

Do Unconscious Beliefs Yield Knowledge?

Luís M. Augusto

(University of Sussex, Dept of Psychology / FCT)

1. Introduction

Epistemology has now been for long widening the gap that separates it from the experimental sciences. I specifically mean sciences such as cognitive neuropsychology and cognitive science, which share their object, cognition, or knowledge, with epistemology. The method by means of which it has been doing so is a negative one par excellence, and one guaranteed to give one the results sought for. I am referring to ignorance, deliberate or accidental, in that epistemology has consistently ignored for many decades important results concerning perception and cognition obtained in these experimental fields. This cannot be due solely to a dislike of experimentation, because epistemologists often resort to it, if only of the purely fictional kind (mental experiments); it thus remains that it must be because their approach targets processes of perception and cognition that are different in some way from those empirically approachable, which does not seem to be the case.

While psychology has, in the last decades, been forced to reconsider many of its basic presuppositions in light of findings issuing from experimental research on cognition, epistemology has stuck to a petrified analysis that wholly disregards the actual conditions of formation and acquisition of knowledge. Of major import has been the work carried out both in perception dysfunctions and by experimentation on learning showing that knowledge is not entirely equatable with consciousness, i.e. one may have knowledge without for that being aware of that possession. We are now at a point when psychology frequently and unproblematically speaks of ‘unconscious knowledge,’ thus frontally opposing epistemology, in whose vocabulary this is a contradiction in terms: the tripartite analysis of knowledge requires (implicitly, if not explicitly) that subjects be conscious both of holding beliefs and of the contents of their beliefs. For those arguing that the fact that this requirement is not explicit does not necessarily make it implicit, the necessary alterations in the tripartite analysis of knowledge

will show how this is not the case. Let us take a proposition p and a belief-holding subject S ; remember that we specifically wish that S be not aware of her/his beliefs and of their contents. The new additions in bold, the analysis of knowledge will resemble something like this:

- (i) p is true.
- (ii) S **is not aware that** s/he believes that p .
- (iii) S is justified in **unconsciously** believing that p .
 S knows that p **while not knowing that s/he knows, or even believes in, p .**

Condition (ii) may be made acceptable by a broader—or vaguer, depending on one’s view—conception that sees belief as willingness to act in a certain way, and condition (iii) may even allow of the peaceful co-presence of the adverbs consciously and unconsciously, but the conclusion will certainly lose much of its apparent solidity, and, for those of a more formal persuasion,¹ even its validity and/or consistency: the ‘ignorance set’ $\{Kp, \neg KKp\}$ (where Kp is read “ p is known”)² is consistent only in the epistemic modal logic **T**, whose most distinctive feature is that it is the modal logic in which the subject has the least access to her/his epistemic situation (in other words, it is the ‘weakest’ system in epistemic terms).³

In this paper, I shall be arguing that unconscious beliefs yield knowledge, and that therefore to speak of unconscious knowledge is perfectly legitimate. My stance is that regarding a positive epistemic status both conscious and unconscious beliefs are on the same footing. This obviously makes the expression “unconscious knowledge” simultaneously superflu-

¹ I am assuming that epistemic logic is of interest for epistemological matters.

² Note that this reading does not stress the passive form: the readings “ p is known” and “a subject/agent knows p ” (also represented as $K_p p$) are synonymous in this context.

³ In more technical terms, **T** has only a reflexive accessibility relation R , according to which $\forall w \in W : wRw$, where w is a possible world in a Kripke structure $M = (W, R, \nu)$ where (W, R) is a frame and ν is a valuation; less technically, a reflex accessibility relation means simply that every world has access to itself. In order to give a brief view of the importance of the accessibility relations in epistemic terms, one should realize that to be able to express full knowledge ($Kp \rightarrow KKp$; *axiom of positive introspection*, or *axiom of self-awareness*) and even knowledge of one’s ignorance ($\neg Kp \rightarrow K\neg Kp$; *axiom of negative introspection*, or *axiom of wisdom*) one needs an accessibility relation that is, respectively, transitive [$\forall w, w', w'' \in W : wRw' \ \& \ w'Rw'' \rightarrow wRw''$] as well as reflexive, and equivalent, i.e. reflexive, transitive, and symmetric [$\forall w, w' \in W : wRw' \rightarrow w'Rw$]; these latter modal logics are **S4** and **S5**. Note that a doxastic modal logic **K45**, the doxastic correspondent to the epistemic modal logic **S5**, will reject even condition (ii) above, because it contradicts what for it is a normal reasoner: $Bp \rightarrow BBp$!

ous and erroneous: my giving positive epistemic status to unconscious beliefs simply means that I hold the thesis that these beliefs yield knowledge simpliciter. However, the expression “unconscious knowledge” is relevant in contexts in which we wish to specify that we are talking about knowledge yielded by unconscious beliefs, and thus I do not propose its elimination.

Back to the canonical analysis of knowledge, I am not arguing for the addition in it of the expressions in bold: I dismiss it altogether as irredeemably flawed. My objective in presenting the corrupted form above was to show that the canonical analysis of knowledge cannot and does not take into consideration either unconscious beliefs or unconscious knowledge; my dismissal is undoubtedly motivated to a great extent by this neglect, or incapacity, but my main reason for this dismissal is that this analysis of knowledge simply is wrong; that it cannot handle unconscious beliefs is just one more proof of its incapability to provide a correct account of knowledge. The argumentation that follows is to be read taking into account my analysis of knowledge, founded on a notion of truth as non-contradiction. Briefly, and merely anticipating the final chapter, I claim that the following are the necessary and sufficient conditions of knowledge:

- (i) S believes <that> p .
- (ii) S is justified in believing <that> p .
- (iii) p is true.⁴

According to this analysis, one knows <that> p when the belief <that> p is justified (= true); it is justified when *it works* in that it does not contradict the external world. This is its ‘proof’; in this consists the ‘provability’ of this analysis.⁵

In the light of this, one may argue that this paper is intended to show that this analysis of knowledge is *the* correct one. That reading, however, would be wrong; as I see it, there might be other analyses of knowledge capable of replacing the canonical one equally as well, or even better. My objective with the present paper is altogether another: I aim to show that

⁴ Condition (iii) is actually irrelevant, as shown below.

⁵ This analysis of knowledge that I propose here will be but insufficiently explained in this paper, reason why I refer the reader to (Augusto, 2008), where s/he will find an adequate elaboration.

one cannot ignore the fact, unearthed by the experimental sciences but defended long ago by less orthodox figures such as Nietzsche and Freud, that knowledge is not the privilege of consciousness. If one accepts this fact, and yet is unable to fit it into her/his elected analysis of knowledge, then one is indeed in trouble. That is not my case, given that, as shall become evident below, my analysis of knowledge unproblematically takes into account unconscious beliefs.

