# Logic and Biology:

## Emotional inference and emotions in reasoning

**Ronald de Sousa**
**University of Toronto**

Penultimate text of paper published in J.E. Adler & L.J. Rips, editors, Reasoning; Studies of Human Inference and its Foundations Cambridge University Press 2008 . pp. 1002-1015

*The unrefined and sluggish mind*

*of Homo Javanensis*

*Could only treat of things concrete*

*and present to the senses*

W.V. Quine

Tropisms and transitions: some leading questions

Before planning and reasoning, there were tropisms. Tropisms have functions and make use of information detected, but they don't, I assume, involve any actual reasoning. Somewhere along the course of evolution, and at some time in any one of us on the way from zygote to adult, some forms of detection became beliefs, and some tropisms turned into reasoned desires. And at some stage—perhaps, if Quine is right, with Homo Javanensis—we became adept at processing information, that yet fell short of the power to abstract and generalize. What selective pressures can we then suppose to have effected in our brains, since then, the innovations required to bring us the capacity for fully abstract and general reasoning?

I take for granted that reasoning is something we do; that much or most of what we do is influenced by emotion; that psychology is interested in everything we do; and that psychology is a branch of biology. These breezy premises raise a number of questions.

1. What kind of connection might there be between biology and rationality?

2. More specifically, how does the normativity of logic relate to its biological origins? Is there not, in the very idea of such a connection, something akin to the naturalistic fallacy?

3. If our capacity for inference is in part a legacy of natural selection, are there specific emotional mechanisms that serve to influence reasoning at the proximate level?

4. If so, what is the connection between those means by which our inferences are policed by our emotions, and the formalisms that govern careful reasoning in science and mathematics?

I shall not address these questions directly. By the end of this paper, however, I hope to have gleaned enough to warrant saying a little about each. On the way there, here is the route I shall follow.

I will begin with prima-facie distinctions among strategic, epistemic and axiological domains of rationality. The domains soon merge at the edges, and the role of emotions in adjudicating border disputes turns out to be both crucial and equivocal, because of the emotions' ability to function as both cause and effect in all three domains. In addition, emotions are ubiquitous in both of what I shall call the mind's "two tracks": some emotions were shaped by natural selection long before the invention of language, but the elicitors and very identity of many emotions now commonly rest essentially on explicit linguistic function. In practice, for example, the "feeling of rightness" plays an indispensable role in our assessment of certain basic inferential forms as acceptable or unacceptable. But a striking characteristic of such inferential forms is that they apparently do not, at least in their pre-linguistic form, reflect any truly topic-neutral power of logic. That comes only with the contribution of explicit linguistic formulations of logical rules, and even then the best we can say about our capacities in that department is that they are uneven. This throws some light, I shall suggest, on the "rationality wars" that have pitted "pessimists" about human rationality against both "optimists" and "meliorists". From those considerations I will turn to a sampling of cases where inferences appear to affected not merely by pre-linguistic but actually by sub-personal processes, including, in some cases, what appear to be very directly chemical factors affecting belief, desire, or emotion as they affect our disposition to get from one belief to another.

## Kinds of emotional rationality

A plausible first approach to rationality distinguishes three forms or domains of its application: the *epistemic*, which aims at maximizing true beliefs and minimizing false ones (neither suffices alone, and no easy formula guarantees both); the *strategic* or practical, which aims at success in action in the light of goals and desires; and the *axiological*, which aims at appropriateness of evaluative emotional response. A running theme in the history of philosophy has consisted in attempts by each of the first two to subsume the other. Socrates famously asserted that the true explanation for bad choices invariably lay with a lack of knowledge. That claim rested on the dubious premise that everyone desires the same thing, namely the Good. From that, Socrates inferred that wrong choices reflect not ill-will, but misinformation. If we remove the dubious premise we can still make a case: if strategic irrationality consists in adopting a counterproductive course of action, it seems reasonable to presume that its counterproductive nature was not known. When I act, I must in some sense believe my action to be the best available. That belief (which doesn't necessarily exclude a simultaneous contrary belief) may be false or irrationally acquired. From that vantage point, epistemic rationality subsumes the strategic.

On the other hand, there is a considerable literature on the "ethics" of belief (Adler 2002). And while opinions differ as to whether believing can be done at will, the whole discussion presupposes that it is at least sometimes something that is *done*. As such it must be subject to considerations of strategic rationality (Levi 1967). Viewed thus, the strategic subsumes the epistemic.

Common cases where the two clash involve self-deception that serve the agent's short or long-term interests. Self-deception is a good thing, it is sometimes claimed, when it contributes to self-confidence, encourages effort, or keeps crippling realizations out of consciousness (Taylor 1989).[1] Emotions standardly play a determining role in self-deception, but they need not enter into the reasoning that motivates it. In some grander styles of self-deception, as advocated by Pascal with his famous wager, or in William James' plea for faith (James 1979) as against Clifford's insistence on the requirement of "adequate evidence" (Clifford 1886), emotions only come in at the implementation stage, once self-deception has been chosen by a dispassionate argument as a maximizing policy.

To see this, let us focus on Pascal. Ignoring theological subtleties,[2] Pascal's wager can be succinctly characterized in Bayesian terms:

> *The expected utility of believing in God is the weighted sum of the utility of living a life of relative deprivation, followed by eternal bliss, and that of living a life of deprivation followed by nothing. The weights are the probability that God (as conceived by the particular theology in question) actually exists and its converse. Even if the probability of God's existence is tiny (but still finite), the resulting expected utilities—respectively positive and negative—are infinitely large.*

It is important to note that while this argument deals with probabilities or Bayesian beliefs, it treats believing or disbelieving themselves as actions. Their expected desirabilities are determined in the same way as that of other actions in terms of their consequences weighted by their likelihoods. And while something like religious terror may well have motivated Pascal's invention of the wager, the argument itself is not overtly driven by emotion. On the contrary, it is

-----------------------------------

[1] I am grateful to Julie Kirsch for this reference.

[2] Ignore, in particular, the fact that the choice of a faith is not a binary one, so that whichever sect one joins, one risks the fate promised by all the others to those who deny the True Faith.

a pure calculation of costs and benefits, intended to counter what might be thought of as the greater prima facie emotional weight of the temptation of present earthly pleasures. It is only when Pascal goes on to give advice about how to implement the goal of believing against evidence that he counsels habit-forming behaviour, which might in turn bring the emotional tone of conviction. So he seems well aware that concentrating on the expected utilities figuring in the calculation would be ineffectual in changing belief.

