A Thoroughly Empirical Approach To Consciousness

Bernard J. Baars
The Wright Institute
2728 Durant Ave.
Berkeley, CA 94704. USA

baars@cogsci.berkeley.edu

PSYCHE, 1(6), August 1994
http://psyche.cs.monash.edu.au/v2/psyche-1-6-baars.html

Keywords: automaticity, Thomas Nagel, conscious, contrastive analysis, awareness, subjective experience, global workspace theory, phenomenology, unconscious

> "The study ... of the *distribution* of consciousness shows it to be exactly such as we might expect in an organ added for the sake of steering a nervous system grown too complex to regulate itself..."
>
> William James, (1890/1984) *The Principles of Psychology*. (Emphasis in the original.)
>
> "For many years after James penned *The Principles of Psychology* ... most cognitive scientists ignored consciousness, as did almost all neuroscientists. The problem was felt to be either purely "philosophical" or too elusive to study experimentally ... In our opinion, such timidity is ridiculous."
>
> Francis Crick and Christof Koch, "The problem of consciousness," *Scientific American*, September, 1992.

Abstract: When are psychologists entitled to call a certain theoretical construct "consciousness?" Over the past few decades cognitive psychologists have reintroduced almost the entire conceptual vocabulary of common sense psychology, but now in a way that is tied explicitly to reliable empirical observations, and to compelling and increasingly adequate theoretical models. Nevertheless, until the past few years most cognitive psychologists and neuroscientists avoided dealing with consciousness. Today there is an increasing willingness to do so.

But is "consciousness" different from other theoretical entities like "working memory" or "mental imagery"? Some argue that under no circumstances can empirical science speak of consciousness *as such*, while others claim that the scientific goal is "knowing what it is like to be a bat" --- to share an organism's conscious experience (Nagel, 1974). The Bat Criterion is ominously reminiscent of the protracted debate on the consciousness of ants and amoebas that caused so much uneasiness in psychology around 1900. It seems to demand that we first solve the mind-body problem *as a condition* of doing sensible science, and thereby creates the risk of endless, fruitless controversy. The endless philosophical debate about consciousness helped trigger the Behaviorist revolution about 1913, which threw out the baby of consciousness with the bathwater of perennial, circular debate. We've been that way; let's not go back to it.

This paper maintains that the position of behavioristic denial is far too restrictive, but that the Bat Criterion is far too demanding --- that in fact, we only need to specify comparable pairs of psychological phenomena that differ only in the fact that one member of any pair is conscious, while the other is not. This "method of contrastive analysis" is a generalization of the experimental method, with consciousness as a variable whose interaction with other psychological and biological phenomena can be assessed in standard ways. As usual in science, this strategy is pragmatic: If it appears to yield sensible results, it can be a stepping-stone toward further understanding (Crick & Koch, 1992; Ericsson & Simon, 1984/1993).

This paper describes five sets of well-established pairs of phenomena that meet these criteria. Others are presented elsewhere, with more of a theoretical interpretation (Baars, 1983, 1988, 1993). Here I simply want to show that *any* adequate theory of conscious experience must satisfy these demanding but achievable empirical constraints.

1. Introduction

1.1 You, the reader, are now conscious of words in your visual focus, but not of the detailed syntactic and semantic analysis which your nervous system is performing quite expertly at this very moment. On reading the word "focus" in the previous sentence, you were very probably unaware of its alternative meanings. (Webster's Dictionary cites 9 different definitions). Likewise, you are not now conscious of a thought or image which came to mind just a few moments ago; of a sound or sight that could be perceived at this very instant, were it not for the fact that reading this sentence competes against awareness of that event; of the feeling of your clothing, the quality of the ambient light, that monotonous background noise, or any other constant, highly predictable stimulus. These currently unconscious thoughts, images and stimuli can often be made conscious "at will", unlike the syntactic and semantic processes mentioned above.

1.2 These examples illustrate the sense of the word "consciousness" we aim to understand: that is, focal consciousness of easily described events. According to Natsoulas (1978), "it is difficult to emphasize sufficiently the fundamental importance of consciousness in this sense. It is arguably our most basic concept of consciousness, for it is implicated in all the other senses."