This will serve for an introduction. My next steps will follow this order: I first give a detailed account of some findings in the experimental sciences that force us to rethink the equation between consciousness and knowledge. I divide this section into two parts, section 2.1. dedicated to data concerning perceptive and cognitive dysfunctions suggesting unconscious (also: implicit; tacit) beliefs and unconscious (also: implicit; tacit) knowledge, and section 2.2. describing experimental work aiming at showing that unconscious cognition is a reality. They are *our* case studies, and this whole section is entitled accordingly. I then initiate my argument proper with section 3.1., a theory of belief and belief ascription, and I fundament it in section 3.2. with a theory of positive epistemic status.

2. Case Studies from the Experimental Sciences

2.1. Perception and Cognition Dysfunctions: Ventral vs. Dorsal Stream

Abundant neuropsychological data suggesting unconscious perception and cognition is today available from studies with humans on many dysfunctions. I chose to use three that are now being studied from the perspective of a double visual pathway, the *dorsal vs. ventral stream*, postulating a parallel conscious and unconscious visual processing in humans. Summarily, evolution would have provided us with two parallel visual streams, the most ancient one concerned with action (the dorsal stream, or the “where”/“how” pathway), and the most recent one aiming at object recognition and analysis of form (the ventral stream, or the “what” pathway; see Goodale & Milner (1992) for an elaboration). As far as the subject of this paper is directly concerned, the dorsal stream is of extreme interest in that it is hypothesized to process its percepts unconsciously, contrarily to the processing carried out by the ventral stream. The percepts of the latter do not necessarily reach consciousness, but they are in prin-

ciple accessible to it, whereas the former are believed to be processed in a wholly unconscious way:

The visual information used by the dorsal stream for programming and on-line control, according to the model, is not *perceptual* in nature. According to our definitions, therefore, it cannot be assessed consciously, even in principle. [...] We maintain that the nature of both dorsal-stream vision and blindsight stand in sharp contrast with visual processing in the ventral stream, *even when that processing fails to reach awareness*. The processing of vision for perception—conscious or unconscious—is, according to our model, restricted to the ventral system. (Milner & Goodale, 2007, p. 776)

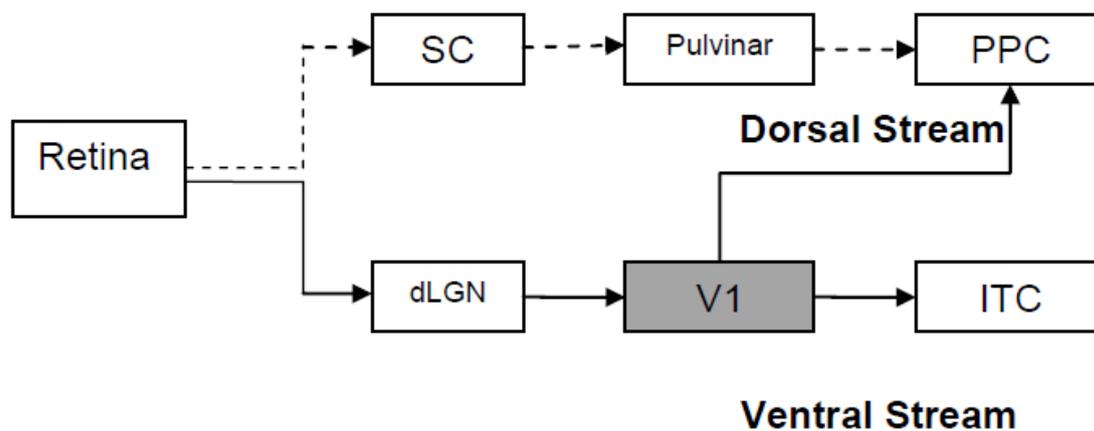


Fig. 1 – The Ventral and Dorsal Streams (SC = superior colliculus; PPC = posterior parietal cortex; dLGN = dorsal lateral geniculate nucleus; ITC = inferotemporal cortex; V1 = primary visual cortex)

It is important to note that the unconscious processing carried out by the dorsal stream is not necessarily a mere automatism, involving, on the contrary, complex learned cognitive processes implying intentions, (see Vakalopoulos, 2005, p. 1185) and, as such, beliefs.

Within this approach, the phenomena exhibited by patients with *blindsight*, *left visuo-spatial neglect*, and *prosopagnosia* are explainable as the sparing of the dorsal pathway in the impairment/removal of the ventral one. In other words, patients with these conditions would act under unconscious beliefs grounded on visual input processed exclusively in the dorsal stream.

Blindsight — More than to an impairment, the term ‘blindsight’ is used to refer to the capability of patients with totally blind regions of their visual field (scotomata) caused by damage to the primary visual cortex (V1) to somehow perceive visual stimuli presented in their blind fields. Although claiming not to perceive stimuli presented in their scotomata, when forced to choose or guess, patients will accurately, or well above chance, not only localize but also identify moving or static stimuli, such as shapes. It is important to emphasize the fact that this forced-choice or forced-guessing strategy is actually essential, because patients report absence of stimuli and may even become annoyed at the insistence shown by the experimenters that they report what for them simply is not there. This total lack of awareness of stimuli has been distinguished from the cases in which ‘awareness’ of a percept *not* of a visual nature is reported; these are now known as Types 1 and 2 of blindsight, respectively.

In 1973, D.B., probably the most well-studied patient with blindsight, underwent an operation in which the striate cortex of the right hemisphere and a part of the adjacent calcarine cortex were removed; this surgical intervention aiming to rid him of severe migraines left intact other areas of the cortex that are usually affected in other lesions; it had as an expected side-effect an hemianopia on the whole of his left field. Years of research on D.B. have shown that, though reporting no consciousness of any, he can accurately reach or point towards stimuli; he makes eye movements when a light is flashed in his visual left field; he correctly distinguishes **X** vs. **O**; he is also capable of discriminating line orientations.⁶

Recently, the range of stimuli perceivable by patients with blindsight was enlarged to accommodate emotional properties, in what has been coined *affective blindsight*. To verify behavioural influence of stimuli in the lack of consciousness of the same, experiments on this particular aspect of blindsight use conditioning techniques and covert responses. Again, the dorsal stream is involved (e.g.: de Gelder et al., 1999; Hamm et al, 2003); supporting the hy-