But surely, the epistemologist will claim, considerations about heavenly bliss and hellish torment are not relevant to the issue of truth. Truth is by definition the formal object of belief, and therefore the standard of a belief's correctness; and only evidence is relevant to truth. (Adler 1999, 268). Therefore treating practical considerations as bearing on the rationality of belief is simple confusion.

Against this, however, Pascal can insist that the purity of truth can't have been an original virtue of belief. It had to be distilled, as it were, from the soup of overall benefit. Knowing the truth is useful in general, to be sure, but it is hasty to think that practical pursuits will massively fail if concern for truth is not paramount. Indeed, it seems to be strategically rational to lie a certain proportion of the time, both at the level of conscious policy and at the level of phylogenetic strategies. Mimetism, in which the markings of one species have in effect been selected to get a "free ride" on the toxicity of another similar-looking species, is effectively deception at the species level. (Sober 1994). Furthermore, as Chris Stephens (2001) as shown, while believing what is true is a good idea in general, there are cases—aptly captured in the slogan "Better safe than sorry," when a signal is wrongly interpreted as indicating the presence of a predator—where the best policy may systematically result in acting on falsehoods as often as on truths.[3]

-----------------------------------

[3] As the Editors of this volume have pointed out to me, it could be claimed that the belief embedded in the higher level of caution need not be the falsehood that 'this is a tiger', but the truth that 'This is sufficiently tiger-like not to be worth the risk.' But since the belief is bundled into the behaviour, it's not clear what evidence could justify one attribution rather than the other. In any case, however the behaviour is

In short, the value of truth—telling it, or believing it—should not be exaggerated, and it is only in a context in which one has already agreed to take the purity of epistemic norms for granted that Pascal's wager can be ruled inappropriate.

What then does distinguish those contexts in which epistemic norms are primary, from those which call for more pragmatic criteria?

This question, as I have argued elsewhere, calls for a meta-level judgment of appropriate value: sometimes a value is purely epistemic, at other times, broadly practical. Since that judgment arbitrates between the epistemic and the strategic, neither mode of rationality can make it without begging the question. When each side accuses the other of an ignoratio elenchi, the decision between them will inevitably depend on how much one *cares* about one and the other. Emotional endorsement is the only ultimate arbiter of the appropriateness of a standard of rationality (de Sousa 2003).

But if emotions are—in any sense, however perverse—to lord it over logic, what can that mean for the validity of logic? Though an arbitrator is not a dictator, and arbitration is called for only in very special circumstances, it would be paradoxical if the faculty we are accustomed to think most disruptive of Reason were crucially to contribute to the determinations of reasons. To resolve that paradox, we should note that at some point in the articulation of the most basic procedural rules, considerations of normative correctness must merge with certain factual necessities. Correct inference, at the most basic level, consists in the regular functioning of basic psychological mechanisms constrained by the power of social sanctions over individuals.

In a moment, I shall adduce a striking example of the depth of that power, not merely over individual emotions, but over the content of perception. (See below, "Six Paradigms", §(vi).) But it must be acknowledged at the outset that social pressure is sometimes—indeed,

interpreted, the "better safe than sorry" policy, typically sacrifices high expected utility for a drastic reduction of risk. In the perspective of evolution, that makes it, by definition, a losing strategy in the long run.

perhaps *often*—deplorably wrong-headed. If the power of social influence results from nothing but the random parasitism of memes, this line of thought is unpromising. The key to doing better lies in acknowledging something deeper than social pressure: an original evolutionary basis for some of the "intuitive" judgments that are codified in social consensus.

## Evolution and rational inference

The need to posit innate predispositions to certain forms of information could be made out with reference to concept acquisition, to classical induction, and to Nelson Goodman's "new" problem of induction (Goodman 1983). But let me illustrate with reference to a narrow and uncontroversial level of rationality: the rationality of logical inference. Clearly, making valid inferences is something we can study to do more effectively but, as famously demonstrated in Plato's *Meno*, it is also something that we know "instinctively." Meno's slave may not know geometry, and he is ever ready to hazard thoughtless answers, but when confronted directly with the consequences of his suggestions he can recognize them as right or wrong—something he could never do with merely empirical information. Similarly, whatever our level of logical sophistication, there are inferences of which we recognize the validity, others which we immediately see as invalid. That fact is no lucky accident. Unless we spontaneously recognized the validity of some basic pattern of inference, such as Modus Ponens or Modus Tollens,[4] no instruction manual could save our inferences from sinking into logical quicksand. That is the lesson Lewis Carroll's story of Achilles and the Tortoise: if we required every applicable rule of inference to be written down as a premise, we would need another rule to tell that the inference was an instance of it, and so on forever. The simplest inference would require us to endorse an infinite number of finite steps (Carroll 1895).

-----------------------------------

[4] These two patterns are not necessarily on a par as a matter of psychological fact. The point made here, like the point made in the *Meno*, holds a priori: it is that unless some patterns of transition are built into the architecture of the brain, no process of reasoning can get going. What those patterns are is an empirical question: in theory—and sometimes in practice—they might include *believe what you are told.*

So it's no mystery that at the most elementary level, the gap closes up between what we do naturally and what valid logical rules prescribe. What does remain unexplained by this, however, is *how* this happy harmony might have come to be pre-established.

This problem is made the more acute, as Mohan Matthen has pointed out, by the fact that most philosophers have given up on the hope of establishing a type-type identity between contentful mental states and neurological ones. So how are we to explain the lucky fact that the causal powers of belief states reflect the logical relations between the propositions believed? If assenting to $p$, assenting to $(p \rightarrow q)$, and assenting to $q$ are merely token identical to the neurological states that implement them, how does it happen that the conjunction of the first two states tend to cause the third? Perhaps those states are somehow tagged with the right syntactic markers in the brain. But if so, we still need to "show why syntax parallels semantics" (Matthen 1989, 563). He then argues that we have "no choice" about adopting some sort of evolutionary explanation. (564), appealing to the now standard notion of an evolutionary function, as arising from selection for some characteristic performance (Millikan 1993). Specifically, Matthen suggests,

> "we might have a special state that has *if p then q* as content, without its being a special purpose state that brings the *q* state into being whenever the *p* state occurs. Such a conditional state would depend for its effectiveness on the existence of higher level modules that were able to execute the logic of conditionals" (Matthen 1989, 567).

