PABLO QUINTANILLA | CARLA MANTILLA | PAOLA CÉPEDA
(editores)

COGNICIÓN SOCIAL Y LENGUAJE

La intersubjetividad en la evolución de la especie y en el desarrollo del niño

Capítulo 4

FONDO EDITORIAL
PONTIFICIA UNIVERSIDAD CATÓLICA DEL PERÚ

Cognición social y lenguaje
La intersubjetividad en la evolución de la especie y en el desarrollo del niño
Pablo Quintanilla, Carla Mantilla y Paola Cépeda (editores)

Av. Universitaria 1801, Lima 32, Perú
Teléfono: (51 1) 626-2650
Fax: (51 1) 626-2913
feditor@pucp.edu.pe
www.pucp.edu.pe/publicaciones

Diseño, diagramación, corrección de estilo
y cuidado de la edición: Fondo Editorial PUCP

Primera edición: mayo de 2014
Tiraje: 500 ejemplares

Hecho el Depósito Legal en la Biblioteca Nacional del Perú Nº 2014-07083
ISBN: 978-612-4146-80-0
Registro del Proyecto Editorial: 11501361400359

Impreso en Tarea Asociación Gráfica Educativa
Pasaje María Auxiliadora 156, Lima 5, Perú

The fragmentation of reasoning

Peter Carruthers
University of Maryland

1. Introduction

Scientists who study human reasoning across a range of cognitive domains have increasingly converged on the idea that there are two distinct systems (or types of system) involved. These domains include learning (Berry & Dienes, 1993; Reber, 1993), conditional and probabilistic reasoning (Evans & Over, 1996; Sloman, 1996, 2002; Stanovich, 1999), decision making (Kahneman & Frederick, 2002; Kahneman, 2003), and social cognition of various sorts (Petty & Cacioppo, 1986; Chaike and others, 1989; Wilson and others, 2000). Although terminology has differed, many now use the labels *System 1* and *System 2* to mark the intended distinction. System 1 is supposed to be fast and unconscious in its operations, issuing in intuitively compelling answers to learning or reasoning problems in ways that subjects themselves have no access to. System 2, in contrast, is supposed to be slow and conscious in its operations, and is engaged whenever we are induced to tackle reasoning tasks in a reflective manner. Many theorists now accept that System 1 is really a *set* of systems, arranged in parallel, while believing that System 2 is a single serially-operating ability.

Moreover, System 1 is generally thought to be unchangeable in its basic operations, to be universal amongst humans, and to be shared (at least in significant part) with other species of animal. It cannot be directly influenced by verbal instruction, and it operates independently of the subject's explicit beliefs. In addition, its operations are thought to be either associative or heuristic in character (or both), epitomized in the phrase «quick and dirty». System 2, in contrast, is thought to be uniquely human, to be malleable, and to differ significantly between people, varying both by individual and by culture. It can be influenced by verbal instruction and can be guided, at least in part, by the subject's beliefs. Furthermore, System 2 is supposed to embody some sort of normatively correct competence (at least to some significant degree, and presumably following appropriate training). The properties generally attributed to the two systems can be seen summarized in figure 1[1].

Figure 1. The properties of the two proposed systems

System 1	System 2
Unconscious/intuitive	Conscious/reflective
Fast	Slow
Parallel, high capacity	Serial, limited capacity
Automatic	Controlled
A set of systems	A single system
Not easily altered	Malleable
Universal amongst humans	Variable (by culture and by individual)
Mostly shared with other animals	Uniquely human
Impervious to verbal instruction	Responsive to verbal instruction
Independent of normative beliefs	Influenced by normative beliefs
Associative and/or heuristic-based	Rule-based and/or embodying normative competence

I shall argue that there is, indeed, a real, scientifically valid, distinction between a set of intuitive, unconsciously operating reasoning systems, on the one hand, and a reflective system whose operations are partly conscious, on the other. But I shall argue that this division fails to line up with many of the other properties generally associated with Systems 1 and 2. In particular, some intuitive systems can be slow, some can be controlled, and some can approach the highest normative standards; so the moniker «quick and dirty» is certainly inappropriate when applied to intuitive reasoning as such. Nor is it true that reflective reasoning always leads to improvement. On the contrary, in some contexts reflection leads to *worse* performance, and there are some tasks where reliance on intuitive reasoning is best. Moreover, reflective reasoning can also employ heuristics. Indeed, although there is (I shall argue) a fixed architecture underlying our reflective capacities, the latter can employ a hodge-podge of different procedures and abilities. In consequence, then, the distinction between System 1 and System 2 processes is not scientifically valid, failing to mark out any natural division within the mind.

[1] The properties listed in figure 1 have been culled from many different sources. Perhaps only a few researchers would accept all of them. But all (until very recently) would accept most. (Some fresh proposals for characterizing the System 1 / System 2 distinction will be considered in section 4.)

Here is how the paper will proceed. Section 2 will outline and motivate an architecture for reflective reasoning, for the most part following Carruthers (2006). On this account, reflective reasoning is real, but it is largely realized in the cyclical operations of unconscious, intuitive, processes (including the set of systems normally regarded as belonging to System 1), rather than existing alongside of the latter. I shall explain how we have independent evidence of the various components of the reflective system, and how most of the components either already exist in non-human animals or are developments of precursors that exist in animals. I shall then go on, in section 3, to argue that despite the reality of the distinction between intuitive and reflective reasoning processes, this does *not* amount to a vindication of anything resembling the full System 1/System 2 distinction as outlined above. On the contrary, I shall show that the latter distinction fails to mark out anything that deserves to be regarded as a *natural kind*. Finally, in section 4, I shall compare the thesis defended in this article with some other recent proposals by leading figures in the field for reconceptualizing the System 1/System 2 distinction, put forward by Evans (2009), Frankish (2009) and Stanovich (2009), in particular.

2. An architecture for reflective reasoning

This section provides an explanation of the distinction between intuitive and reflective processes, and shows how the latter can be partly realized in cyclical operations of the former.

2.1. The puzzle of reflective reason

We begin with some *prima facie* puzzles about the postulated reflective reasoning system. The latter is presumably meant to be a pan-human capacity of some kind; and yet at the same time its operations are dependent upon learning and vary a great deal by culture and individual. How can this be? One possible answer is that there is an innately channeled reasoning system of some sort which is then molded and completed on the basis of experience. Compare the human motor system. This, too, is pan-human in its architecture, and presumably contains a significant amount of innate structure (think of sneezing and yawning, for the clearest examples, but innately channeled behaviors probably also include walking and talking). Yet it is also built to be a learning system, in which new motor modules are constructed each time someone acquires a novel skill (Ghahramani & Wolpert, 1997; Manoel and others, 2002). Hence the mature state of the motor system will likewise vary greatly by culture and by individual, depending on the skills that the people in question possess.

