of scientific thinking. One is a capacity for creative thought, which has already been discussed. This is because scientists will frequently need to display creativity in thinking up hypotheses to explain their data, or in arriving at novel ways to test their theories.

Donald (2001) argues that our scientific capacities are themselves a cultural product. There is a great deal of truth in such a claim. Many of the elements of "scientific method" have emerged slowly over the last few centuries, and are now learned by children and young adults, giving rise to novel inferential abilities. And people growing up in a scientific culture will think and reason very differently from someone growing up as a hunter-gatherer, especially if they go on to become scientists themselves (Henrich, Heine, and Norenzayan 2010). There is thus a clear sense in which the contemporary human mind is, to a significant degree, constructed by culture, resulting in further cultural change and cultural products.

These claims comport nicely with an idea that is now widely accepted by those who work on the psychology of reasoning, namely that humans possess two different *types* of cognitive systems for thinking and reasoning (Evans and Over 1996, Sloman 1996, 2002, Stanovich 1999, Kahneman 2002; see also Carruthers 2009 for extended discussion). Most consider that what is generally now called "System 1" is really a collection of different systems that are fast and unconscious, operating in parallel with one another. The principles according to which these systems function are, to a significant extent, universal to humans, and they aren't easily altered (e.g., by verbal instruction). Moreover, the principles via which System 1 systems operate are, for the most part, heuristic in nature ("quick and dirty"), rather than deductively or inductively valid. It is also generally thought that most of the mechanisms constituting System 1 are evolutionarily ancient and shared with other species of animal.

System 2, on the other hand, is generally believed to be a single system that is slow, serial, and conscious. The principles according to which it operates are variable (both across cultures and between individuals within a culture), and can involve the application of valid norms of reasoning. These System 2 principles are malleable and can be influenced by verbal instruction, and they often involve normative beliefs (that is, beliefs about how one *should* reason). Moreover, System 2 is generally thought to be uniquely human. And it is plain that System 2, as instantiated in a trained scientist, will be very different indeed from the System 2 of a hunter-gatherer.

While it is no doubt true that much of our capacity for scientific thinking is socially learned, I believe that there is one key component (in addition to creativity) that long pre-dates the scientific era, and which is probably innate. This is our capacity for inference to the best explanation. Let me elaborate.

Most philosophers of science agree that scientists employ a set of tacit principles for choosing between competing theories—that is, for making an inference to the *best* explanation of the data to be explained. The most plausible way of picturing this, is that contained within the principles employed for *good* explanation are enough constraints to allow one to rank more than one explanation in terms of goodness. While no one any longer thinks that it is possible to codify these principles, it is generally agreed that the good-making features of a theory include such features as the following: *accuracy* (predicting all or most of the data to be explained, and explaining away the rest); *simplicity* (being expressible as economically as possible, and with the fewest commitments to distinct kinds of fact and process); *consistency* (internal to the theory); *coherence* (with surrounding beliefs and theories, meshing together with those surroundings, or at least being consistent with them); *fruitfulness* (making new predictions and suggesting new lines of enquiry); and *explanatory scope* (unifying together a diverse range of data).

Liebenberg (1990) makes a powerful case that at least some of the elements of scientific thinking (including inference to the best explanation) are present among hunter-gatherers, displayed most clearly in their tracking of prey when hunting. As he remarks, it is difficult for a city-dweller to appreciate the subtlety of the signs that can be seen and interpreted by an experienced tracker. Except in ideal conditions (e.g., firm sand or a thin layer of soft snow) a mere capacity to recognize and follow an animal's spoor will be by no means sufficient to find it. Rather, a tracker will need to draw inferences from the precise manner in which a pebble has been disturbed, say, or from the way that a blade of grass has been bent or broken; and in doing so he will have to utilize his knowledge of the anatomy, detailed behavior, and patterns of movement of a wide variety of animals. Moreover, in particularly difficult and stony conditions (or in order to save time during a pursuit) trackers will need to draw on their background knowledge of the circumstances, the geography of the area, and the normal behavior and likely needs of the animal in question to make educated guesses concerning its likely path of travel.

Most strikingly for our purposes, successful hunters will often need to develop speculative hypotheses concerning the likely causes of the few signs available to them, and concerning the likely future behavior of the animal; and these hypotheses are subjected to extensive debate and further empirical testing by the hunters concerned. When examined in detail these activities look a great deal like science, as Liebenberg (1990) demonstrates. First, there is the invention of one or more hypotheses (often requiring considerable imagination) concerning the unobserved (and now unobservable) causes of the observed signs, and the circumstances in which they may have been made. These hypotheses are then discussed and evaluated on the basis of their accuracy, coherence with background knowledge, simplicity, and explanatory and predictive power. One of them may emerge out of this debate as the most plausible, and this can then be acted upon by the hunters, while at the same time searching for further signs that might confirm or count against it. In the course of a single hunt one can see the birth, development, and death of a number of different "research programs" in a manner that is at least partly reminiscent of theory-change in science (Lakatos 1970).