⁶ As a matter of fact, all these capabilities were shown in the first studies by Weiskrantz and colleagues (see Weiskrantz, 1986); studies have since been refined. For a recent assessment of blindsight, see Cowey (2004).

pothesis of an affective blindsight is data indicating that the amygdala or, more specifically, the SC-pulvinar (see fig. 1) visual pathway to the amygdala does without a cortical representation (i.e. consciousness) in categorizing and responding to emotional stimuli, namely fear (Liddell et al., 2005).⁷

Since the dubbing of the condition, blindsight has proved to be a source of controversy with much opposition falling particularly on the possibility of spared islands of visual cortex to explain the ‘residual awareness’ verified in the patients observed (e.g.: Campion et al., 1983). A. Ptito and S. E. Leh (2007) have just dismissed this obstacle by testing on patients with blindsight who have undergone hemispherectomy, the complete removal or deaf-feration of a whole cerebral hemisphere. In the cases with which they worked, the occipital lobe, where V1 is located, had either been removed or disconnected from the rest of the brain, thus contradicting the possibility of islands of spared normal vision, and confirming the hypothesis that preserved pathways into the extraestriate cortex bypassing V1—i.e., the dorsal stream—are enough for perception, though not for consciousness.

Left visuo-spatial neglect — A form of extinction, an attentional disorder following brain damage in which patients are not aware of visual stimuli on the side of space opposite the damaged hemisphere, this dysfunction is characterized by the failure to perceive not only objects on the left side but *the* left visual side of space.⁸ For instance, in both spontaneous drawings and copies, patients completely leave out the left side of no matter what they draw or copy; this is particularly observable in drawings/copies of watches, in which either only the numbers 12 and 1 to 6 are drawn as in their usual positions, or the numbers 1 to 12 are all located on the right side of a circle. Other salient signs of this dysfunction are shaving only the right side of the face, eating only the food on the right side of the plate even when hungry, and writing only on the right side of a page (see Halligan & Marshall, 1998 for abundant examples). This being so, it is not surprising that when asked to identify objects or shapes on their left side, they more often than not fail to do so correctly. However, studies (e.g.: McIntosh et al., 2004; Rice et al., 2006) have consistently shown that when asked to reach between obstacles they claim not to see or cannot identify, their trajectories, like those of

⁷ For a study on the role of the amygdala in emotional learning not directly connected with affective blindsight but confirming the same hypothesis see Morris et al. (1998).

⁸ It is important to specify this within the umbrella terms *hemagnosia*, *hemineglect*, *hemispacial neglect*, etc., that include dysfunctions in other sensory modalities.

normal subjects, take these obstacles into account; also unsurprisingly, when asked to point midway between two objects, they perform very poorly. This dysfunction is strategically (Schindler et al., 2004) compared with *optic ataxia*, impaired behaviour in space shown by patients with parietal lesions: patients with this impairment do not take obstacles into account in their reaching trajectories, though they can point midway between two obstacles in a normal way. This double dissociation seems to corroborate the hypothesis of the existence of the two visual systems, dorsal and ventral, or ‘vision for action’ and ‘vision for perception’ (Milner & Goodale, 2007, p. 774).

It is important to try to realize to what extent there is *no* consciousness of the left visual side of space in this impairment, and no one better to tell us than someone affected by it; P.P., a patient quoted by Halligan & Marshall (1998, p. 360), says:

I think they thought I was definitely, deliberately not looking to the left. It was painful looking to the left. ... People think you are not looking... you are neglecting to look but it's not there. If it's not there you are not neglecting it.

Prosopagnosia — Also known as *face-blindness*, this is an extremely impairing condition both in psychological and sociological terms in which patients fail to recognize individual faces, even those of family members, friends, and often their own. Interestingly enough, prosopagnosics can know that they are looking at a face, and they can perfectly identify a nose, eyes, etc, but are unable to recognize the owners; this is in tune with other findings that reveal other types of stimuli that they can recognize whose tokens they cannot identify: these are automobiles, clothes of the same type and general shape, similarly looking (shape and volume) foods, and specific animals within a group; for instance, a prosopagnosic will identify an object as being a car without being able to recognize her/his own; a farmer will clearly know what a cow is without being able to recognize the individual cows s/he owns and sees daily (cf. Damásio, 1985, p. 134). It is important to remark that neither *other* cognitive skills nor complex visual abilities show any degradation whatsoever besides an acquired *achromatopsia* (loss of colour vision) that commonly accompanies prosopagnosia.

Associated with bilateral lesions involving the central visual system in the mesial occipitotemporal region, this impairment is of particular interest for functionalist hierarchical/in

parallel approaches (e.g.: Bruce & Young, 1986) that hypothesize the break-down of the recognition process at some specific module or unit. The current trend is to see prosopagnosia as a material-specific memory impairment, namely the failure to activate memories relative to specific visual stimuli (e.g.: Damásio, 1985).

What is of import for us is the fact that though, as seen, prosopagnosics fail to overtly recognize the faces of people they know, measures of covert recognition, namely skin conductance responses, have consistently shown that the unrecognized faces do indeed cause emotional arousal in patients (see Bauer, 1984; Tranel & Damásio, 1988; for a double dissociation, see Tranel et al., 1995). To my knowledge, no studies in prosopagnosia have been conducted aiming at showing behaviour beyond the automatic or reflex level, but the results obtained with measures of covert responses are enough to allow us to speak of perception and behaviour without consciousness. Once again, the sparing of the dorsal visuolimbic pathway vs. the impaired ventral stream appears to lie at the root of this unconscious recognition (cf. Bauer, 1984, p. 465f).

2.2. Unconscious Cognition

The experimental research being carried out today on unconscious cognition follows in the footsteps of that started at the end of the 19th century as studies “on small differences of sensation” (e.g.: Peirce & Jastrow, 1884; Dunlap, 1900) and ‘suggestion,’⁹ which evolved into experiments on subliminal perception (subception) and on “behaviour without awareness” (e.g.: Lazarus & McCleary, 1951; Adams, 1957); grosso modo. This culminates in the early 60s in the putting together of the findings in this field, and especially in the realization that awareness actually restricts perception and cognition. An important study (Spence & Holland, 1962) presented the following results: a) registration of stimuli is independent of awareness; b) the effect of impoverished (subliminal) stimuli varies inversely with their intensity; c) impoverished stimuli follow laws independent from those that rule conscious perception; finally, d) awareness of word stimuli restricts their effect on recall of other associated words.

⁹ This comprised a scientific interest in hypnosis. See for instance Binet (1900) and Sidis (1898).