What is somewhat more problematic is to understand how the reflective reasoning system can be guided by verbal instruction, and how it can be partly dependent upon people's *beliefs* about how they should reason. For the system in question is a cognitive one, issuing in transitions between one set of belief states and another, or generating decisions from a given set of beliefs and goals. And these cognitive transitions are presumably realized in some computational process or other. But how can verbal instruction and normative beliefs have the power to reach into the «innards» of this system, altering the state-transitions that are employed? How can a subject's beliefs have the capacity to re-write the computational algorithms that are used? Since verbal instructions and normative beliefs have their paradigmatic effects on *action*, it looks as if the reflective reasoning system must somehow be action-dependent; and it must likewise be capable of being intentionally controlled. Indeed, this is what I shall shortly argue.

The deepest difficulty for anyone wishing to defend the reality of a distinction between intuitive and reflective reasoning processes, however, is to understand how the two are related to one another, especially when viewed from an evolutionary perspective. For each seems to replicate, to a significant degree, the functionality of the other, and is concerned with the same types of task. Each issues in new beliefs on the basis of evidence, and each issues in decisions about what to do when faced with choices. How can this be? How could the (relatively short) period of hominid evolution have issued in such wholesale change and duplication? And if there were pressures for improvements in reasoning and decision-making (such as an increased need for cognitive flexibility), then why didn't they lead to modifications in the existing systems, rather than to the construction of an entirely different sort of system *de novo*? For that is the way in which evolutionary processes generally work.

How one answers these questions will obviously depend quite a bit on what one thinks about the character of non-human cognitive processes. For the more the latter resemble distinctively-human forms of thinking and reasoning, the deeper the problem. And in contrast, if one thought that all non-human animal behavior could be explained as resulting from various forms of associative conditioning, then the puzzle would largely disappear. For the benefits of rule-governed, systematically structured, forms of thinking and reasoning are obvious, especially in respect of flexibility and opportunities for one-shot learning. And it might be argued that these evolved in the hominid lineage subsequent to the evolution of the human language faculty, which provided the sort of structured representational system necessary for rule-governed reasoning to make its appearance (Bickerton, 2009). These questions will be addressed in section 2.2.

2.2. Unreflective reasoning in non-human animals

Contrary to the proposed explanation of the evolutionary need for System 2 reasoning, cognitive processes involving one-shot learning and structured representational states are *rife* within the animal kingdom. Bees, for example, can extract information about the direction and distance of a food source from a single presentation of a waggle dance. And experiments involving bee navigation demonstrate that they use both vector information (deriving either from their own dead reckoning computations or the dances of other bees) and landmark information gained during their own exploratory flights from the nest before they begin their lives as foragers (Menzel and others, 2005; De Marco & Menzel, 2008). They can use this information flexibly in the service of multiple goals (searching for nectar, pollen, water, etcetera, returning to the hive, or dancing for other bees), suggesting very strongly that their decision making involves computations over structured representational states (Carruthers, 2006).

Moreover, Gallistel has demonstrated that conditioning behavior itself is best explained in rule-governed computational terms, rather than in terms of associative strengths (Gallistel & Gibbon, 2001; Gallistel & King, 2009). He points out there are many well-known conditioning phenomena that are extremely puzzling from an associationist perspective, but that fall out quite naturally from a computational account. To give just a single example: the number of reinforcements that are necessary for an animal to acquire an intended behavior is unaffected by mixing *un*reinforced trials into the learning process. One set of animals might be trained on a 1:1 schedule: these animals receive a reward every time that they respond when the stimulus is present. But another set of animals might be trained on a 10:1 schedule: here the animals only receive a reward once in every ten trials that they respond when the stimulus is present. Still it will, on average, take both sets of animals the same number of rewarded trials to acquire the behavior. It will take the second set of animals *longer* to acquire the behavior, of course. If it takes both sets of animals 40 rewarded trials to acquire the behavior, then the first set might learn it in 40 trials, whereas the second set will take 400. But the number of *reinforcements* to acquisition is the same. This is extremely puzzling from the standpoint of an associationist. One would expect that all those times when the stimulus *isn't* paired with a reward ought to weaken the association between stimulus and reward, and hence make learning the intended behavior harder. But it doesn't, just as Gallistel's computational model predicts.

Indeed, Gallistel and others (2001) demonstrate that animals in conditioning experiments who are required to respond to randomly changing rates of reward are able to track changes in the rate of reward about as closely as it is theoretically possible to do. Thus both pigeons and rats on a variable reward schedule from two different alcoves will match their behavior to the changing rates of reward.

There is a lever in each alcove, each set on a random reward schedule of a given probability. But these probabilities themselves change at random intervals. It turns out that the animals respond to these changes *very* rapidly, closely tracking the random variations in the immediately preceding rates. They certainly aren't averaging over previous reinforcements, as associationist models would predict. On the contrary, the animals' performance comes very close to that of an ideal Bayesian reasoner. And the only model that can predict the animals' behavior is one that assumes that they are capable of calculating the ratio of the two most recent intervals between rewards from the two alcoves.

I conclude, therefore, that the evolutionary challenge to the distinction between System 1 and System 2 is sustained: since non-human animals engage in unreflective processes that can be both flexible and rule-governed, it is puzzling how a distinct system for reflective reasoning could ever have evolved. What I shall suggest in section 2.3. is that there is a sense in which it did not. Rather, reflective reasoning is realized in cycles of operation of unreflective forms of cognition, building upon and exapting the resources provided by the latter.

2.3. The mental rehearsal of action

In outline, the proposal is that reflection operates like this: action-schemata are selected and activated, and are mentally rehearsed (with overt action suppressed); this gives rise to conscious images which are globally broadcast (in the manner of Baars, 1988) and thus made available as input to the full suite of intuitive systems; the latter draw inferences from them, activate relevant memories, and issue in emotional reactions; during decision making the somatic consequences of the latter are monitored (in the manner of Damasio, 1994) and motivations are adjusted up or down accordingly; and the result is a whole new cognitive and affective environment influencing the selection of the next mentally rehearsed action (or in some cases, issuing in overt action). We have robust evidence of each of the components appealed to in this account, and each is very likely present in non-human animals, as I shall now briefly explain.

The creation of «forward models» of the expected sensory consequences of activated action-schemata is now quite well understood, as is its original function in facilitating fast on-line correction of action as it unfolds (Wolpert & Ghahramani, 2000; Wolpert & Flanagan, 2001; Wolpert and others, 2003). When attended to, these forward models issue in conscious motor imagery, as well as imagery of other sorts; and they can also do so when actions are mentally rehearsed with overt performance suppressed (Jeannerod, 2006). The result is sequences of motor images, visual images, or auditory images (often in the form of so-called «inner speech», when the actions rehearsed are speech actions), which serve as the conscious components of reflective thought.

It is also quite widely agreed that «global broadcasting» underlies the conscious status of conscious experiences and images (Baars, 2002; Dehaene & Naccache, 2001; Dehaene and others, 2003; Baars and others, 2003). Attended perceptual and imagistic representations are made accessible as input to a wide range of other cognitive systems for drawing inferences, for evoking and forming memories, for decision making, and for generating emotional and motivational responses. Moreover, there is evidence that motor cortex is active in the creation and transformation of visual images (Kosslyn, 1994; Ganis and others, 2000; Richter and others, 2000; Kosslyn and others, 2001; Lamm and others, 2001), and we also know that during episodes of inner speech not only the language comprehension areas of the brain but also language *production* areas and associated regions of motor cortex are active (Paulescu and others, 1993; Shergill and others, 2002). So we have evidence that activations of motor schemata are used to drive the sequences of conscious images found in reflective thinking. In addition, Damasio and colleagues have amassed a great deal of evidence of the crucial role that monitoring of emotional reactions to globally broadcast images plays in normal human decision-making (1994, 2003).