Given that they are universal among humans, it is hard to see how the principles of inference to the best explanation could be other than substantially innate (Carruthers 1992). For they are certainly not explicitly taught, at least in hunter-gatherer societies. While nascent trackers may acquire much of their background knowledge of animals and animal behavior by hearsay from adults and peers, very little overt teaching of tracking itself takes place. Rather, young boys will practice their observational and reasoning skills for themselves, first by following and interpreting the tracks of insects, lizards, small rodents, and birds around the vicinity of the campsite, and then in tracking and catching small animals for the pot (Liebenberg 1990). They will, it is true, have many opportunities to listen to accounts of hunts undertaken by the adult members of the group, since these are often reported in detail around the campfire. So there are, in principle, opportunities for learning by imitation. But in fact, without at least a reasonably secure grasp of the principles of inference to the best explanation, it is hard to see how such stories could even be so much as *understood*. Those principles are never explicitly articulated, yet they will be needed to make sense of the decisions reported in the stories; and any attempt to uncover them by inference would need to rely upon the very principles of abductive inference that are in question.

COOPERATION AND NORMS

Sterelny (2003) emphasizes the importance of cooperation among humans, both ancient and modern. This is a pervasive feature of all human cultures, and it has probably played a vital role throughout the course of human evolution. Early forms of it might have included banding together against predators to protect carcasses in the open savannah, sharing of meat, and various kinds of specialized labor and exchange. No doubt, also, successful cooperation played a vital role in agonistic conflicts between groups. What is distinctive of human cooperation isn't just its ubiquity, however. It is also the fact that it is regulated and enforced by systems of norms. Like the kinds of material culture that have formed the focus of most of our discussion up to now, distinct normative systems tend to distinguish different cultures from one another. They are also reliably transmitted between generations.

All human societies are pervaded by beliefs about what people may, may not, and must do (Brown 1991). Indeed, all human activities fall into one of these three categories—they are either permitted, forbidden, or required. And people in all cultures are apt to become angry or indignant if someone does something that is forbidden or fails to do something that is required. They are also apt to become angry when others in the group fail to be suitably disapproving of such breaches of normative rules, thereby engaging in a form of meta-punishment. Moreover, people everywhere experience guilt-like emotions if they themselves should happen to do something that is forbidden, or fail to do something that is required. Furthermore, the motivations involved are intrinsic ones (not merely grounded in fear of punishment or ostracization), since people will pay significant costs to punish others for breaching social norms, even when the punishers have nothing to gain—for example, because those punished are anonymous and won't be interacted with again (Fehr and Gachter 2002).

Reviewing much of the available anthropological and developmental literature, Sripada and Stich (2006) argue that humans possess a distinct multipart normative reasoning and motivational faculty. One component of the system is concerned with the acquisition and storage of the norms that are current in the surrounding culture. This is by no means a trivial task. For even where explicit instruction is given—e.g., "You mustn't hit your little sister"—the subject generally has to figure out what the real rule is that lies in the background. Another component of the system activates the norms that are appropriate in a given context, and deduces what actions or inactions are necessary to remain in compliance. And then a further component is motivational, issuing in emotions like indignation and guilt. If Sripada and Stich are correct, then it appears that we need to add a further item to our catalogue of the pillars of cumulative culture.

How plausible is it that normative thinking and feeling might be explicable in terms of one or more of the other distinctively human systems distinguished so far? I take it to be obvious that neither language, nor creativity, nor skilled behavior, nor naïve physics, nor any combination thereof is sufficient to explain our normative natures. Tomasello et al. (2005b) claim, however, that both cooperation and the normative aspects of culture can be explained as emerging out of the initial motivation to share mental states, which they think lies at the bottom of our mindreading abilities. Even if we were to grant, though, that the motivational strand in shared intentionality would be sufficient to explain human cooperation (which is actually far from obvious, for the full extent of the latter), there aren't the materials, here, with which to explain normativity. It is one thing, for example, to want (intrinsically) to cooperate with someone, and quite another thing to think and (especially) feel that you *must* do so. And it would be left mysterious why failures to cooperate should result in *guilt* (as opposed to mere disappointment) in one's own case, let alone *indignation* in the case of another.