The study of behaviour without awareness was very much restricted when only overt behaviour was involved, but with the advances in psychophysiology it became possible to ‘indirectly’ verify unconscious perception in the physiological responses produced by a variety of unconscious psychological states via a plethora of methods and appliances (e.g.: GSR, fMRI, EEG, etc). In this vein is the research now being carried out in the very recent *somatic marker hypothesis*, a theory defending that in humans decisions are made securing advantageous effects to a great extent—though not wholly—unconsciously, and this due to somatic markers, physiological affective states associated to certain stimuli; it, too, postulates a crucial role to the amygdala, and studies have been made with patients with lesions both in this and in the ventromedial prefrontal cortex. A. R. Damásio, his conceiver, has carried out studies both showing the existence of such markers and the effects that lesions on the amygdala and on the prefrontal cortex might have in this context; thus they can be integrated in the vast experimental research aiming at proving the existence of unconscious cognition. In a different vein, though sharing the same objectives, is the research using language, namely artificial grammars, prompted in the late 60s by A. S. Reber; this investigation explicitly aims at showing that there actually is something like implicit, or unconscious knowledge, and it therefore appeals more often to consciousness, actively taking part in the effort to conceptualize it (e.g.: Dienes & Perner, 2001; 2003).

The Somatic Marker Hypothesis — It is important to emphasize the fact that this is not a straight-away hypothesis on unconscious cognition; it simply is stated in such terms and has prompted such research that it contributes to research on unconscious knowledge. According to his conceiver, the hypothesis was originated by “intriguing observations” in patients with focal damage to the prefrontal region, namely to its ventral and medial aspects, who displayed a pattern of abnormal decision making simply not verified in the premorbid period. As Damásio himself summarizes the hypothesis,

The key idea in the hypothesis is that ‘marker’ signals influence the process of response to stimuli, at multiple levels of operation, some of which occur overtly (consciously, ‘in mind’) and some of which occur covertly (non-consciously, in a non-minded manner). The marker signals arise in bioregulatory processes, including those which express themselves in emotions and feelings, but are not necessarily confined to those alone.

This is the reason why the markers are termed somatic: they relate to body-state structure and regulation even when they do not arise in the body proper but rather in the brain's representation of the body.¹⁰ (Damásio, 1996, p. 1413)

Of direct interest to us is the supporting thesis that this process has qualitatively different results depending on whether it is overt (conscious) or covert (unconscious): in the latter case, the somatic state is a biasing signal *influencing* the cognitive process (more specifically through a non-specific neurotransmitter system such as dopamine); if overt, then the process influences cognition at a conscious level (Damásio, 1996, p. 1415).

Connected to this still controversial hypothesis (e.g.: Colombetti, 2008) is the Iowa Gambling Task (Bechara et al., 1999; see also Bechara et al., 1994): subjects are confronted with four decks of cards, two of them (the 'bad' decks *A* and *B*) yielding immediate gain but larger future loss, and the other two (the 'good' decks *C* and *D*) yielding lower immediate gain but a smaller future loss. This is intended to test the extent to which subjects are capable of making decisions with a view to securing long term gain and avoiding long term loss. While aiming at showing the disparity of behaviour between normal control subjects, a group of patients with bilateral damage to the ventromedial prefrontal cortex, and yet another group of patients with bilateral damage to the amygdala, the results that interest us more immediately are precisely those verified with the normal control subjects: a physiological measure (skin conductance response, SCR) showed that normal subjects, when becoming experienced with the task, started to generate SCRs prior to the selection of the cards; remarkably, when choosing a card from the 'bad' decks, their SCRs were almost twice as pronounced as when choosing from a 'good' deck. A previous study (Bechara et al., 1997) had led to the distinction of four periods of performance throughout the task, from the first to the last trial: these were the pre-punishment, pre-hunch, hunch, and conceptual periods; it was now verified that during the pre-hunch period normal subjects already showed a substantial rise in SCRs, and that even those few (30%) who did not eventually reach the conceptual period performed in an advantageous way. Bechara and Damásio conclude from this that "conscious knowledge alone is not sufficient for making advantageous decisions," (Bechara & Damásio, 2005, p. 348) and that the unconsciously biased cognitive process dictated in great measure this advantageous behaviour. Damásio does not use a technical philosophical definition of knowl-

¹⁰ What he calls the "as if body loop."

edge, seeming to adopt instead one that equates knowledge with *available* information. This is clear in his quadripartite division of knowledge, ranging from the mere innate and acquired knowledge to what appears to be a conceptual, scientific level, and in his postulation that *availability* of knowledge is not necessarily the same as *accessibility*, the representation in images and the ‘becoming minded’ of the results of all motor responses (Damásio, 1996, p. 1414).

Unconscious Cognition — For reasons internal to psychology that show how strong an influence behaviourism left in the field,¹¹ the terms ‘conscious’ and ‘unconscious’ tend to be substituted more often than not by ‘explicit’ and ‘implicit,’ respectively. Thus, talk is usually of ‘implicit knowledge’ and ‘implicit cognition,’ but the above are synonymous for our purposes.¹² In this context, research on unconscious cognition is since its inception connected to investigation on memory, namely on the distinction between an implicit (also: procedural) and an explicit (also: declarative)¹³ memory, though this connection, or cross-reference has only recently been unambiguously formulated (Berry & Dienes, 1991). A. S. Reber’s first studies on unconscious cognition in the late 60s (Reber, 1967; 1969), studies that were explicitly aimed at proving that indeed such a kind of cognition is a reality, namely as far as language learning is concerned, set the tone for a vast research with finite artificial grammars that has since then followed suit. And that has since been refined in methodology and theory: concerning the former, and namely in face of strong opposition, procedures and measuring methods have been refined and well-defined criteria have been established and adopted, (cf. Destrebecqz & Peigneux, 2005); as for the theory, this, depending on the different assumptions of the diverse researchers, reduces to the claim that there indeed is implicit, or unconscious knowledge. This claim, however, does not merely state that there is unconscious knowledge when the learning process fails to reach consciousness: the contention is that it is *qualitatively* different from explicit, or conscious cognition, not being, in principle, accessi-

¹¹ In fact, not only did behaviourism dismiss the psychoanalytic postulation of the/an unconscious, but it also shunned the very notion of consciousness.

¹² The term ‘tacit’ is also seen as a synonym for ‘implicit.’