There is no reason to think that any of the components appealed to in this account is uniquely human (with the exception of speech and inner speech, of course). Indeed, there is some reason to think that apes might occasionally engage in reflective decision-making involving mental rehearsals of action, issuing in so-called «insight» phenomena (Carruthers, 2006). Moreover, we can be confident that by the time of *Homo ergaster*, some 1.4 million years ago, our ancestors were regularly using mental rehearsal of action in their reasoning and decision making (2006). For we know that there is no other way of producing the symmetrical hand-axes and blades of that era, using variable and unpredictable materials (Gowlett, 1984; Pelegrin, 1993; Mithen, 1996; Schlanger, 1996; Wynn, 2000). Stone knappers need to be able to plan several strikes ahead, in each case visualizing both the intended blow and its anticipated effects.

2.4. Reflection explained

Note that the above account has the resources to explain a significant number of the properties generally attributed to System 2, while also avoiding the puzzles about its existence raised in section 2.1. Because globally broadcast images are conscious, that element of each cycle of mentally rehearsed action will also be conscious (while the cognitive activity that precedes and follows the broadcast image will generally be *un*conscious). And because mental rehearsal activates and co-opts the resources of the various intuitive reasoning systems, its overall operations are likely to be significantly slower than most of the latter. Nor is there any special difficulty in explaining how reflective reasoning could have evolved. For rather than existing alongside of intuitive reasoning systems while performing many of the same functions, reflection is partly realized in cycles of operation of the latter, utilizing pre-existing mechanisms and capacities. All that had to evolve were (1) a language system and (2) a disposition to engage in mental rehearsals of action on a routine basis.

Moreover, because action selection in general is under intentional control and can be influenced by normative belief and verbal instruction, so can the operations of the described reflective system. We can *choose* to engage in mental rehearsal (often unconsciously), just as we choose to engage in any other form of action. And just as with other forms of action, some sequences of rehearsal can be produced smoothly and automatically, resulting from previous practice (think, here, of doing an addition sum in your head). Others can be guided by beliefs about how one *should* reason, sometimes by activating a stored memory of a previous instruction (when faced with a version of the Wason selection task, for example, I might rehearse the sentence, «In order to evaluate a conditional, I should look for cases where the antecedent is true and the consequent false», or I might form a mental picture the standard truth-table for the conditional). And of course with each iteration of mentally rehearsed action the various System 1 systems that «consume» the globally broadcast images become active, sometimes producing an output that contains or contributes towards the solution.

3. Against the System 1/System 2 distinction

The ideas presented in section 2 (when suitably worked out and developed; see Carruthers, 2006, 2009, for more in this direction) appear sufficient to establish the reality of the distinction between intuitive and reflective reasoning, while explaining how the latter is realized in cycles of mental rehearsal. The present section will argue that this is not, however, a vindication of the System 1/System 2 distinction as it has normally been understood. Quite the contrary. This will be demonstrated, in part, via examination of some recent data that fall outside the normal range that reasoning researchers might typically consider.

3.1. Heuristics can be «rational»

One powerful critique of the System 1/System 2 distinction is implicit in the «simple heuristics» research program of Gigerenzer and others (1999)[2]. For they challenge the

[2] It should be noted, however, that some reasoning researchers have always understood the System 1/ System 2 distinction in such a way as to leave room for the point developed in this section. Thus, Evans and Over (1996), in particular, distinguish between two different kinds of rationality that are applicable

assumption that intuitive processes are always «irrational». Although these processes might fail to comply with the normative standards laid down by philosophers and logicians, they argue that standards of rationality need to be relativized to the cognitive powers and life-demands of those who employ them. Human beings have to form beliefs and take decisions in real time (often seconds or minutes; rarely longer), and they have to do so given limited and noisy information, and with severely limited computational powers. In such circumstances compliance with rational norms might be impossible. And indeed, when people *try* to take decisions in a way that complies with normative standards —as did Darwin when attempting to decide whether or not to get married, forming a weighted list of the pros and cons— they often fail, frequently overturning the results of their own calculations in favor of a «gut feeling» (Gigerenzer, 2007). Moreover, when heuristic methods are pitched against «normatively correct» ones, they often prove remarkably successful, and are frequently more robust than the latter in generalizing successfully to new environments (Gigerenzer and others, 1999).

Everyone allows, of course, that heuristic reasoning methods can lead to errors, and that humans often think and act in ways that are downright foolish. But instead of focusing on the distinction between «error-prone» heuristics and a supposed set of ideally rational norms, it is more fruitful to investigate the circumstances in which a given heuristic works well and those in which it doesn't. This has issued in the concept of «ecological rationality» (Gigerenzer, 2000), which is the idea that particular reasoning heuristics will be well-adapted to certain cognitive environments but not others.

It is worth emphasizing, however, that proponents of the simple heuristics program need not deny the distinction between intuitive and reflective reasoning. And they had better *not* do so, for there is no doubt that humans often engage in conscious forms of reflection and that this can sometimes lead to improved outcomes. What they insist on (and rightly so) is just that the intuitive/reflective distinction fails to line up with the irrational/rational distinction (at least when the latter is properly naturalized).

3.2. Intuitive reasoning can be «slow»

Recall that System 1 processes are supposed to be *fast* by comparison with reflective forms of reasoning and decision making. No doubt this is true of many intuitive systems. But it surely isn't true of all. Consider the processes that issue in romantic love, or romantic attachment.

---

to the two systems, one of which is ecological rationality (leading to successful outcomes in many circumstances) whereas the other is normative rationality.

Although some people believe in love at first sight, there is no doubt that in general the process takes considerably longer —often days or weeks, and sometimes months or years. Moreover, although people sometimes reflectively weigh up the good and bad points of a prospective partner, this certainly doesn't always happen, and it is doubtful how effective it is when it does. Yet there is good reason to regard the process in question as one of reasoning. For psychologists have documented the multiple cues that influence romantic attachment, which are gathered over time and appear to be integrated with one another somehow to issue in an eventual intuition of the form, «This person is for me». These range from sensitivity to cues of kindliness (such as being nice to dogs and children) and intelligence (such as displaying a sense of humor), to the influence of physical attractiveness (including indicators of youthfulness in women and bodily symmetry in both sexes), to the unconscious detection of pheromones, and chemical information about the other person's immune system obtained from saliva through kissing (Barrett and others, 2002; Buss, 2005). And in addition, of course (as readers of Jane Austen will be well aware), our own level of attachment is strongly influenced by signals of interest from the prospective partner.