1.3 Cognitive metatheory suggests that behavioral measures can be used to infer underlying explanatory constructs (Baars, 1986b). There have been those who argue that psychology should not be permitted to postulate explanatory constructs at all (e.g. Skinner, 1974; Ryle, 1949), but they are a dwindling minority. Given our current metatheory, cognitive theories typically postulate some system that performs information processing -- the internal manipulation of representations <1>. This approach has proved fruitful in dozens of areas including perception, memory, language, and imagery. The information-processing language is widely accepted in psychology and neuroscience, in linguistics, parts of philosophy, and computer science. The question is whether conscious and unconscious processes can be usefully treated as inferred constructs expressed in this language.

1.4 Several cognitive views of consciousness have appeared so far. Some stress the close connection between consciousness and limited-capacity mechanisms, including selective attention and working memory (e.g. Posner and Rothbart, 1991; Mandler, 1975, 1984; Baddeley, 1993). Others relate it to the control of action or executive functions (e.g. Shallice, 1972, 1978; Norman and Shallice, 1986; Hilgard, 1992). On the basis of his classic experiment on masked visual priming, Marcel (1983a) has suggested that consciousness is "an attempt to make sense of as much data as possible at the most functionally useful level ... (it) requires a constructive act whereby perceptual hypotheses are matched against information recovered from memory and serves to structure and synthesize information recovered from different domains."

1.5 All these contributions are significant. My own work has emphasized combining such evidence into a single, integrative type of organization called "contrastive analysis," which compares conscious vs. unconscious processes across numerous experimental domains. A vast amount of data along these lines is already available. Building on a contrastive analysis of the functional capabilities of conscious and unconscious processes, I have presented an integrative theory called Global Workspace (GW) Theory, a broad architectural notion of the nervous system in which conscious events are viewed as mental representations that have at least three necessary properties (a) they are *global*, so that their contents are disseminated throughout the nervous system; (b) they are *internally consistent*, because conflicting processes will compete each other out of consciousness; and (c) they are *informative*, in that they trigger widespread adaptation among specialized processors in the nervous system. This list of three features is not complete, but will give some idea of the productivity of the contrastive analysis approach. (See Baars, 1983; 1988, chapter 11; 1993ab; Newman & Baars, 1993) GW theory incorporates all the standard views of consciousness cited above very readily.

1.6 This paper will not go into these theoretical claims in any detail. Rather, it concentrates on the question of evidence. What kind of data could entitle theorists to claim that some construct corresponds to "consciousness" *as such*?

2. The search for criterial evidence

2.1 A great variety of evidence may have a bearing on consciousness -- hypnotizability, perceptual reaction time, memory measures, psychophysical scaling -- but the relationships between these measures and "consciousness per se" are probably indirect. Which measures are most relevant to understanding consciousness seems to be a great source of confusion. A review of the vast literature on event-related brain potentials concludes pessimistically that, although there are well-known ERP correlates of perception, decision-making, discrimination, and voluntary action, it is unclear whether any of them say anything about consciousness (Donchin, McCarthy, Kutas, and Ritter, 1983)! The sticking point seems to be that there is no explicit agreement on what to look for in trying to understand "real" consciousness. But of course the raw observations won't tell us anything until we have some notion of what to look *for*. To be sure that we are really dealing with consciousness, what do we need to know?

2.2 In philosophy of mind, Nagel (1974) has epitomized the problem of consciousness with the question: "'What is it like to be something?' It is certainly not like anything to be a brick or a hamburger; it certainly is like something to be you or me; and it seems to be like something to be a bat or a dog or a dolphin, if only we could figure out what. (This) question ... may in the end be the wrong question to ask, but it excellently captures the intuitions that constitute the challenge to a theory of consciousness. Until one's psychological ... theory explains how it can be like something to be something ... one's theory will be seriously incomplete." We will call Nagel's criterion the Bat Criterion of consciousness.

2.3 This paper will argue that at this moment in scientific history the Bat Criterion is, in Nagel's own words, "the wrong question to ask." It raises far more difficult issues than we can handle all at once. There is a much simpler source of criterial evidence to help define the type of empirical findings which any theory of consciousness must explain. Let us take these points one by one.