This suggestion raises three questions. The first is whether an evolved mechanism such as the one Matthen envisages could be pre-linguistic. The second is about the role of emotion in the processes in question. The third, which I shall refer to the *problem of scope*, is whether the mechanisms in question are sufficiently general to cover the full range of possible inferences.

To address the first question, we should first note that logic aspires to operate on topic-neutral form. It presupposes abstraction, requiring some form of representation that is devoid of specific reference. But it is hard to see how a pre-linguistic representation might have the required generality. Hence it may be more promising to reject the question, on the supposition that organisms that do not have language don't actually ever *need* topic-neutral mechanisms. It may be, as Peter Carruthers among others has suggested, that among the specific virtues of language, is the capacity to bridge and connect information derived from encapsulated pre-linguistic modules. On that view, language holds the monopoly on the capacity for abstract or topic neutral reasoning (Carruthers 2002).

This idea fits in with much evidence for a two-tiered organization of our cognitive faculties, a "two-track mind". Track One comprises relatively modular response patterns; The more abstract, language-bound patterns of Track Two sometimes supersede the others, but often both remain in competition. Prima facie, it would seem that simply in virtue of involving older structures in the brain, emotions might play an important role in the former, but not in the latter. In an early formulation of the hypothesis that brain structures originating at different stages of evolution have overlapping functions, Paul Maclean (1975) identified the limbic system as implicated in both emotional functions and cognitive ones; but only the cortex is involved in language capacity. More recent versions of the view are less cut-and-dried, but there is accumulating evidence for the basic hypothesis of the existence of two systems of mental processing, grounded in structures set up both by evolution and learning on the one hand, and involving explicit linguistic processing on the other. The two systems sometimes compete in generating beliefs, wants, and plans. The idea has appeared many times in psychology (Evans 2003), and its consequences have been most thoroughly explored by Keith Stanovich (2004), who lists twenty-three versions of it. His own version lists "automaticity, modularity, and heuristic processing" as the features characteristic of Track One, while Track Two is described as "rule-based, often language based, computationally expensive." (Stanovich 2004, 34-36).

All this suggests, in answer to my second question, that, insofar as emotions can be attributed to beings without language, they belong primarily in Track One. But the two-track mind hypothesis is not committed to drawing the line by opposing the phylogenetic origins of the faculties involved to those that arise by learning, including language learning. Many emotions require language for their full specification, and language reaches deep into the emotional brain. This is obvious from a glance at the power of political rhetoric. Equally obvious, however, is the fact that the emotional power of political rhetoric doesn't work by stimulating our capacity for logical inference. Instead, such uses of language proceed by triggering what George Lakoff calls "frames" and what I have called "paradigm scenarios" (Lakoff 2002); (de Sousa 1987). Frames and paradigm scenarios are fundamentally emotional in their mode of operation, and tend to wrest control from logical reasoning altogether.

## The rationality debate

It may be objected that the best-known examples of systematic irrationality involve word problems, as detailed by (Kahneman, Slovic and Tversky 1982), do not appear to involve emotions. How does this square with the idea that the irrationality is explained in part by emotional aspects of the two-track mind?

Let's look at how this question might apply to one of the best known cases, the Wason test. In the classic version of this, subjects are shown four cards, showing 3, D, 7, K. They are told that all cards have a letter on one side and a number on the other, and they are asked what cards must be turned over to verify that *if a card has D on one side it has 3 on other.* Most people get this wrong, in a puzzling variety of ways. Yet most people can easily solve a problem about how to make sure that under-age customers are not drinking alcohol. This latter problem has the same abstract form: in either case, the question can be represented abstractly as requiring verification of a statement of the form 'if $p$ then $q$', where the possible moves comprise turning over a $p$ card, a $\sim p$, a $q$, or a $\sim q$. Given that falsification of 'if $p$ then $q$' is secured iff $p \& \sim q$ is true, it should be clear that the actions required are turning over just the $p$ and the $\sim q$ cards. In the case of the drinking problem, that means inspecting the ID of those drinking beer and checking the drinks of those without ID. In the abstract card version, it means turning over the D and the 7. Why does the former seem so much easier than the latter? And what does it have to do with the use of language?

There have been a number of attempts to reject the head the problem off at the pass. One recent argument is made by Jerry Fodor, who denies the basic assumption that the two versions of the problem have the same logical form. Fodor notes that in cheater detection, the subject is charged with conditionally enforcing an imperative, whereas the card-turning version requires verification of a conditional statement (Fodor 2000, 101–104). To be sure, checking on *q* just in case *p* is not the same task as checking the truth of *if p then q*. The former task requires nothing unless p is true. But that accounts only for one typical mistake, which consists in checking only the *p* card.[5] But Fodor's ingenious attempt to pry the two tasks apart is a red herring. For among the vast number of variant experiments to which the Wason test has been submitted, some cases that could be represented as requirements or permissions fail to yield a majority of correct judgments, unless they specifically involve possible violations due to cheating (Cosmides and Tooby 1992, 199–205). Consensus on a clear diagnosis is still lacking; what is clear, however, is that a number of different content-dependent factors appear to be involved in determining whether problems of the same abstract form are more or less easily solved. The difficulty of accessing and applying the purely abstract schema is well established.

We can draw two morals, which at first blush appear somewhat inconsistent. First, that we are sometimes poor at reasoning about abstract problems, compared to our ability to deal with some of their specific concrete instantiations; second, that when we insist on applying formal rules of inference strictly, we can get it right and convincingly expose the error. So is explicit reasoning systematically irrational or not? What seems to be happening in the Wason

-----------------------------------

[5] Intriguingly, this mistake is more common among mathematicians, who are much less likely to make the mistake often made by other subjects, which is to turn over both the D and the 3 (Inglis, Simpson and Watson 2005).

test is that when the question is posed, it commonly fails to trigger the truth-table for the material conditional even when that schema has been studied in formal logic class. Lacking the clue that will route it to the cheater-detection schema, subjects fall back on some more accessible course, such as attending to the cards mentioned in the problem ("matching bias"), leading to the choice of turning over 3 and D, or confusing 'if' with the biconditional. If an emotion is involved, it might come under the heading of intellectual sloth. An abstract word problem requires an analytical Track Two strategy, and that is harder to access if some familiar and more easily available schema appears ready to hand.