A quite different form of intuitive reasoning that is slow rather than fast is popularly known as «sleeping on it». It is a familiar occurrence in daily life that one's reflective reasoning about some problem has become «stuck» —one is unable to see a way to a solution—. So, one lays the problem aside, either literally going to sleep for the night or occupying oneself with other tasks. But then sometimes after an interval the solution suddenly and unexpectedly emerges into consciousness. It is natural to think that one must have continued reasoning about the problem, unconsciously, during the interim. Admittedly, this isn't the only interpretation possible. For it may be that the reason one had got stuck with the problem in the first place resulted from adopting an inappropriate mental «set», which framed and constrained one's conscious reflection. What happened during the time when one was asleep or doing other things, then, may have been that this initial set was *forgotten*. With that constraint out of the way, one is then able to approach the problem afresh, often resulting in success.

No doubt this sort of «set shifting» explanation is sometimes appropriate (Schooler & Melcher, 1995). But it can't explain those cases where a solution emerges fully-formed into consciousness without any prior attention to the problem. For if all that happens in the interval following one's previous attempts is that one forgets the initial line of approach to the problem that had led one astray, then, one would expect that one would thereafter have to think *reflectively* about the problem once again in order to achieve a solution. But this isn't always what happens. And indeed, we will see towards the end of section 3.3. that there is direct experimental evidence that unconscious reasoning often takes place in the intervals during which one's conscious mind is otherwise occupied.

3.3. Intuitive reasoning is better than reflection for some tasks

One set of data supporting the claim made in the title of this section is provided by Wilson and others (1993). They offered their subjects a choice between a number of different posters to take home to display on their walls. All subjects were allowed some time to study the posters on offer before making their choice. But one group were asked to articulate the positive and negative features of the various options, thus forcing them to be reflective in making their selection. The experimenters then followed up with each of the subjects a week later, asking them how happy they were with their new poster. Those who had chosen unreflectively expressed significantly greater satisfaction; and when asked how much money would be necessary to persuade them to sell their choice of poster back to the experimenters, those in the unreflective condition asked for double the amount demanded by those who had made their choice following explicit reflection.

This seems to be a case where unreflective choice is better than reflection, at least in the sense of leading to greater subjective satisfaction with the chosen object. Wilson and others (1993) hypothesize that the effect occurs because reflection will inevitably focus most attention on those good-making and bad-making features that are easily expressible in language, as well as perhaps narrowing attention to fewer features more generally. And the moral seems to be that when faced with a complex multi-faceted choice one should pay attention to as many features as one can, but one shouldn't attempt to make one's decision via conscious reflection and reasoning; rather, following study of the choices on offer one should «trust one's gut feeling».

Even more decisive evidence of the benefits of unconscious over conscious forms of decision making has been provided by Dijksterhuis and colleagues (Dijksterhuis, 2004; Dijksterhuis and others, 2006). In one set of experiments subjects were asked to make a complex choice under one of three conditions. For example, they might be asked to choose the best of four cars when provided with ten items of information about each (cost, gas mileage, and so on). The information was arranged so that one of the items was best (with say eight positive features and two negative ones) and another was worst (with eight negative features and two positive ones), whereas the other two choices were of intermediate quality (in some experiments the good-making and bad-making features were calibrated against subjects' own evaluations made at another time).

In one condition subjects were asked to make an immediate choice having been presented with all the information; in a second condition they were told to reflect for a few minutes before making their choice; and in the third condition they were provided with a demanding distracter task for the same amount of time, which would have made conscious reflection impossible. What the experimenters found is that intuitive reasoning is significantly better than immediate choice, suggesting that the process in question is fairly slow. But they also found that intuitive reasoning is significantly better than conscious reflection (there was no qualitative difference between immediate and reflective choosing, suggesting that in these circumstances reflection brings no benefit).

Indeed, Dijksterhuis and colleagues have been able to demonstrate that successful intuitive reasoning is goal-dependent (Bos and others, 2008)[3]. They have thereby demonstrated that such reasoning isn't always automatic, in contrast with almost every theorist who writes about the System 1/System 2 distinction (Evans, 2009; Stanovich, 2009). In one experiment subjects studied complex information about four different makes of car, as before. One set of subjects was asked to make a choice after a period of reflection, while another group was told that they would make their choice after completing an attentionally-demanding task, as previously. But now the third group was told that the relevant phase of the experiment was over, hence blocking the formation of any goal of choosing between the four cars. But after they had completed the same attention-occupying task as the second group, they were in fact presented with the same choice. There were no significant differences between the first and third groups. Only the subjects who had the opportunity to reason intuitively about their choice and *also* had the goal of choosing a car did better than those who chose reflectively. (Note that this result is inconsistent with a «set shifting» explanation, since the third group would have had the same opportunity to forget their initial characterization of the problem as did the second group.)

Even more remarkably, Bos and others (2008) were able to demonstrate that intuitive reasoning is guided by quite specific goals. They asked another set of subjects to study complex information about two kinds of object (cars and potential room-mates). One group of subjects was then told, before undertaking an attentionally-demanding task, that they would be choosing between the cars, whereas the other group was told that they would be choosing between room-mates. But then both groups were actually thereafter asked to choose the best car *and* the best room-mate. Members of both groups showed the benefits of unconscious reflection, but *only* with respect to the items that they had been given the goal of selecting between.

3.4. Intuitive reason can approximate to the highest normative standards

Not only can intuitive reasoning be markedly better than reflection, in some circumstances, but some unconscious processes can be highly optimal, approximating

[3] They have also shown that intuitive reasoning issues in better organization of memory than does conscious reflection (Dijksterhuis & Nordgren, 2006).

to the most stringent normative standards. One example of this has already been discussed in section 2.2. This is the data from Gallistel and others (2001) concerning rate estimation in rats and pigeons, showing that the animals were tracking the randomly changing rates about as closely as it is theoretically possible to do (I presume that these animals were making their estimates intuitively, and weren't engaging in conscious reflection!). But an even more dramatic example is provided by Balci and others (2009). They tested swift and intuitive assessments of risk, using similar experiments in both humans and mice, with very similar results (I shall discuss only the human data here, however, since our main topic is dual systems of reasoning in human beings).

The human subjects were set the task of «capturing» an object in one of two positions on a computer screen for a reward. There were two types of trial, short latency and long latency, the probability of each of which varied from one series of trials to the next. If the trial was a short one, the target could be captured in the left-hand position within two seconds of the stimulus onset; if the trial was a long one, the target could be captured in the right-hand position during the third second. Subjects were therefore required to estimate the optimum time to switch from the short-latency strategy to the long-latency strategy. This estimate depends upon two factors. One is the objective chance that the interval would be either short or long (this was set by the experimenters in each series of trials). The other is the accuracy of each subject's own estimate of elapsed time (which varies from person to person, but which is normally in the region of ±15%). Balci and others (2009) were able to compute the optimum switch time for each subject, combining both sets of probabilities. This was then compared with actual performance. Subjects came within 98% of optimum performance. Moreover, very little learning was involved (either for humans or for the mice). In most series of trials, subjects were just as successful during the first tenth or the first quarter of the series as they were during the final tenth or the final quarter. This is a remarkable result, especially given the repeated finding that humans are so poor at reasoning about probabilities in explicit tasks.