¹³ The terminology ‘procedural vs. declarative’ is openly founded on a distinction between knowledge-*that* and knowledge-*how* that has no philosophical relevance, in my view. Studies on amnesia seem to favour this distinction, in that procedural memory appears to be intact in amnesia, while declarative memory seems to be impaired (see for instance Graf et al., 1984). However, I believe that this can be explained by the fact that unconscious knowledge is more robust, namely for evolutionary reasons (see for instance Reber, 1989, p. 232).

ble.¹⁴ In other words, unconscious cognition just is believed to be altogether different from conscious learning; that head-injured patients may show impairments in tests of unconscious knowledge seems to corroborate this hypothesis (Barker et al., 2006).

Roughly, studies measure the extent to which subjects exposed to finite artificial grammars (usually alphabet-based) show actual knowledge of the same in spite of the fact that they are not consciously learning the grammar. For instance, a subject may be confronted with lists of grammatical sentences or formulas of the artificial languages in question and later, confronted with new sentences or formulas, asked to distinguish the grammatical from the non-grammatical ones. Results consistently show that subjects perform above chance despite claiming that they are just guessing. It appears to be the case that the subjects simply do not know that they know, i.e., they lack metaknowledge about themselves (e.g.: Dienes & Perner, 2002).

This research carried out with artificial grammars, as well as other studies involving implicit verbal knowledge, is fundamental to dismiss the obstacle that unconscious knowledge is merely—logically—inferable knowledge (Dennett, 1983), or knowledge that albeit unconscious may be brought to consciousness in some way. In fact, one of the advantages of using artificial grammars is that their complexity (see fig. 2) simply precludes any possibility of the subjects actually learning their rules in a conscious manner in the time allotted to the tests.

¹⁴ This radically distances the meaning of ‘implicit knowledge’ in psychology and in knowledge management in that for the latter implicit knowledge is seen as a manna waiting to be made explicit!

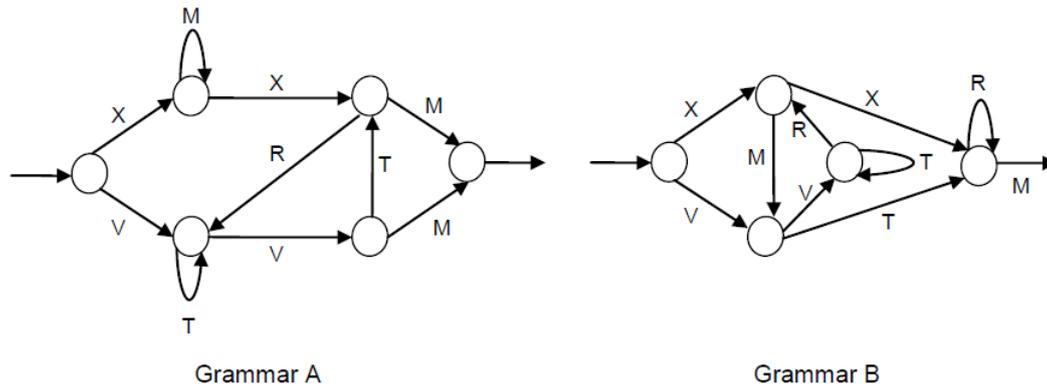


Fig. 2 – Two synthetic grammars, both in Dienes et al. (1995); grammar A follows Reber (1969). A few examples of strings allowed by grammar A follow: xmxrtvtm, vtvtrm, xmmxrvm, vtvtm, xxrvtm, etc.

3. Unconscious Knowledge and Epistemic Status

3.1. *Belief and Belief Ascription*

As seen in the beginning of this paper, to speak of unconscious knowledge is to attribute a positive epistemic status to unconscious beliefs. The attribution is trickier than it might appear at first sight, as shown, because when seeing unconscious beliefs in this light one is claiming what for many is a paradox: that a subject may not know that s/he knows p , or even believes <that> p . But we have evidence that strongly suggests that this is not a paradox: that subjects taking part in experimental studies in unconscious cognition more often than not actually ‘get it right’ or show automatic signs of correct recognition in spite of claiming not to have a clue of what is going on, and that patients with lesions to the ventral visual stream are capable of somehow discriminating stimuli—as shown in navigation with negotiation of obstacles and in automatic emotional responses—while claiming that they do not perceive the stimuli in question seems to provide evidence that there indeed is knowledge involved.

This, of course, depends on a conception of positive epistemic status, and on a theory of belief and belief ascription without which not. Let me thus start with the latter.

Humans, like all living organisms, act and react according to stimuli from both the external environment and the ‘self,’ that is, living organisms establish reciprocal relations between themselves and the environment (e.g.: an internal stimulus signalling hunger will prompt an organism to search for food; the perception of food will activate the hunger instinct). What seems to set humans ‘apart’ from other organisms is the fact that they can verbalize both the stimuli presented to them and the effects those stimuli have on them: they hold conscious beliefs about the world and themselves. But this is, so to say, the tip of the iceberg, given that more often than not we simply act without formulating any beliefs regarding our actions; moreover, many of our actions, including social behaviour, remain unknown to ourselves (e.g.: electrodermal activity, pupular dilation, hormone production, etc).

Given that the vast majority of our actions, conscious or unconscious, simply do not require that a conscious belief be held, but in view of the fact that when probed we can verbalize the beliefs behind many of them, it seems more appropriate to see belief as primarily a willingness to act in a certain way, and only secondarily, and accessorially, as a propositional attitude of the kind *S A that p*, where *S* is a subject, *A* is the attitude (believes, thinks, etc), and *p* is a sentence expressing a proposition. Namely, beliefs take this form when in questions of doubt, or in other special contexts (logic, for instance), in which cases a belief <that> *p* is the psychological attitude of holding *p* as true.

If we accept, in humans, this equation between belief and willingness to act in a certain way, then we must accept that all our behaviour, excluding perhaps reflex behaviour in states of deep coma or in other non-responsive states, is in principle grounded on beliefs, whether they be held consciously or unconsciously. In other words, we accept that behaviour is a *function* of belief. Thus, we might feel inclined to hold that beliefs are *causal* in that a displayed behaviour, or a propensity to behave in a certain way, reflects a particular belief.

Besides, it is very much evident that beliefs are *species-specific* (e.g.: that we know of, humans alone have ‘higher’ mathematical beliefs), as well as *social*, or *cultural* (e.g.: eating from the floor is not commonly held as a permissible behaviour for humans and, even when alone, people just do not eat from the floor); in other words, as a belief-holding species, *homo sapiens* has a vast yet finite repertoire of attitudes or behaviours that issue from commonly held/holdable beliefs. This means that beliefs are *sharable* and, therefore, *observable*, directly and/or indirectly. Given this, we can ascribe beliefs to agents who do not specify

them whether because they do not want or need to, or because they are not aware of holding any particular beliefs.