Topic-neutrality is a defining characteristic of logic, in which arguments can be assessed independently of the reference of their terms . It is closely related to the idea of universality: for if validity can be assessed without regard to subject matter, then nothing is in principle beyond our comprehension. Yet one might reasonably doubt whether the capacities we now have, even when boosted by language, are able to span topics far removed from those likely to be treated by, or useful to, "the unrefined and sluggish mind of homo Javanensis." That is the problem of scope: are there inherent limits to our mental capacities that forever bar us from understanding some things? This idea has taken many forms. Here is a generic version: *our mental capacities have evolved under the selective pressure of a limited range of practical problems; so we have no good reasons to trust them when we venture beyond that restricted range into theoretical speculation.* The Whorf hypothesis—that a specific language may be equipped to express only a specific range of thoughts (Whorf 1973)—, and Colin McGinn's "mysterian" suggestion that there might be an intrinsic limit to the capacity of the mind to understand itself (McGinn 1982), are well-known variants of the scope problem. But its classic formulation comes from Descartes:

> For the proper purpose of the sensory perceptions given me by nature is simply to inform the mind of what is beneficial or harmful.... I misuse them by treating them as reliable touchstones for immediate judgments about the essential nature of the bodies located

outside us. (*Med. VI*)

The very precision of Descartes' version of the problem hints at a solution. The answer to the problem of scope lies not in the senses alone, but in our capacity to link sensory data to scientific models and thereby to mathematics. The possibility of discovering or constructing advanced mathematical structures cannot have evolved *as such*. Our remote ancestors are unlikely to have left significantly more offspring on the basis of abilities that could only have been manifested in the past two thousand years. The capacity to think about higher mathematics has to be what Gould called a spandrel or an exaptation (Gould and Lewontin 1979). But then what Eugene Wigner has called "the unreasonable effectiveness of mathematics in the natural sciences" (Wigner 1960) becomes a telling piece of evidence. The success of mathematics in solving problems that have nothing to do with immediate survival, coupled with its applicability in the construction of entirely new technological devices, affords credible support for the view that mathematical inference gives us some access to an objective world. Mathematics and science, however, do not (despite an apparent exception taken up in the next section) progress far without language. This leaves us free to suppose that while our pre-linguistic ancestors and cousins mastered an abundance of modular inferences based on analog computation and applicable in specific domains, they could not manipulate topic-neutral abstractions embodied in digital systems of representation.

That the first track deals in analog computing while the second track is the domain of digital representation is strikingly confirmed by work on non verbal counting and calculation in animals and infants, which suggests that animals are capable of primitive forms of ordering, adding, subtracting and even some forms of multiplication and division (Gelman and Cordes 2001); (Gallistel and Gelman 2000). Unsurprisingly, however, there is no evidence for a pre-linguistic domain-neutral representation system.

The topic-neutrality of language does not solve the problem of scope. It tells us nothing about the completeness of the potential knowledge to which we could give linguistic expression.

The mysterians might still be right about the ultimate limits of knowledge.[6] But it does lead us to expect a dichotomy between the range of problems that can be expressed and solved in the explicit language of science and mathematics, and those that are best approached in terms of the "intuitive" methods of Track One. The problem of the role of emotions in reasoning is particularly acute here, because it is obviously at the level most relevant to the activities and needs of our ancestors in the environment of evolutionary adaptation (EEA) that emotions are most likely to have preempted, or set up effective precursors to, explicit rational thought. As triggers to fast, pre-reflective and relatively organized responses to urgent situations, emotions are most likely to be involved in the inferential and decision-making short-cuts that were, at least in the EEA, highly cost effective if not always such as to guarantee correct answers. That would place them in the domain of Track One. By the same token, emotions are notoriously likely to block rational analysis. They can get in the way of the sort of calculation that might be required to devise more elaborate solutions to problems different in nature from those typically encountered in their contexts of adaptation. On the other hand, Antonio Damasio (1999) has described neurological evidence that brain lesions in the frontal lobes, by affecting patients' capacity for normal emotional responses, result in profound disruptions of practical rationality even while leaving unimpaired the ability to solve Track Two problems on the analytical and verbal level. Just why that is, Damasio leaves unclear; but we can speculate that the emotions are crucial to rational decision-making in several ways. They routinely control the agenda of practical reasoning by modifying the salience of different aspects of life situations; they narrow down to manageable size the unmanageable vastness of our options at any given moment of choice; and they may incline us to one side or another of alternatives that promise comparable expected utilities but different levels of risk (de Sousa 1987). They may also be essential to our ability to commit ourselves to stick with policies that affect our long-term interests (Ainslie

-----------------------------------

[6] My late Hegelian colleague Emil Fackenheim, puffing on his pipe in the Common Room in the days when North American philosophers were not yet not prosecuted for that egregious crime, once enunciated this Deep Thought: "The aim of philosophy is to find the limits of knowledge..." (puff) "... and then transcend them."

2001). Furthermore, specifically 'epistemic' emotions such as wonder, doubt, curiosity, surprise, or the "feeling of rightness" spur the quest for analytic rigour typical of Track Two processing.

It is in that light, I suggest, that we should view the "rationality wars" that have opposed, in recent decades, "pessimists" against "Panglossians" about human rationality. These terms are those of Keith Stanovich (2004). Among the pessimists are the contributors to (Kahneman, Slovic and Tversky 1982), who take themselves to have demonstrated a number of ways in which human reasoning is systematically flawed. The best known Panglossians are the contributors to (Gigerenzer, Todd and ABC Research Group 1999; Gigerenzer and Selten 2001) who argue that the examples exploited by Tversky, Kahneman and their colleagues are either artifacts of misleading test conditions, or instances of "fast and frugal" strategies, fine-tuned by natural selection, which are actually superior to those methods thought by the pessimists to embody rational thought.[7] Against both, Stanovich presents himself as a "meliorist", advocating an acknowledgement of the shortcomings of ancient and particularly of non-verbal strategies of decision and reasoning, followed by the careful cultivation of improved methods of explicit reasoning that take advantage of the sophisticated language developed for logic, mathematics and science. This approach is grounded in the acknowledgement of our two-track minds. But the meliorist attitude mustn't be interpreted as placing emotions exclusively in the first track, nor as requiring that they should play no role in reasoning. On the contrary, that role is deep and pervasive, and by no means limited to the pre-verbal strategies that might have been selected for us in the course of the EEA.