A very different kind of normatively correct intuitive reasoning system is proposed by Sperber and Mercier (2009), which is designed for public argumentation. They initially motivate the need for such a system on evolutionary grounds, having to do with the epistemic vigilance that is necessary to sustain successful systems of communication. They point out that speech is a cooperative activity, and like all forms of cooperation it is vulnerable to being parasitized and undermined by cheaters and free-riders. People therefore need to be alert to the possibility of being misled, and should have a healthy skepticism about what others tell them, especially in circumstances where it is plain that conflicts of interest are involved.

This creates pressure, in turn, for would-be communicators to lay out reasons why the communication should be accepted by their audience; which in turn creates pressure on hearers to evaluate those reasons and to formulate counter-reasons where necessary. So we should predict the emergence of a specialized argumentation system to facilitate these tasks[4].

Sperber and Mercier (2009) then amass a broad body of evidence in support of their proposal. Part of this consists in an extensive set of studies by a variety of researchers demonstrating that people are much better at reasoning in argumentative contexts (and this isn't just a product of greater motivation, since paying subjects to reason well in individual reasoning tasks has little effect). For example, although subjects are notoriously bad at employing modus tollens arguments in standard paper-and-pencil tasks, when people want to attack the views of others they are actually very good at employing such arguments (Pennington & Hastie, 1993). Even more striking, Moshman and Geil (1998) had one group of subjects attempt Wason conditional reasoning tasks individually. Only 9% succeeded. Another set of subjects were put into groups and asked to solve the same tasks, and a massive 70% of the groups succeeded. It seems that subjects have the capacity to recognize correct solutions when proposed by others. Indeed, it seems that such solutions can emerge out of a process of discussion. For when groups were formed using only subjects who had previously failed at the task on an individual basis, 30% of those groups were nevertheless able to come up with the correct solution. Moreover, discussion involving dissent has been shown to be a crucial determinant of group performance (Schulz-Hardt and others, 2006).

In addition, and in contrast with the extensive studies that demonstrate poor reasoning in subjects when working on tasks individually, naturalistic studies of persuasion have found that people in argumentative contexts are remarkably good at distinguishing good arguments from bad ones and at generating good arguments for themselves (Petty & Wegener, 1998; Neuman and others, 2006). Especially striking is the fact that even young children can engage successfully in argumentation (Stein & Albro, 2001) and can spot argumentative fallacies (Baum and others, 2007).

---

[4] Confusingly, Sperber and Mercier themselves describe the argumentation system as a *reflective* one, on the grounds that it is often engaged when we explicitly and consciously evaluate the arguments of others or propose counter-arguments ourselves. This may be so. But since the internal operations of the system are always unconscious it is plain, I think, that it should be classified as an intuitive system. For *all* intuitive systems are capable of issuing in conscious outputs, and many may help to guide cycles of conscious reflection. Indeed, when presented with someone's argument, the argumentation system will swiftly and unconsciously generate an intuition about the strength of that argument, together with intuitions about its likely weak spots. This fits the normal characterization of an intuitive system.

3.5. Reflection can employ heuristics

So far this section has built a cumulative case that it is a mistake to describe intuitive reasoning systems as «quick, automatic, and dirty». For some of them can be slow; some can be controlled; some are more successful than reflective reasoning when dealing with the same types of task; and some can approximate to the highest normative standards. I now turn attention to reflective kinds of reasoning, questioning the extent to which such reasoning processes possess the properties generally attributed to System 2. The first point I want to make is that reflective reasoning, just as well as intuitive reasoning, can have a heuristic character.

While reflection can sometimes involve appeal to valid norms of reasoning, it can also involve the use of heuristics (this point is now fully accepted by some two-systems theorists, especially recently; see Frankish [2004], Evans [2009], and Stanovich [2009]). Some of these can be good and useful, and some can be bad. The «sleep on it» heuristic, which is often consciously and reflectively employed by people in our culture, might be a good example of the first sort. Moreover, if Sperber and Mercier (2009) are right about the existence of an intuitive argumentation system, then one might predict that a useful heuristic to employ when engaged in private reflective reasoning about some problem would be to imagine that the context is an argumentative one (for example, by engaging an imaginary opponent). For by simulating argument, one will provide the cues necessary for the argumentation system to become active. And then each cycle of mental rehearsal will be evaluated by that system, thus facilitating the transition to the next step in one's reasoning in a normatively correct manner.

No doubt some conscious reasoning heuristics may result from individual learning and theorizing. But many will be culturally acquired (the «sleep on it» heuristic is surely an example of the latter sort). And a moment's reflection is sufficient to demonstrate that many culturally sanctioned reasoning heuristics have been quite bad. Consider the common practice in many cultures of consulting an oracle, such as the entrails of a newly killed chicken, when making an important decision. It is plain to us that decisions made on this basis will be at chance. And it is equally obvious that people in such a culture would be much better off not reflecting about the decision at all, given that reflection will call to mind the «look at the entrails» heuristic, which is then likely to be followed[5].

[5] Note there is also a quite different sense in which reflection can involve heuristics. For insofar as reflection often activates and depends upon intuitive reasoning processes, then to that extent it will involve any heuristics that are employed within those processes. What I have in mind in section 3.5., in contrast, are heuristics that are consciously activated and deployed during reflection.

3.6. The fragmentation of reflective reasoning

The present section argues that reflective reasoning processes are by no means unitary in character. On the contrary, some can be based upon normative beliefs and can lead to good performance if the beliefs in question are correct (or to poor performance if the beliefs are incorrect). Some are skill-based, mediated by learned and internalized action sequences. And some depend upon the modular systems that mediate transitions within the architecture. I shall discuss each in turn. But I should stress that I don't mean to imply that these three types of process are mutually exclusive. On the contrary, some instances of reflective reasoning will involve a combination of at least two of them, and some will involve all three.

Recall that reflection, on the account provided in section 2, is action-based. It can therefore be influenced in any of the ways that action can. And in particular, it can be guided by the subject's beliefs (especially normative beliefs about how one *should* reason). One effect of taking a course in logic or probability theory, for example, is that one forms a number of beliefs about the ways in which reasoning should be conducted. As we noted earlier, when asked to evaluate a conditional one might pause to reflect, perhaps asking oneself the explicit question, «How does one evaluate a conditional?». When globally broadcast, this might call to mind (whether consciously or unconsciously) the acquired belief, «A conditional is only false when the antecedent is true and the consequent false». This belief provides the information needed to guide the next stage of reflection, leading one to think, in the context of a Wason selection task, for example, «So I should turn over the P and the $\sim Q$ cards». Likewise when asked to assign probabilities to a series of statements, pausing to reflect might call to mind the belief, «A conjunction cannot be more probable than one of its conjuncts». This might then be sufficient to override some of one's initial intuitions with respect to the cases.