We are now ready to formulate the general principle of a theory of belief ascription based on the properties above of beliefs (to review: beliefs are causal/functional, species-specific, social/cultural, sharable, and (in)directly observable):

When an 'approved' belief-holding subject S belonging to a specific community acts in a specific way x because s/he believes <that> p, S, a member of the same community and assumed to be also an 'approved' belief-holding subject, acting in a similar way x, ceteris paribus, may be said to also believe <that> p.*

Although one may err more or less marginally in the ascription of beliefs, there is always a one-to-one correspondence between behaviour and belief, at least at a basic level, that allows us of legitimately ascribing beliefs to other agents. Take for instance this situation: Brenda was reading a book when the door bell rang. She went to the door and opened it. She might have opened the door for a vast number of *other* reasons (she was expecting someone; she was actually going out; etc) but among all those possible beliefs relevant for Brenda's opening the door is the belief that one can open doors/doors can be opened. This allows us to complete the principle above with an appendix meant to ward off much antagonism:

Acting in a specific way x does not necessarily entail that S only believes that p, but it entails that S also believes that p.

Armoured with this theory, we can now go back to our case studies and view them from a belief ascription perspective: we ascribe to a blindsight patient to whom an **X** is presented on a screen and, when forced, correctly 'guesses,' the implicit belief that the shape on the screen is an **X**; to the patient suffering from left visuo-spatial neglect that shows to be capable of avoiding obstacles when reaching for an object the unconscious belief that there

are obstacles to avoid is ascribed; we attribute to the patient with prosopagnosia whose SCRs show significant increases when presented with the faces of close relatives and friends the unconscious belief that the owner of the face is familiar. In sum, we attribute to patients with lesions to the ventral visual stream the faculty of forming and holding beliefs of which they are not nor can they be aware, thanks perhaps to the preserved functions of the dorsal visual stream. Similarly, we ascribe participants in the Iowa Gambling Task who early in the task, long before any hunch concerning what might be going on, begin to choose the ‘good’ decks the implicit belief that those are the decks to choose with a view to long term gain; and we attribute to participants in studies with artificial grammars who correctly ‘guess’ above chance which sentences are (non-)grammatical the unconscious belief that they actually are so.

But do they have knowledge?

3.2. *Positive Epistemic Status*

It all depends on one’s concept of positive epistemic status. For my part, I argue that a belief has positive epistemic status when it is a justified true belief. So far, nothing appears to distinguish my stance from that of a vast number of philosophers that take knowledge to be justified true belief. However, I see the canonical tripartite analysis of knowledge, commonly accepted by them with more or less reservations, as irredeemably flawed and incapable of providing us with a satisfactory—if not correct—account of knowledge.

I am not the first to think that this analysis of knowledge is wrong; E. Gettier (1963) famously showed that there is no rescue for it because the concept of justification expressed in its condition (iii) simply might have nothing to do with its condition (i), the truth condition. In other words, true beliefs may fail to yield knowledge because of a wrong justification. Gettier gave two examples to illustrate this: in the first, Smith, an applicant for a job, has the true belief *e* that the man who will get it has ten coins in his pocket, a belief that he has inferred from the belief *d* that Jones, another applicant for the job, will get it, because he was not only told so by the president of the company, but he has also counted the coins in Jones’ pocket and verified that there are precisely ten coins there. But Smith, who, unbeknownst to himself, has ten coins in his pocket, too, actually gets the job. The example is indeed florid,

but Gettier's point is that Smith's belief that the man who will get the job has ten coins in his pocket, though true, does not yield knowledge, because Smith is not justified in holding it as true. And this is so because he inferred it from the false belief *d*. The second example is even more florid, so much so that I invite the reader to have a go at its original formulation, but what matters to us is that, again, Smith has a true belief, and yet that belief does not count as knowledge in his case due to, again, a failure in justification; and again, the justification is of a solely logical type: in the first example, a wrong inference was the culprit; in the second, what is at stake is logical entailment of the form $P \rightarrow (P \vee Q)$ where *P* is actually false but the entailment is nevertheless valid.¹⁵

As claimed, I, too, see knowledge as justified true belief, but instead of trying to rescue an analysis of knowledge that I see as deontic and transcendental (meaning: it falls prey to a concept of a priori truth, or truth come what may), namely by trying to save condition (iii), I propose, firstly, that the order of the conditions be changed, and secondly, that their number be shortened. In fact, since what has been dubbed *the Gettier problem*, the vast majority of efforts has been concentrated in securing condition (iii) by 'degettierizing' it; this opened an abyss not only between those who defend that justification is internal and those who see it as external, but also between those who explicitly require justification and those who only implicitly do so. What follows is the *summary* illustration of how things stand as far as this issue is concerned:

¹⁵ I am here using material implication as more or less equivalent to entailment; actually, the former, rather than \models , seems to be intended by Gettier.

	Internalism	Externalism
Justification required	J-Internalism: justification is directly recognizable; at any time a subject is in position to know that his/her belief is justified (Chisholm, 1977; 1989)	J-Externalism: justification is not directly recognizable; it is not the case that at any time a subject can know that his/her belief is justified J-Reliabilism: a subject is justified in holding a belief iff her/his belief was produced by a reliable cognitive process (in a way that degettierizes it) (Plantinga, 1988; Nozick, 1981; Goldman, 1979)
No justification required	(Internalism just is the stance that claims that justification is internal, therefore: Ø	K-Externalism & K-Reliabilism: the same as J-Externalism and J-Reliabilism with the wrong conviction that it is not so, i.e., that they do not require <i>justification</i> for knowledge (thus the K) (Dretske, 1989)

Table 1: Knowledge and Justification

As I see it, though these efforts have perhaps contributed to highlight even more problems concerning knowledge, they have all without exception been in vain in very practical terms in that they failed to offer a notion of justification and knowledge that can be adopted by other fields in which knowledge matters. The fact is that they have already adopted one, and by ‘they’ I mean both the empirical sciences and the more formal-mathematical ones, and they did it because they could not advance without it. Because among those working in those areas the more philosophically-informed realize that epistemology reclaims knowledge as *its* subject, and a surface reading of the literature will show that there is no consensus—i.e., epistemology has failed to provide a positive account of knowledge—, they cleverly opt for a

tacit concept of knowledge in their disciplines. Commonly, knowledge is for them the information that holds, or works, and they go on checking their hypotheses against the external world and boasting to possess knowledge in all instances in which their beliefs appear to match the world. What I mean to say is that the battleground that opposes externalists and internalists simply is of no interest for science, because it does not matter at all, for the scientist, whether the justification of our beliefs lies outside or inside us. Probably, different specific problems will require different external and/or internal justifications, if any of these, without for that changing the nature of knowledge.