The difficulty most subjects have in solving abstract word problems such as the Wason test is slight compared to those observed long ago by Aleksander Luriia in conversations with illiterate Russian peasants. What impeded his subjects from performing simple Modus Ponens seemed to be an inability to focus on just the suppositions embodied in the premises of an

-----------------------------------

[7] A recent best-selling book takes a similar line, though radically downplaying the role of emotion (Gladwell 2005).

argument. They always brought more information to the problem than was contained in the problem they were asked to consider. This prevented them from marshalling the information expressed in the premise of the argument. When presented with *There are no camels in Germany*; *The city of B. is in Germany*, *Are there camels there or not?* some subjects reasoned that there probably were camels there if B was a large city, and if not, it was because it was too crowded. When asked for what "didn't belong", or was "not alike" in the list of words *saw*, *hammer*, *hatchet*, and *log*, one response was: "They all work together and chop the log". When Luriia persisted in trying to get them to attend to the fact that all except logs can be called 'tools', the informant replies:

> "Yes, you could, except a log isn't a tool. Still, the way we look at it, the log has to be here. Otherwise, what good are the others?[8]

Luriia noted that his informants were apparently incapable of repeating the problems put to them. And yet,, as Wilson comments:

> An untutored peasant will be perfectly clear in practice that if milk is healthy for humans and this is milk it is healthy for humans, and that if something looks like ordinary milk but is fatal to humans it is not ordinary milk. But she may not be able to repeat the corresponding syllogism or to draw the right inference on command.

It seems that the problems, when stated in words, trigger neither verbal schemata nor situational ones. Luriia's account doesn't allow us to judge whether emotions play a role in these difficulties; but it seems to me that it does suggest the power of what (Lakoff 2002) calls "frames" and I have called "paradigm scenarios": basic narratives reminiscent of situations and responses experienced in early age, that remain significant in someone's life, and that typically

-----------------------------------

[8] (Luriia 1976, 58–59). I'm indebted to Catherine Wilson for the reference and a discussion in an unpublished talk to the Canadian Philosophical Association (Wilson 2000).

evoke tendencies to specific sequences of action. The triggering of such a scenario seems to me sufficient, in many circumstances, to warrant speaking of emotion. But it is not entirely clear whether the triggering of such a frame or scenario should be construed as necessarily involving an emotion or not.

Six Paradigms.

To try to cast a little light on this question, I shall presently turn to five examples of reasoning or inference in which emotion is somehow involved, but where its exact role is hard to define. They are cases where we can observe, if not quite understand, the role that sub-personal factors, including chemical agents commonly implicated in emotional states, affect in a surprisingly direct way what we would otherwise regard as inferences driven solely by considerations of validity or evidence. But first, a reminder that non-verbal proto-reasoning is ubiquitous, but that in many forms it needn't involve emotion in any way.

*(i) Animal Computation.*

In the sort of skilled physical activity that we share with other animals, subpersonal computation plays an indispensable role. This does not explicitly involve emotions. A cartoon by Harry Hargreaves, published in *Punch* in 1960, summed this up rather well. A kingfisher looks down at a fish in the water, and the caption records its thoughts: "target in sector 3; speed, 2 knots; angle of deflection, 25$^{o}$ . . . . Who'd be a kingfisher?" Apart from the explicit form of the bird's thought, Hargreaves's cartoon is highly realistic about animals' powers of computation. Studies of ant navigation show that Desert ants track their distance from home by means of some sophisticated way of measuring the horizontal ground distance they have travelled even when their actual path involved a hilly journey of greater absolute length (Wohlgemuth, Ronacher and Wehner 2001). Pharaoh ants appear equally ingenious in their method of reorienting themselves to their nests. They rely on an ability to identify the angle of branching paths out from their nests, typically 60° when facing the nest, but 120° when going the other way (Jackson, Holcombe and Ratnieks 2004). But such capacities are by no means confined to "lower" species. Humans playing tennis—or more exactly their brains—solve complicated and efficient Bayesian equations, involving both acquired expectations and fast updating of feedback information (Körding and Wolpert 2004). This is done at the subpersonal level, and achieves a precision that couldn't possibly be attained by anyone attempting to solve the problem explicitly without the aid of sophisticated computers. These cases illustrate First Track processing at its most impressive, involving neither explicit reasoning nor emotion, yet producing highly accurate results. Such feats make it irresistible to postulate sophisticated computing mechanisms working at the sub-personal level, deep beneath the level of any consciousness or explicit formulation.

Unlike these first examples, which are intended merely as reminders of the vast terrain of sub-personal "inference" involved in the ordinary physical motions of humans and other animals, the next series of examples all involve emotion.

*(ii) The Capgras syndrome.*

The first involves the emotion of *feeling of familiarity,* or rather a stubborn feeling of *unfamiliarity* that resists ordinary recognition. Typical patients afflicted with the Capgras syndrome persist in believing that a person close to them — wife, or father — is an impostor.[9]

According to Ramachandran, the best explanation for this strange disorder is that a direct link normally exists between the facial recognition mechanism and the areas controlling the appropriate emotional responses (particularly the amygdala). The sight of a close relative—a parent, in the case of Ramachandran's patient Arthur—normally triggers an affective response, which is itself subject to a "familiarity" evaluation. In Arthur's case, the direct link to the area in charge of generating the affective response is missing. As a result, the affective response to his father is not produced. This sets up a incongruity between the strictly cognitive familiarity check

-----------------------------------

[9] See (Ramachandran and Blakeslee 1998);. (Pacherie 2005; Mangan 2001) The next two paragraphs draw from my comments on Mangan's paper in the same online issue of *Psyche*.

that applies to the face and the missing familiarity check applied to the expected affective response. The Capgras delusion can then be construed as the conclusion of a perfectly reasonable inference (though of course one that is neither conscious nor explicit): *I get a characteristic thrill when my father appears; I'm not getting that now; therefore the person before me is not my father.* On the other hand, *he looks exactly like my father. Therefore he is an impostor, a stranger who looks just like my father.* This hypothesis is particularly neat in its capacity to explain why it is that the "impostor" delusion occurs only with persons to whom the person is close: typically parents or spouses. It doesn't occur with mere acquaintances, because in most cases of recognition a more or less indifferent emotional reaction is normal, not aberrant. (It also doesn't normally occur over the telephone, which doesn't implicate the same pathways of facial recognition.) If something like this is correct, it would imply that the emotional aspect of recognition is subject to an independent familiarity marker. Where the person recognized is both familiar and affectively significant, both markers are involved in the required ID check.

Two things are worth noting about this case. First, although we can make sense of it by construing it as a kind of inference, it is not experienced as an inference, but as intuitive conviction. Second, while the "feeling of rightness" acts as a marker, it doesn't present itself as a marker of correct inference *as such*. There are other cases, such as obsessive-compulsive disorder (OCD), however, where the feeling of rightness does just that.