These are examples where reflection will lead to improved performance, because the beliefs accessed via reflection are true. But of course training can also issue in beliefs that are false, either because of errors in the transmission process, or because the teachers themselves have false beliefs about how one should reason. The history of the philosophy of science, for example, is littered with cases where philosophers have endorsed normative standards for the conduct of scientific reasoning —including naïve inductivism and Popperian falsificationism— which would have seriously impeded scientific progress had they been adopted. Fortunately for us, most scientists received their training from other scientists, and not from philosophers of science, so the normative beliefs in question did little damage (for an argument that significant damage was done in psychology, however, especially in studies of children's scientific reasoning, see Koslowski, 1996).

Since reflection is action-based, it can also involve habitual sequences that have been synergized together, in the same way that physical skills can involve such sequences. These can be acquired through training and practice, often in the context of formal schooling. This is surely true of the sequences of mental rehearsal involved when doing multiplication or long division sums, for example. But reasoning skills can also be acquired through imitation of the overt reasoning of other people. Much of the learning that goes on in universities is probably of this sort. Philosophy instructors exhibit, in their reasoning during class, patterns of thought that they hope their students will copy. And likewise faculty involved in scientific lab meetings will be exhibiting patterns of reasoning and problem solving to the students who are present.

The most straightforward way for reflection to operate, however, is by harnessing and engaging the resources of the intuitive systems that are the consumers of globally broadcast events (Shanahan & Baars, 2005). For example, by asking myself a question I might evoke information contained in one of the intuitive systems that I didn't know that I had, and that might not otherwise have become activated (Frankish, 2004). Another instance of the same phenomenon was discussed in section 3.5., on the assumption that Sperber and Mercier (2009) are correct about the existence of an intuitive argumentation system. Reflectively simulating an argument with an imagined opponent will cause this system will become active, thereby leading to an improvement in one's reasoning.

Moreover, the deployment of intuitive systems through cycles of reflection is central to Damasio's (1994, 2003) account of the operations of reflective practical reason. By mentally rehearsing a proposed action, images of that action become globally broadcast to the full suite of intuitive systems. Some of these elaborate the image to include its immediate consequences. But the results are also received as input by the emotional and motivational systems which respond accordingly, issuing in a set of bodily/emotional reactions. These are monitored by the subject, adjusting the motivation towards the proposed action upwards or downwards. We know from cases of people with frontal brain damage who have lost this monitoring capacity (but whose *theoretical* reasoning about action and the costs and benefits of different options can be fully intact) that this sort of practical reflection is crucial for leading a successful life (Damasio, 1994). But it is by no means infallible. On the contrary, it is subject to a number of fallacies and biases, such as overestimating the pain that we will feel through the loss of some major goal, for instance failure to be granted tenure (Gilbert, 2005).

4. Comparisons and conclusion

I conclude that there is no pair of reasoning systems (nor pair of system types) that lines up with all or most of the properties traditionally ascribed to System 1 and System 2. What really exists is a distinction between a set of intuitive systems and a reflective (mental rehearsal involving) architecture. This secures *some* of the properties of the original distinction (unconscious versus consciousness-involving, impervious to instruction versus malleable, and impenetrable by normative belief versus capable of being controlled by such beliefs). But some of the subsystems or processes within each can display normatively correct reasoning, each often makes use of heuristics, some intuitive systems can be as slow or slower than reflection, and some are goal-dependent rather than automatic. Moreover, some intuitive systems are better than reflective processes on the same task. In addition, although all reflective reasoning employs a common *architecture* of mental rehearsal and global broadcast, very different reasoning strategies and processes can be employed within that architecture. Some of these can be heuristic based, and some can be quite unreliable. Indeed, the very same architecture is also employed for *non*-reasoning processes, such as personal fantasy and mere «idle thinking».

I suggest, therefore, that while the distinction between intuitive and reflective processes is real, the distinction between System 1 and System 2 reasoning systems is not. Reasoning researchers should abandon the latter conceptual framework, and should instead focus on investigating the nature and variety of processes of both intuitive and reflective kinds. However, others, too, have recently suggested a reconceptualization of the System 1/System 2 distinction. I shall briefly consider the main two alternatives, contrasting them with my own proposal before concluding.

Evans (2009) now argues that System 1 processes fall into two distinct types. There are those that *compete with* System 2 processes, operating alongside them and issuing in intuitions, judgments, or decisions. And there are those that *influence* the operations of System 2, or *provide input* to System 2, while not generating the same sorts of outcomes in their own right. He also argues that we need to recognize a third system, which receives input from Systems 1 and 2 and makes decisions between them[6]. I have no quarrel with these suggestions. But in respect of System 2 itself, Evans argues that it can be identified with the set of processes that involve working memory (which can otherwise differ from one another in many of the respects outlined in section 3). Although this suggestion contains an important element of truth, it doesn't quite capture the distinction that we need. For anyone answering

[6] Stanovich (2009), too, draws a related but different distinction within System 2 processes, between the computational powers of the conscious system and its reflective beliefs and dispositions.

a question in an intuitive, *un*reflective, way will nevertheless be employing working memory to process the task instructions and maintain the ensuing representations long enough for the intuitive systems to generate an answer. What is distinctive of reflective processes is that *actions* should be mentally rehearsed, thereby issuing in appropriate sequences of imagery within working memory. In short, reflection is an *activity*, whereas working memory, as such, is not. But in any case (since processes that involve working memory are a heterogeneous lot) Evans' recent position is consistent with one of the main claims argued for here, namely that there is no unified character to the set of processes labeled «System 2».

Frankish (2009), in contrast, characterizes the key difference between System 1 and System 2 in terms of *control*, as does Stanovich (2009). But we have already had occasion to note in section 3.3. that some *intuitive* processes can be controlled as well, at least in the sense that they are goal-dependent. Moreover, not all reflective processes are controlled by the agent's beliefs and decisions, surely. For recall that reflection is *action*-based. Yet we know that actions can be triggered and guided independently of conceptual thought and decision-making. People will thoughtlessly mirror one another's behavior in social situations, for example (Jeannerod, 2006), and the mere sight of object affordances (such as the «to-be-grasped» shape of a hammer or coffee mug) will trigger activation of the appropriate motor schemata, which then need to be actively *suppressed* to prevent them from becoming acted upon (Shallice and others, 1989; Jacob & Jeannerod, 2003). Moreover, the work of Bargh and colleagues demonstrates that action schemata (such as walking in the manner of an old man) can be activated and executed by suitable conceptual priming without giving rise to any *belief* that the subject is an old man or to any corresponding intention (Bargh and others, 1996, 2001; Bargh, 2005). We should predict, then, that reflective processes can be just as automatic as overt actions can be. Moreover (and just as one might expect, on my account), in individuals suffering from obsessive-compulsive disorder, not only behavior but also reflective thought processes are undertaken compulsively, outside of the control of the agent (Leckman and others, 1997).

I conclude that human reasoning is best characterized in terms of a distinction between a varied set of intuitive systems, on the one hand, and a cognitive architecture that enables us to engage in reflective thinking and reasoning, on the other (with the latter operating in heterogeneous ways). If one of the goals of science is to discover what natural kinds there are in the world —in the sense of homeostatic property clusters with unifying causal etiologies (Boyd, 1989, 1991)— then cognitive scientists would be well-advised to abandon the System 1/System 2 conceptual framework. The human mind is messier and more fine-grained than that[7].