But justification matters, because that is what validates our beliefs, i.e., makes them true. I propose that a belief has positive epistemic status when it is a justified true belief, and that a belief is true when it is justified. This is to say that truth is a consequence of justification. Thus, in a first approach, I propose that the order of the conditions of the tripartite analysis of knowledge be altered in the way already seen in the beginning of this paper:

- (i) S believes <that> p .
 - (ii) S is justified in believing <that> p .
 - (iii) p is true.
- S knows p .

Now for the justification proper: I hold that *a belief is justified when it does not contradict the external world*.¹⁶ Bear in mind, however, that the fact that the external world is involved does not make of this just another externalist claim, because here it is not the world that is in question, but the beliefs held concerning it: only beliefs can be justified or unjustified, and therefore true or false. And they are true when they *work*. Regarding the world qua world, there might be a vast amount of facts about it that might be true... if conceived as beliefs. This entails that beliefs are not to be separated from the subjects who hold them and from the situations in which they are held, what is the same as to say that no proposition is a priori true and thus true come what may, as the canonical analysis of knowledge wants us to believe. Moreover, that beliefs are justified and therefore true does not mean that whatever it is that

¹⁶ I am here claiming that all beliefs respect the external world; beliefs about the self are but beliefs concerning a part of the world, which is external at large.

they hold about the world is actually true *of the world*: it might or it might not, and there simply is no possible way of ever finding that out.¹⁷

If one wants to see this as an advocating of epistemic luck, by all means: indeed, it is irrelevant whether our epistemic justification is grounded on a non-contradiction of the external world or on the capturing of *its truth*. But while the latter is an obscure metaphysical concept that attributes to our beliefs the property of being a copy of reality, or of at least somehow corresponding to it, the former invokes solely the clear and workable concept of non-contradiction. This is not the logical concept encapsulated in the sentence $\neg(P \ \& \ \neg P)$, which expresses incompatibility between two or more propositions; it is more akin to the intuitionistic notion according to which a contradiction arises when a proposition has no proof; as for our ‘proof conditions,’ they are the external world, given that our statements/beliefs respect it, and it alone.¹⁸

The importance of the examples of unconscious knowledge given above is precisely that of helping to eliminate from the analysis of knowledge the metaphysical obscurity brought on by a notion of a priori truth that implicitly, if not explicitly, requires that a subject be conscious both of his/her being in a state of believing and of the contents of the beliefs held. Given that they obviously fail to fulfil the consciousness requirement while fulfilling the belief ascription conditions, they strongly suggest that if we cling to this deontic-

¹⁷ The astute reader will have realized by now that I am doing without the concept of truth altogether: in fact, truth, in the perspective defended here, just is justification, and therefore truth as a condition simply vanishes from the analysis of knowledge that has now the abbreviated form as follows:

- (i) S believes <that> p .
- (ii) S is justified in believing <that> p .
 S knows p .

I am not for this repudiating my claim that knowledge is justified true belief; as seen, truth just is justification.

¹⁸ It is important to emphasize the fact that the concept of non-contradiction meant by me is *not* an intuitionistic concept, being merely inspired by the way contradiction (falsity) is seen by intuitionists. Briefly, for intuitionism, contradiction and negation are connected in the sense that $\neg A$ is a shorthand for $A \rightarrow \perp$, i.e., there is no proof for A , and the assumption that there is one leads to a contradiction. Now, let P be any arbitrary belief; then, $\neg P$ is the realization that P entailed a contradiction; this realization is attained via a process of empirical—or empirical-like—verification. See (Augusto, 2008).

transcendental analysis of knowledge, then we leave unconsciously motivated behaviour—most human behaviour, that is—unaccounted for in doxastic and epistemic terms. This would entail that both doxastic and epistemic situations are special circumstances in human action.

They are not, as any basic study on memory in normals will show. It is precisely because the experimental sciences to a great extent equate knowledge with memory (knowledge is information that can be acquired, stored, and retrieved) that they unproblematically speak of unconscious knowledge. And once this metaphysical obscurity is removed, the way is paved for a reformulation of a concept of justification that is capable of integrating unconscious beliefs in a positive analysis of knowledge.

4. Conclusion

To sum up: in order to have positive epistemic status, a belief <that> *p* must be justified in that it does not contradict the external world. This is the only requirement. That the belief <that> *p* is held consciously, or unconsciously—as the vast majority of our beliefs are held—is of no consequence regarding their epistemic status. Perhaps all human activity aims to conduce to knowledge, and conscious doxastic states are the exception, not the rule, for that end. To deny unconscious beliefs the same epistemic status of conscious beliefs is to persist in an analysis that has shown itself incapable of giving a positive account of knowledge. This is only possible within an epistemology that neglects the experimental sciences that share its object.

I wish to thank Fundação para a Ciência e a Tecnologia (FCT) for the postdoctoral fellowship that funded the research conducing to this paper, as well as the University of Sussex for the excellent working conditions provided to me as a Visiting Research Fellow in the Department of Psychology and in the Centre for Research in Cognitive Science (COGS). I also wish to thank Zoltán Dienes for reading the whole paper and making valuable suggestions.

REFERENCES

- Adams, J. K. (1957). Laboratory studies of behavior without awareness. *Psychological Bulletin*, 54:5, 383-405.
- Augusto, L. M. (2008). Putting the horse before the cart: A pragmatist analysis of knowledge (forthcoming).
- Barker, L. A., Andrade, J., Romanowski, C. A. J., Morton, N. & Wasti, A. (2006). Implicit cognition is impaired and dissociable in a head-injured group with executive deficits. *Neuropsychologia*, 44, 1413-24.
- Bauer, R. M. (1984). Autonomic recognition of names and faces in prosopagnosia: A neuropsychological application of the Guilty Knowledge Test. *Neuropsychologia*, 22:4, 457-69.
- Bechara, A., Damásio, A. R., Damásio, H. & Anderson, S. W. (1994). Insensitivity to future consequences following damage to the human prefrontal cortex. *Cognition*, 50, 7-15.
- Bechara, A., Damásio, H., Tranel, D. & Damásio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, 275, 1293-5.
- Bechara, A., Damásio, H., Damásio, A. R. & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *The Journal of Neuroscience*, 19:13, 5473-81.
- Bechara, A. & Damásio, A. R., (2005). The Somatic Marker Hypothesis: A neural theory of economic decision. *Games and Economic Behavior*, 52, 336-372.
- Berry, D. C. & Dienes, Z. (1991). The relationship between implicit memory and implicit learning. *British Journal of Psychology*, 82, 359-373.
- Binet, A. (1900). *La Suggestibilité*, Paris: Schleicher.
- Bruce V. & Young, A. W. (1986). Understanding face recognition. *British Journal of Psychology*, 77, 305-27.
- Campion, J., Latto, R. & Smith, Y. M. (1983). Is blindsight an effect of scattered light, spared cortex, and near-threshold vision? *The Behavioral and Brain Sciences*, 6, 423-486.
- Chisholm, R. (1977). *Theory of Knowledge*, 2nd ed., Englewood Cliffs: Prentice Hall. (3rd ed., 1989).
- Colombetti, G. (2008). The Somatic Marker Hypothesis, and what the Iowa Gambling Task does and does not show. *The British Journal for the Philosophy of Science*, 59:1, 51-71.