(*iii) Obsessive-Compulsive Disorder.*

OCD affects specifically not the sense of familiarity but the feeling of "rightness" itself. OCD may be seen as resulting from some sort of disconnection of the normal emotion of rightness in relation to recent memory of having taken necessary precautions. The relevant emotions here would be specifically *epistemic* emotions. As Chris Hookway has pointed out, epistemic emotions have been almost wholly neglected in the literature but constitute an extremely important aspect of our ability rationally to reason our way to new beliefs. If I didn't experience *doubt*, I wouldn't launch on an inquiry in the first place. If I didn't have the feeling of *rightness* about an inference, I wouldn't rely on it. If I didn't have the feeling of *conviction* about a conclusion, I wouldn't infer it (Hookway 1998). The patient suffering from OCD lacks some of those normal feelings. OCD has traditionally been taken to be a neurotic syndrome calling for psychoanalytic diagnosis and therapy. But the fact that some of these cases are apparently capable of clearing up under the influence of a targeted drug such as Prozac (Kramer 1993) suggests that this apparent complexity is an illusion. As in the case to which I turn in the next paragraph, it seems that the feeling of rightness is an emotion can be triggered or at least facilitated by a simple chemical agent, and in turn determine the presence or absence of conviction in a particular proposition.

*(iv) The Chemistry of Trust*.

In a recent article in *Nature* widely reported in the press, researchers at the University of Zurich have shown that "intranasal administration of oxytocin, a neuropeptide that plays a key role in social attachment and affiliation in non-human mammals, causes a substantial increase in trust among humans." (Kosfeld, Heinrichs, Zak, et al. 2005). Their results also support the conclusion that "the effect of oxytocin on trust is not due to a general increase in the readiness to bear risks. On the contrary, oxytocin specifically affects an individual's willingness to accept social risks arising through interpersonal interactions." (ibid). The experimental set-up in each run of the experiment involved two subjects, an "investor" and a "trustee". Both received 12 monetary units, and each unit invested by the investor was tripled by the experimenter. Thus if the investor handed over all of his 12 units, the trustee how had 48, comprising his original 12 plus the tripled investment. He could then return any amount to the investor.

Trust is a nice bridge emotion between the strictly epistemic and the strategic; for it inclines us to believe identifiable propositions—in this case, the proposition that investment in this particular trustee would prove profitable for the investor— but only in the context of a transaction envisaged with a person. A *specific* acceptance of *social* risk was ingeniously distinguished from a *generic* increase in risk tolerance. This was done by comparing the original set-up with a situation in which the investor was told that a computer, not a person, would determine what return if any the investor would get. In that situation, the oxytocin had no significant effect. This seems to show that oxytocin didn't simply shift the estimate of risk down a notch; rather, it worked specifically on the emotional component of trust. Another significant control contrasted the effects of oxytocin on the investor with its effect on the trustee: the latter was nil. This showed that the causal factor responsible for the effect wasn't a general increase in benevolence or good feeling. For that would presumably have led also to larger returns from the trustees.

The authors note that there is substantial independent evidence "that oxytocin promotes prosocial approach behaviour by inhibiting defensive behaviours" (675). In the light of this known effect of oxytocin on social approach in other mammals, they tend to minimize its specific effect on belief: "the investors given oxytocin show more trusting behaviour but do not hold significantly different beliefs about the trustworthiness of others." That is paradoxical, if we assume that in either case the behaviour of the investor follows a roughly Bayesian strategy. It can be partly though not wholly explained, according to the authors, by appealing to an evaluative rather than a strictly cognitive appraisal: what the chemical has done is help the investors "overcome their betrayal aversion in social interactions." Still, the *consequence* of the diminished "betrayal aversion" is equivalent to a change in the probability measure of the expectation of return. So we have here a kind of primitive, purely causal case of direct biological influence over a process that is, we might say, functionally equivalent to an inference, even though no explicit inference is made.

If we represent the process in question as a Bayesian calculation,

$$V = \Sigma_{i=1-n}(p_i \times v_i)$$

the fact that the estimate of risk [p] is not directly affected suggests that “betrayal anxiety” feeds into the desirability factor [v] rather than the probability estimate [p]. Yet there is no explicit inference. We have only the behavioural upshot to go on. So we can think of the emotion as feeding into the practical inference without necessarily assuming that such an 'inference' can, like an explicit Bayesian argument, be split into "belief" and " desire" components. Insofar as the Bayesian model fits, however, emotion seems to be targeting the desirability rather than the belief component.

In any case, as might be expected, the effect of the oxytocin is not determining. It may contribute to a Svengali effect, but cannot guarantee its success and could hardly be credited with one all by itself. (Specifically, the median amount entrusted by investors who had absorbed oxytocin was 25% higher than those sprayed with a placebo.) So what needs to be made more precise is the nature of the relation between that sort of direct chemical influence on inference, on the one hand, and the influence that common sense attributes to other emotions in cases of bona fide valid inference, on the other.

*(v) Cognitive foraging.*

One more illustration of the surprisingly direct influence of emotional chemistry on inference is worth noting. In a recent issue of *Nature*, Jonathan Cohen and Gary Aston-Jones (2005) look at some findings by Angela Yu and Peter Dayan (2005) on the application to science of the trade-off between exploration and exploitation. Exploitation of known resources is safe but likely to yield diminishing returns. On the other hand, giving up well-trodden paths for the sake of exploration may yield a jackpot of discovery, but is inherently risky.

That trade-off is well known to students of foraging. An ant faced with an established path may either follow it, in the expectation of finding food where many others have already found it, or else strike out in an unexplored direction. The latter option is risky but will pay off, if not for the individual at least for the colony, when the original sources of food are exhausted (Johnson 2003). This is a good example, then, of a mechanism first instantiated at the most basic level of foraging decisions. What is surprising is that it can be directly applied in the context of sophisticated scientific cognitive strategies, where it appears still to be controlled by a combination of chemical triggers.

Subjectively, the tension between the relative security of "normal science" and the excitement of a potentially fruitful paradigm shift is experienced as an urbane sort of struggle between fear and greed. What Yu and Dayan found is that the balance between the tendency to explore and the tendency to exploit in the cognitive domain are apparently regulated in part by specific neuromodulators, controlling respectively the kind of uncertainty that arises from the bearing of a signal and the kind of uncertainty that arises from the reliability of the signal:

> Acetylcholine signals expected uncertainty, coming from known unreliability of predictive cues within a context. Norepinephrine signals unexpected uncertainty, as when unsignaled switches produce strongly unexpected observations. These uncertainty signals interact to enable optimal inference and learning in noisy and changeable environments. This formulation is consistent with a wealth of physiological, pharmacological, and behavioral data implicating acetylcholine and norepinephrine in specific aspects of a range of cognitive processes.
>
> (Yu and Dayan 2005, 681).