[7] I am grateful to Keith Frankish and Hugo Mercier for their comments on an earlier draft of this paper.

References

Baars, Bernard (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Baars, Bernard (2002). The Conscious Access Hypothesis: Origins and Recent Evidence. *Trends in Cognitive Science*, 6, 47-52.

Baars, Bernard, Thomas Ramsoy & Steven Laureys (2003). Brain, Consciousness, and the Observing Self. *Trends in Neurosciences*, 26, 671-675.

Balci, Fuat, David Freestone & Charles Gallistel (2009). Risk Assessment in Man and Mouse. *Proceedings of the National Academy of Sciences*, 106, 2459-2463.

Bargh, John (2005). Bypassing the Will: Toward Demystifying the Nonconscious Control of Social Behavior. In Ran R. Hassin, James S. Uleman and John A. Bargh (eds.), *The New Unconscious* (pp. 37-58). Oxford: Oxford University Press.

Bargh, John, Mark Chen & Lara Burrows (1996). Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype Priming on Action. *Journal of Personality and Social Psychology*, 71, 230-244.

Bargh, John, Peter Gollwitzer, Annette Lee-Chair, Kimberly Barndollar & Roman Troetschel (2001). The Automated Will: Nonconscious Activation and Pursuit of Behavioral Goals. *Journal of Personality and Social Psychology*, 81, 1014-1027.

Barrett, Louise, Robin Dunbar & John Lycett (2002). *Human Evolutionary Psychology*. New Jersey: Princeton University Press.

Baum, Laura, Judith Danovitch & Frank Keil (2007). Children's Sensitivity to Circular Explanations. *Journal of Experimental Child Psychology*, 100, 146-155.

Berry, Diane & Zoltan Dienes (1993). *Implicit Learning*. New Jersey: Erlbaum.

Bickerton, Derek (2009). *Adam's Tongue: How Humans Made Language, How Language Made Humans*. New York: Hill and Wang.

Bos, Maarten, Ap Dijksterhuis & Rick van Baaren (2008). On the Goal-Dependency of Unconscious Thought. *Journal of Experimental Social Psychology*, 44, 1114-1120.

Boyd, Richard (1989). What Realism Implies and What it Does Not. *Dialectica*, 43, 5-29.

Boyd, Richard (1991). Realism, Anti-foundationalism, and the Enthusiasm for Natural Kinds. *Philosophical Studies*, 61, 127-148.

Buss, David (ed.) (2005). *The Handbook of Evolutionary Psychology*. New Jersey: Wiley.

Carruthers, Peter (2006). *The Architecture of the Mind*. Oxford: Oxford University Press.

Carruthers, Peter (2009). An Architecture for Dual Reasoning. In Jonathan Evans and Keith Frankish (eds.), *In Two Minds: Dual Processes and Beyond* (pp. 109-127). Oxford: Oxford University Press.

Chaiken, Shelly, Akiva Liberman & Alice H. Eagly (1989). Heuristic and Systematic Processing within and Beyond the Persuasion Context. In James S. Uleman and John Bargh (eds.), *Unintended Thought* (pp. 212-252). New York: Guilford Press.

Damasio, Antonio (1994). *Descartes' Error*. London: Papermac.

Damasio, Antonio (2003). *Looking for Spinoza*. San Diego: Harcourt.

De Marco, Rodrigo & Randolf Menzel (2008). Learning and Memory in Communication and Navigation in Insects. In Randolf Menzel and John H. Byrne (eds.), *Learning Theory and Behavior* (pp. 477-498). Amsterdam: Elsevier.

Dehaene, Stanislas & Lionel Naccache (2001). Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework. *Cognition*, 79, 1-37.

Dehaene, Stanislas, Claire Sergent & Jean-Pierre Changeux (2003). A Neuronal Network Model Linking Subjective Reports and Objective Physiological Data during Conscious Perception. *Proceedings of the National Academy of Science*, 100, 8520-8525.

Dijksterhuis, Ap (2004). Think Different: The Merits of Unconscious thought in Preference Development and Decision Making. *Journal of Personality and Social Psychology*, 87, 586-598.

Dijksterhuis, Ap & Loran Nordgren (2006). A Theory of Unconscious Thought. *Perspectives on Psychological Science*, 1, 95-109.

Dijksterhuis, Ap, Maarten Bos, Loran Nordgren & Rick van Baaren (2006). On Making the Right Choice: The Deliberation-Without-Attention Effect. *Science*, *311*(5763), 1005-1007.

Evans, Jonathan (2009). How Many Dual-Processing Theories do We Need? One, Two, or Many? In Jonathan Evans and Keith Frankish (eds.), *In Two Minds: Dual Processes and Beyond* (pp. 33-88). Oxford: Oxford University Press.

Evans, Jonathan & David Over (1996). *Rationality and Reasoning*. Hove: Psychology Press.

Frankish, Keith (2004). *Mind and Supermind*. Cambridge: Cambridge University Press.

Frankish, Keith (2009). Systems and Levels: Dual-System Theories and the Personal-Subpersonal Distinction. In Jonathan Evans and Keith Frankish (eds.), *In Two Minds: Dual Processes and Beyond* (pp. 89-107). Oxford: Oxford University Press.

Gallistel, Charles & John Gibbon (2001). Time, Rate and Conditioning. *Psychological Review*, 108, 289-344.

Gallistel, Charles, Terence Mark, Adam King & Peter Latham (2001). The Rat Approximates an Ideal Detector of Rates of Reward: Implications for the Law of Effect. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 354-372.

Gallistel, Charles & Adam King (2009). *Memory and the Computational Brain*. Singapore: Wiley-Blackwell.

Ganis, Giorgio, Julian Keenan, Stephen Kosslyn & Alvaro Pascual-Leone (2000). Transcranial Magnetic Stimulation of Primary Motor Cortex Affects Mental Rotation. *Cerebral Cortex*, 10, 175-180.

Ghahramani, Zoubin & Daniel Wolpert (1997). Modular Decomposition in Visuomotor Learning. *Nature*, 386, 165-193.

Gigerenzer, Gerd (2000). *Adaptive Thinking*. New York: Oxford University Press.

Gigerenzer, Gerd (2007). *Gut Feelings*. New York: Viking Press.

Gigerenzer, Gerd, Peter Todd & The ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.

Gilbert, Daniel (2005). *Stumbling on Happiness*. New York: Vintage Books.

Gowlett, John (1984). Mental Abilities of Early Man: A Look at Some Hard Evidence. In Robert Foley (ed.), *Hominid Evolution and Community Ecology* (pp. 167-192). Nueva York: Academic Press.

Jacob, Pierre & Marc Jeannerod (2003). *Ways of Seeing*. New York: Oxford University Press.

Jeannerod, Marc (2006). *Motor Cognition*. New York: Oxford University Press.

Kahneman, Daniel (2003). A Perspective on Judgment and Choice: Mapping Bounded Rationality. *American Psychologist*, 58, 697-720.