- Cowey, A. (2004). Fact, Artefact, and Myth about Blindsight. *The Quarterly Journal of Experimental Psychology*, 57A:4, 577-609.
- Damásio, A. R. (1985). Prosopagnosia. *Trends in Neurosciences*, 8, 132-5.
- Damásio, A. R. (1996). The Somatic Marker Hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London B*, 351, 1413-20.
- Dennett, D. C. (1983). Styles of mental representation. *Proceedings of the Aristotelian Society*, 83, 213-26.
- Destrebecqz, A. & Peigneux, P. (2005). Methods for studying unconscious learning. *Progress in Brain Research*, 150, 69-80.
- Dienes, Z., Altmann, G. T. M., Kwan, L. & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21:5, 1322-1338.
- Dienes, Z. & Perner, J. (2001). When knowledge is unconscious because of conscious knowledge and vice-versa. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, 1-4 August, Edinburgh, Scotland. Mahwah, NJ: Lawrence Erlbaum Associates, 255-260.
- Dienes, Z. & Perner, J. (2002). A theory of the implicit nature of implicit learning. In R. M. French & A. Cleeremans (eds.), *Implicit Learning and Consciousness: An Empirical, Philosophical, and Computational Consensus in the Making?* New York: Psychology Press, 68-92.
- Dienes, Z. & Perner, J. (2003). Unifying consciousness with explicit knowledge. In L. Cleeremans (ed.), *The Unity of Consciousness*. Oxford: Oxford University Press (pp. 214-232).
- Dretske, F. (1989). The need to know. In M. Clay & K. Lehrer (eds.), *Knowledge and Skepticism*. Boulder: Westview Press, 89-100.
- Dunlap, K. (1900). Effect of imperceptible shadows on the judgments of distance. *Psychological Review*, 7, 435-453.
- de Gelder, B., Vroomen, J., Pourtois, G. & Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *NeuroReport*, 10, 3759-63.
- Gettier, E. L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23, 121-3.
- Goldman, A. (1979). What is Justified Belief? In G. S. Pappas (ed.), *Justification and Knowledge*, Dordrecht: Reidel.

- Goodale, M. A. & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15:1, 20-25.
- Graf, P., Squire, L. R. & Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10:1, 164-178.
- Halligan, P. W. & Marshall, J. C. (1998). Neglect of awareness. *Consciousness and Cognition*, 7, 356-380.
- Hamm, A. O., Weike, A. I., Schupp, H. T., Treig, T., Dressel, A. & Kessler, C. (2003). Affective blindsight: Intact fear conditioning to a visual cue in a cortically blind patient. *Brain*, 126, 267-275.
- Lazarus, R. S. & McCleary, R. A. (1951). Autonomic discrimination without awareness: A study of subception. *Psychological Review*, 58, 113-122.
- Liddell, B. J., Brown, K. J., Kemp, A. H., Barton, M. J., Das, P., Peduto, A., Gordon, E. & Williams, L. M. (2005). A direct brainstem-amygdala-cortical 'alarm' system for subliminal signals of fear. *Neuroimage*, 24:1, 234-43.
- McIntosh, R. D., McClements, K. I., Schindler, I., Cassidy, T.P., Birchall, D. & Milner, A. D. (2004). Avoidance of obstacles in the absence of visual awareness. *Proceedings of the Royal Society of London B*, 271, 15-20.
- Milner, A. D. & Goodale M. A. (2007). Two visual systems re-viewed. *Neuropsychologia*, 46:3, 774-85.
- Morris, J. S., Öhman, A. & Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, 393, 467-70.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Peirce, C. S. & Jastrow, J. (1884). On small differences of sensation. *Memoirs of the National Academy of Sciences*, 3, 73-83.
- Plantinga, A. (1988). Positive Epistemic Status and Proper Function. In J. E. Tomberlin (ed.), *Philosophical Perspectives 2*, Atascadero, CA: Ridgeview Publishing Company.
- Ptito, A. & Leh, S. E., (2007). Neural substrate of blindsight after hemispherectomy. *The Neuroscientist*, 13:5, 506-18.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 6, 855-63.
- Reber, A. S. (1969). Transfer of syntactic structure in synthetic languages. *Journal of Experimental Psychology*, 81:1, 115-119.

- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118:3, 219-235.
- Rice, N. J., McIntosh, R. D., Schindler, I., Mon-Williams, M., Démonet, J.-F. & Milner, A. D. (2006). Intact automatic avoidance of obstacles in patients with visual form agnosia. *Experimental Brain Research*, 174, 176-188.
- Schindler, I., Rice, N. J., McIntosh, R. D., Rossetti, Y., Vighetto, A. & Milner, A. D. (2004). Automatic avoidance of obstacles is a dorsal stream function: evidence from optic ataxia. *Nature Neuroscience*, 7:7, 779-784.
- Sidis, B. (1898). *The Psychology of Suggestion*, New York: Appleton.
- Spence, D. S. & Holland, B. (1962). The restricting effects of awareness: A paradox and an explanation. *Journal of Abnormal and Social Psychology*, 64:3, 163-174.
- Tranel, D. & Damásio, A. R., (1988). Non-conscious face recognition in patients with face agnosia. *Behavioural Brain Research*, 30, 235-49.
- Tranel, D., Damásio, H. & Damásio, A. R. (1995). Double dissociation between overt and covert face recognition. *Journal of Cognitive Neuroscience*, 7:4, 425-32.
- Vakalopoulos, C. (2005). A theory of blindsight—the anatomy of the unconscious: a proposal for the koniocellular projections and intralaminar thalamus. *Medical Hypotheses*, 65, 1183-90.
- Weiskrantz, L. (1986). *Blindsight: A Case Study and Implications*, Oxford: Oxford University Press.