They go on to remark that there seem to be "a class of attentional cueing tasks that involve both neuromodulators and shows how their interactions may be part-antagonistic, part-synergistic." (ibid.) And of course those sorts of situations are typically experienced, in humans, as giving rise to emotional states: the "fear" of risk; the "lure" of the unknown; the "disappointment" generated by scientific prospects that don't pan out. What Yu and Dayan's discovery seems to be telling us, is that a chemical mechanism underlies, in part, both the phenomenology of emotion and the process of what we assume to be high-level decision-making. What they don't tell us, which raises an intriguing question, is just what relationship there is between those two aspects of brain chemistry and their felt and functional consequences in the subjective experience of the quest for invention and discovery.

*(vi) The emotional price of nonconformism.*

In some famous experiments on the power of social conformism done in the 1950's, Solomon Asch had found that when asked to make a judgment of a visual quantity, some 40% of subjects went along with the false judgment of stooges posing as fellow subjects (Asch 1962). In a new variant on these experiments, a group of researchers explored "the potency of social pressure in inducing conformity, how information that originates from humans, versus inanimate sources, alters either perception or decision making and the neural basis for such changes" (Berns, Chappelow, Zink, et al. 2005, 6). Here again, there are intriguing questions about the role of emotion, and about the precise locus of the inference to the wrong conclusion. Using fMRI data, Berns et al. found a "highly suggestive... lack of concomitant activity changes in more frontal areas," where one might have expected activity if the subject's judgment had resulted from a decision to override one's own judgment in favour of that of the majority. The surprising aspect of their findings is that no special cognitive activity was detected in the cortex of those who conformed to others' false opinion. Instead, "the effects of social conformity are exerted on the very same brain regions that perform the task". In other words, far from being an inference required by the need to resolve cognitive dissonance ("So many others can't be wrong, I must revise my verdict"), the influence of others' judgments seems to act *directly on perception*. The distorting effect of conformity did not require any calculation of costs and benefits: it was those who saw and stood up for an independent truth who endured emotional cost. This finding is as intriguing as it is discouraging, for the emotions involved (though the authors make no attempt to pin-point these in our repertoire of normal emotions) do not seem to be among those that we would spontaneously label 'epistemic emotions'.

## Skeptical Concluding Remarks

The mere activation of neuromodulating chemicals, as instantiated in the findings I have sketched, can't be assimilated to the presence of an emotion. One reason for this is that emotions are phenomena that belong to the *personal* level; the activation of neuromodulators is a *subpersonal* phenomenon. Some of what we are learning about the involvement of specific parts or functions of the brain in reasoning, illustrated by cases (ii)-(vi) above, implicates just such subpersonal factors, but it is not clear to what extent we are justified in inferring, from the fact that some of the same chemical factors are involved in emotion, that emotions, as commonly conceived, are involved in reasoning. Nevertheless, I shall risk some tentative answers to my leading questions.

*1. What kind of connection might there be between biology and rationality?*

Social pressure alone won't suffice to guarantee the normative correctness of our inclinations to draw inference. A sort of core of basic procedures—perhaps including Modus Ponens—had to be installed by natural selection, in the sense that we have a native disposition to effect transitions from one belief to another in accordance with such principles in specific domains. Thus far only Track One processes need be involved. But when these transitions are codified by rules of language, logic or mathematics, we can see them in their full generality as well as provide an explicit and conclusive argument for their validity. In practice, however, commonly used inference patterns do not necessarily become more reliable. The reason is that the fit between the "native" dispositions to acquire beliefs and their implementation in explicit language is not itself part of that system of mechanisms on which natural selection has put its certificate of warranty.

2. *Is there not, in the very suggestion that there might be a connection between logic and biology, something akin to the naturalistic fallacy?*

There is a prima facie presumption of functionality to any heritable disposition the complexity of which makes it unlikely to be accidental. But it is crucial to remember that what's been put in place by natural selection, however useful to our ancestors in the EEA, may be not be worthy now of any evaluative endorsement. Justification can't be infallibly grounded. It has to run in circles, the capacity of which to inspire respect depends directly on their size. Ultimately we trust our epistemic emotions to tell us when the circle is big enough for comfort. In that way, Track Two's analytic mechanisms must submit to the judgment of emotions.

3. *If our capacity for inference is indeed in part a fruit of natural selection, are there specific emotional mechanisms that serve to guide our inference at the proximate level?*

The surprising lesson to be learned from the samples I have cited from recent psychological and brain research is that in some cases a relatively abstract inference is triggered by what appears to be a fairly simple chemical agent. It would obviously be greatly exaggerated to conclude from the Zurich experiments that trust was simply triggered by oxytocin, or from those of Yu and Dayan that strategic research decisions were determined by noradrenaline. But these experiments are part of a accumulating body of evidence that suggests that emotional factors, more obviously linked to non-cortical brain and hormonal activities, are important to our judgments of what inferences are or are not acceptable.

4. *What is the connection between those means by which our inferences are policed by emotions, and the formalisms that guide our most careful explicit inferences in science and mathematics?*

To answer this question, I have suggested that we must take seriously the Two Track hypothesis: we are endowed with two only partly connected and sometimes conflicting levels of processing. Emotions are involved in both; and while it is not surprising to find their role in Track One processes closely tied to the effects of various sub-personal factors, including neuromodulators, what is more surprising is that such factors also appear to be implicated in Track Two reasoning.

In the final analysis, the increasing precision of our understanding of the brain mechanisms underlying the actual and the normatively correct practices studied by psychology and epistemology may blur the image I have sought to sketch. The idea that one should be able to distinguish specific contributions of emotion to inference, and indeed that there is always a clear answer to the question of what inferences have been drawn, presupposes that we can set out clear lines of demarcation between the mind's two tracks, between emotional and merely evaluative determinants of decision-making, and between the influence that brain chemicals exert on reasoning and the effects they have on emotions. But the fine-grained picture, when it emerges, may overwrite the lines drawn in the sand by the presuppositions of our questions.

## References


Adler, J. E. 1999. The ethics of belief: Off the wrong track. *Midwestern Studies in Philosophy* 23:267-85.

------. 2002. *Belief's own ethics*. Cambridge, MA: MIT Press.