Kahneman, Daniel & Shane Frederick (2002). Representativeness Revisited: Attribute Substitution in Intuitive Judgment. In Thomas Gilovich, Dale Griffin and Daniel Kahneman (eds.), *Heuristics and Biases* (pp. 49-81). Cambridge: Cambridge University Press.

Koslowski, Barbara (1996). *Theory and Evidence*. Cambridge: The MIT Press.

Kosslyn, Stephen (1994). *Image and Brain*. Cambridge: The MIT Press.

Kosslyn, Stephen, William Thompson, Mary Wraga & Nathaniel Alpert (2001). Imagining Rotation by Endogenous versus Exogenous Forces: Distinct Neural Mechanisms. *NeuroReport*, 12, 2519-2525.

Lamm, Claus, Christian Windtschberger, Ulrich Leodolter, Ewald Moser & Herbert Bauer (2001). Evidence for Premotor Cortex Activity during Dynamic Visuospatial Imagery from Single Trial Functional Magnetic Resonance Imaging and Event-Related Slow Cortical Potentials. *Neuroimage*, 14, 268-283.

Leckman, James and others (1997). Symptoms of Obsessive-Compulsive Disorder. *American Journal of Psychiatry,* 154, 911-917.

Manoel, Edison, Luciano Basso, Umberto Correa & Go Tani (2002). Modularity and Hierarchical Organization of Action Programs in Human Acquisition of Graphic Skills. *Neuroscience Letters*, 335, 83-86.

Menzel, Randolf and others (2005). Honey Bees Navigate According to a Map-Like Spatial Memory. *Proceedings of the National Academy of Sciences*, 102, 3040-3045.

Mithen, Steven (1996). *The Pre-History of the Mind*. London: Thames and Hudson.

Moshman, David & Molly Geil (1998). Collaborative Reasoning: Evidence for Collective Rationality. *Thinking and Reasoning*, 4, 231-248.

Neuman, Yair, Michael Weinstock & Amnon Glasner (2006). The Effect of Contextual Factors on the Judgment of Informal Reasoning Fallacies. *Quarterly Journal of Experimental Psychology*, 59, 411-425.

Paulescu, Eraldo, Chris D. Frith & Richard Frackowiak (1993). The Neural Correlates of the Verbal Component of Working Memory. *Nature*, 362, 342-345.

Pelegrin, Jacques (1993). A Framework for Analyzing Prehistoric Stone Tool Manufacture and a Tentative Application of Some Early Stone Industries. In Arlette Berthelet and Jean Chavaillon (eds.), *The Use of Tools by Human and Non-Human Primates* (pp. 302-314). New York: Oxford University Press.

Pennington, Nancy & Reid Hastie (1993). Reasoning in Explanation-Based Decision-Making. *Cognition*, *49*(1), 123-163.

Petty, Richard & John Cacioppo (1986). The Elaboration Likelihood Model of Persuasion. Leonard Berkowitz (ed.), *Advances in Experimental Social Psychology.* Volume 19. New York: Academic Press.

Petty, Richard & Duane Wegener (1998). Attitude Change: Multiple Roles for Persuasion Variables. In Daniel Gilbert, Susan Fiske and Gardner Lindzey (eds.), *The Handbook of Social Psychology* (pp. 323-390). New York: McGraw-Hill.

Reber, Arthur (1993). *Implicit Learning and Tacit Knowledge*. New York: Oxford University Press.

Richter, Wolfgang and others (2000). Motor Area Activity during Mental Rotation Studied by Time-Resolved Single-Trial (fMRI). *Journal of Cognitive Neuroscience*, *12*(2), 310-320.

Schlanger, Nathan (1996). Understanding Levallois: Lithic Technology and Cognitive Archaeology. *Cambridge Archaeological Journal*, *6*(2), 75-92.

Schooler, Jonathan & Joseph Melcher (1995). The Ineffability of Insight. In Steven Smith, Thomas Ward and Ronald Finke (eds.), *The Creative Cognition Approach* (pp. 97-133). Cambridge: The MIT Press.

Schulz-Hardt, Stefan, Felix Brodbeck, Andreas Mojzisch, Rudolph Kerschreiter & Dieter Frey (2006). Group Decision Making in Hidden Profile Situations: Dissent as a Facilitator for Decision Quality. *Journal of Personality and Social Psychology*, *91*(6), 1080-1093.

Shallice, Tim, Paul Burgess, Frederick Schon & Doreen Baxter (1989). The Origins of Utilization Behavior. *Brain*, *112*(6), 1587-1598.

Shanahan, Murray & Bernard Baars (2005). Applying Global Workspace Theory to the Frame Problem. *Cognition*, *98*(2), 157-176.

Shergill, Suckwinder and others (2002). Modulation of Activity in Temporal Cortex during Generation of Inner Speech. *Human Brain Mapping*, *16*(4), 219-227.

Sloman, Steven (1996). The Empirical Case for Two Systems of Reasoning. *Psychological Bulletin*, *119*(1), 3-22.

Sloman, Steven (2002). Two Systems of Reasoning. In Thomas Gilovich, Dale Griffin and Daniel Kahneman (eds.), *Heuristics and Biases* (pp. 379-396). Cambridge: Cambridge University Press.

Sperber, Dan & Hugo Mercier (2009). Reasoning as a Social Competence. In Jon Elster and Hélène Landemore (eds.), *Collective Wisdom* (pp. 368-392) Cambridge: The MIT Press.

Stanovich, Keith (1999). *Who is Rational? Studies of Individual Differences in Reasoning*. New Jersey: Erlbaum.

Stanovich, Keith (2009). Distinguishing the Reflective, Algorithmic, and Autonomous Minds: Is it Time for a Tri-Process Theory? In Jonathan Evans and Keith Frankish (eds.), *In Two Minds: Dual Processes and Beyond* (pp. 55-88). Oxford: Oxford University Press.

Stein, Nancy & Elizabeth Albro (2001). The Origins and Nature of Arguments. *Discourse Processes*, *32*(2-3), 113-133.

Wilson, Timothy, Samuel Lindsey & Tonya Schooler (2000). A Model of Dual Attitudes. *Psychological Review*, *107*(1), 101-126.

Wilson, Timothy and others (1993). Introspecting About Reasons Can Reduce Post-Choice Satisfaction. *Personality and Social Psychology Bulletin*, *19*(3), 331-339.

Wolpert, Daniel & Randall Flanagan (2001). Motor Prediction. *Current Biology*, *11*(18), 729-732.

Wolpert, Daniel & Zoubin Ghahramani (2000). Computational Principles of Movement Neuroscience. *Nature Neuroscience*, 3, 1212-1217.

Wolpert, Daniel, Kenji Doya & Mitsuo Kawato (2003). A Unifying Computational Framework for Motor Control and Social Interaction. *Philosophical Transactions of the Royal Society of London B*, *358*, 593-602.

Wynn, Thomas (2000). Symmetry and the Evolution of the Modular Linguistic Mind. In Peter Carruthers and Andrew Chamberlain (eds.), *Evolution and the Human Mind* (pp. 113-139). Cambridge: Cambridge University Press.