Just as one would predict if humans possessed some sort of innately channeled normative faculty, there is plenty of evidence that a capacity to reason about rules and obligations is an early emerging one, cross-culturally (Cummins 1996, Harris and Núñez 1996, Núñez and Harris 1998, Harris, Núñez, and Brett 2001). Three-year-old and 4-year-old children are highly reliable at identifying cases where someone has broken a rule; and they are also very good at distinguishing between intentional and accidental noncompliance (categorizing only the former as "naughty"). Moreover, they do this not only in connection with rules imposed by an authority (e.g., a parent or teacher), but also when reasoning about agreements negotiated among the children themselves. And as one might expect, deontic *concepts* are acquired even earlier still. Psycholinguistic research shows that children as young as 2 years of age can make appropriate use of modal terms like "must" and "have to" (Shatz and Wilcox 1991). This latter fact is especially impressive, since few philosophers think that normative concepts can be defined, either ostensively or in terms of others.

FURTHER ADAPTATIONS TO CULTURE

A number of the pillars of culture discussed above are, at the same time, adaptations to social environments in which culture is already important to fitness. Sophisticated forms of mindreading, for example, are surely an adaptation to increasingly complex social groups made possible by simpler varieties of intention-reading. Likewise, the normative faculty discussed earlier would be an adaptation to social groups in which both cooperation and punishment of defectors was becoming progressively more common. But then in addition, of course, as humans became an increasingly cultural species, much of what they needed to learn would have had to be learned from other people—whether by copying their behavior or by believing what others tell them. And this sets up a novel adaptive problem for learners: whom to learn *from*.

One strategy is to imitate the behavior and adopt the beliefs of the majority in your community. Mathematical modeling demonstrates that this is adaptive in a significant range of circumstances (Henrich and Boyd 1998). For individual learning is difficult, error prone, and costly; whereas if the community is a stable one, its beliefs and practices are likely to be adaptive on the whole. This leads Richerson and Boyd (2005) to the claim that there should be a species-unique motivational mechanism in humans that makes us *desire to conform*. Or rather, the mechanism will compute the most frequent variants of behavior and belief in one's environment, and deliver a desire *to behave/believe like that*. And indeed, there are a whole raft of experiments and data in social psychology that demonstrate people's desire to conform, and its powerful effects on behavior (Myers 1993).

Another strategy that one can adopt is to imitate the behavior and assume the beliefs of those who are most successful. Here, too, mathematical modeling shows that there are many circumstances in which the strategy can be adaptive (Boyd and Richerson 1985). And this prompts Henrich and Gil-White (2001) to argue that we have evolved a species-unique motivational mechanism that makes us want to resemble the prestigious. We attempt to get close to those who are prestigious, and we try to learn from them and to model our behavior on theirs. This gives rise to a social dynamic in which those of us who grant prestige to others allow them certain privileges, in return for which they allow us to get close enough to learn from them. And here, too, there is an experimental literature supporting the claim that we have a strong tendency to want to be similar to those who are prestigious (Rogers 1995, Henrich and Gil-White 2001).

In addition to the numerous pillars of culture already identified, therefore, we probably need to add a couple more. These are less fundamental than the others. Much accumulation of culture would surely still be possible without them. But they nevertheless have a significant impact on the ways that cultural practices change and get transmitted over time.

CONCLUSION

I have argued that authors such as Bickerton (1995) and Tomasello (1999) are mistaken in claiming that there is just a single major new adaptation underlying human cognitive uniqueness. Even confining our attention to the uniquely human capacity for cumulative culture (and ignoring the many other respects in which humans are distinctive) there are a significant number of evolutionarily novel cognitive systems underlying it. In fact, there are at least eight. (1) There is a mindreading system for figuring out and ascribing mental states like belief. (2) There is a language acquisition device, focused on syntax in particular. (3) There are mechanisms that underlie the development of creative forms of thinking and behaving. (4) There are specialized motor control systems that enable fine-grained motor control and support the acquisition of novel skills. (5) There is a naïve physics system that enables us to comprehend simple forces and physical structures. (6) There is a mechanism that enables people to make inferences to the best among two or more competing explanations of some phenomenon. (7) There is a multi-component system for normative learning, reasoning, and motivation. And (8) there are motivational systems that support targeted social learning of cultural practices and cultural information. We can conclude, therefore, that the cognitive underpinnings of the distinctive human capacity for cumulative culture are multi-faceted, and by no means simple.

Acknowledgments

An early version of this chapter was delivered and discussed at the Evolution of Mind, Brain, and Culture workshop held at Penn Museum in September 2007. I am grateful to all the participants—and especially to my commentators, Thierry Chaminade and Philip Chase—for their advice and feedback. I am also grateful to two anonymous reviewers for their comments and suggestions on an earlier draft, and to Gary Hatfield and Holly Pittman for their advice.