Ainslie, G. 2001. *Breakdown of will*. Cambridge: Cambridge University Press.

Asch, S. 1962. *Social Psychology*. Englewood Cliffs, N.J.: Prentice-Hall.

Berns, G. S., et al. 2005. Neurobiological correlates of social conformity and independence during mental rotation. *Biological Psychiatry*, online 22 June 2005.

Carroll, L. 1895. What the Tortoise said to Achilles. *Mind* 4:278-80. Available online at http://www.ditext.com/carroll/tortoise.html.

Carruthers, P. 2002. The cognitive functions of language. *Behavioral and Brain Sciences* 25(6):657-74; 705-10.

Clifford, W. K. 1886. The ethics of belief. In *Lectures and essays (2nd. Ed)*, ed. L. Stephen and F. Pollock. London: Macmillan.

Cohen, J. D., and G. Aston-Jones. 2005. Decision amid uncertainty. *Nature* 436 (28 July):471-72.

Cosmides, L., and J. Tooby. 1992. Cognitive adaptations for social exchange. In *The adapted mind: Evolutionary psychology and the generation of culture*. The adapted mind, ed. J. H. Barkow, L. Cosmides, and J. Tooby, 163-228. Oxford: Oxford University Press.

de Sousa, R. 1987. *The rationality of emotion*. MIT Press, a Bradford Book.

------. 2003. Paradoxical emotions. In *Weakness of will and practical irrationality*, ed. S. Stroud and C. Tappolet, 274-97. Oxford; New York: Oxford University Press.

Evans, J. S. B. T. 2003. In two minds: Dual process accounts of reasoning. *Trends in Cognitive Science* 7(10):454-59.

Fodor, J. 2000. *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge, MA: MIT Press, A Bradford book.

Gallistel, C., and R. Gelman. 2000. Non-verbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences* 4(2):59-65.

Gigerenzer, G., and R. Selten, eds. 2001. *Bounded Rationality*, 320. Cambridge, MA: MIT Press.

Gigerenzer, G., P. Todd, and ABC Research Group. 1999. *Simple heuristics that make us smart*. New York: Oxford University Press.

Gladwell, M. 2005. *Blink: The power of thinking without thinking*. New York: Little, Brown.

Goodman, N. 1983. *Fact, fiction, and forecast*. 4th ed. Cambridge, MA: Harvard.

Gould, S. J., and R. L. Lewontin. 1979. The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London* B 205:581-98.

Hookway, C. 1998. Doubt: Affective states and the regulation of inquiry. *Canadian Journal of Philosophy, Supplementary Volume* 24:203-26.*Pragmatism*, vol. 24, ed. C. Misak.

Inglis, M., A. Simpson, and D. Watson. 2005. Mathematicians, Eye-Movements and the Selection Task. *Under Review*.

Jackson, D. E., M. Holcombe, and F. L. Ratnieks. 2004. Trail geometry gives polarity to ant foraging networks. *Nature* 432:907-09.

James, W. 1979. The will to believe. Ed. F. H. Burkhardt. In *The will to believe: And other essay in popular philosophy.* Cambridge, MA: Harvard University Press.

Johnson, N. L. 2003. The development of collective structure and its response to environmental change. *Semiotics, Evolution, Energy, and Development* 2(3).

Kahneman, D., P. Slovic, and A. Tversky, eds. 1982. *Judgment under uncertainty: Heuristics and biases*. Cambridge and New York: Cambridge University Press.

Kosfeld, M., et al. 2005. Oxytocin increases trust in humans. *Nature* 435(2), 2 June:673-76.

Körding, K., and D. M. Wolpert. 2004. Bayesian integration in sensorimotor learning. *Nature* 427(6971), 1915 January:244-47.

Kramer, P. D. 1993. *Listening to Prozac: A psychiatrist explores antidepressant drugs and the remaking of the self.* London, New York: Penguin Books.

Lakoff, G. 2002. *Moral politics: How liberals and conservatives think*. Chicago: University of Chicago Press.

Levi, I. 1967. *Gambling with Truth*. New York: Alfred A. Knopf.

Luriia, A. R. 1976. *Cognitive Development: It's Cultural and Social Foundations*. Cambridge, Massachusetts: Harvard University Press.

MacLean, P. D. 1975. Sensory and perceptive factors in emotional functions of the triune brain. In *Emotions: Their parameters and measurement*, ed. L. Levi. New York: Raven Press.

Mangan, B. 2001. Sensation's ghost: The non-sensory "fringe" of consciousness. *Psyche* 7. Http://psyche.cs.monash.edu.au/v7/psyche-7-18-mangan.html.

Matthen, M. 1989. Intentional parallelism and the two-level structure of evolutionary theory. In *Issues in evolutionary epistemology*, ed. C. A. Hooker and K. Halweg.

McGinn, C. 1982. *The character of mind*. Oxford: Oxford University Press.

Millikan, R. G. 1993. In defense of proper functions, White Queen psychology and other essays for Alice, 13-29.

Pacherie, E. 2005. Perception, emotions and delusions: The case of the Capgras delusion. Paris: Unpublished.

Sober, E. 1994. The primacy of truth-telling and the evolution of lying. In *From a biological point of view: Essays in evolutionary biology*. Cambridge Studies in Philosophy and Biology, 71-92. Cambridge: Cambridge University Press.

Stanovich, K. 2004. *The robot's rebellion: Finding meaning in the age of Darwin*. Chicago: Chicago University Press.

Stephens, C. 2001. When is it selectively advantageous to have true beliefs? sandwiching the better safe than sorry argument. *Philosophical Studies* 105:161-89.

Taylor, S. 1989. *Positive illusions: Creative self-deception and the healthy mind*. New York: Basic Books.

Whorf, B. L. 1973. *Language, thought, and reality*. Ed. J. B. Carroll. Foreword by S. Chase. Cambridge, Massachusetts: M.I.T. Press.

Wigner, E. P. 1960. The unreasonable effectiveness of mathematics in the natural sciences. *Communications in Pure and Applied Mathematics* 13:1-14.

Wilson, C. 2000. Human rationality" as a misplaced composite: A response to the controversy between optimists and pessimists. Canadian Philosophical Association. Edmonton.

Wohlgemuth, S., B. Ronacher, and R. Wehner. 2001. Ant odometry in the third dimension. *Nature* 411:795-98.

Yu, A. P., and P. Dayan. 2005. Uncertainty, neuromodulation, and attention. *Neuron* 46:681-92.