2.4 First, for scientific work, is the Nagel question a useful one to ask today? The view that consciousness is essentially the domain of access of the personal self has a good deal of intuitive appeal. It certainly fits our ordinary language in describing conscious experiences -- practically all statements about conscious experience use personal pronouns. as in "I saw a pussycat," "You are imagining that pain in your tummy," "She smelled a rat," etc. And no doubt we must eventually deal with the notion of self in an adequate scientific way. But all of those points do not make their argument *strategically* wise in our present state of understanding. Scientists who try to work with the Nagel criterion would have to define not just one, but two major theoretical constructs simultaneously -- consciousness *and* self. And if we succeeded in doing so, we would be left with two constructs which could not be defined independently of each other. We could never ask, for example, whether "self" could vary independently of conscious experience -- whether infants at an early stage of self-development experience a red ball in the same way an adult does. That sort of question depends on the conceptual and empirical separation of consciousness and self, and the Bat Criterion does not permit that.

2.5 Finally, the kinds of questions that emerge from the Nagel criterion -- What is it like to be a bat, a dog, a dolphin? -- are ominously reminiscent of the protracted arguments about the consciousness of ants and amoebas that caused so much trouble in psychology around 1900. Not that such "bat" questions are forever closed to scientific inquiry; but they certainly do not provide us with a modest, workable and consensus-building approach to the problem.

2.6 There is a simpler way. To operationalize a scientific construct we must at least be able to specify its conditions of occurrence and non-occurrence. In Newtonian mechanics we must be able to state the conditions of occurrence of gravitational force; in specifying a new chemical element or sub-nuclear particle we need to specify the conditions under which it can and cannot be found. In the case of consciousness, we can do this by carrying out a "contrastive analysis" of comparable conscious and unconscious events. By comparing two things that are very similar except for the fact that one is conscious while the other is not, we can hone in on just those elements that are uniquely associated with consciousness. Thus we should be able to specify the distinctive properties of conscious processes *as such*.

2.7 The distribution of consciousness

a) This course follows William James' recommendation that we consider the *distribution* of consciousness in attempting to understand its workings in the nervous system (see epigraph). It is much like the experimental method, with consciousness as the variable of interest while keeping everything else as constant as possible.

b) Quite a few authors have already discussed contrasting pairs of phenomena, though generally not as a method to be applied systematically across numerous domains, in the way advocated in this paper. Nevertheless, Marcel (1983ab) and Libet (1977) have investigated pre-perceptual vs. perceptual processes in terms of consciousness as the critical variable; likewise, blindsight has been explored as a source of insight into consciousness compared to normal visual object perception (Weiskrantz, 1988; Natsoulas, 1982)). There is by now a vast literature on subliminal priming, selective attention, habituation, processing of ambiguity, etc.

2.8 Again, the *general* point developed here about the question of criterial evidence has not been made elsewhere as far as I know. It is important to define the approach explicitly, because once we do so we can apply it to existing bodies of evidence in perception, selective attention, processing of ambiguities, sleep vs. waking, etc., which have been amassed over decades without reference to consciousness. This great body of evidence does, in fact, bound and constrain any viable theory of consciousness. Approached it in this spirit, we can use the data to establish a solid and wide-ranging empirical basis for theory.

2.9 In what specific cases can we carry out such a contrastive analysis? Tables 1-5 list a number of empirical contrasts between pairs of comparable conscious and unconscious events. This paper attempts a fairly complete catalogue of such contrasts, with an eye

toward describing as simply as possible the empirical basis for understanding conscious experience.

2.10 There is much to say about each contrastive pair, of course, but here we will focus on evidence that both the conscious and unconscious members of a contrastive pair can be viewed as "representational processes." If both members of the contrasted pair can be viewed in this way, they can be said to exist in the same domain of discourse (Ryle, 1949); i.e., the difference between them which causes only one of them to be conscious is not due to one being a representational event and the other not. Such an argument sharpens up the question, what are the distinctive properties of conscious processes? We will describe a set of such properties at the end of this paper <2>.

2.11 The classification of the contrasts shown in Tables 1-5 is not cast in concrete, of course: it is convenient but not exhaustive, and the categories overlap somewhat. The contrasting pairs may be arranged in different ways for different purposes. Nevertheless, these tables indicate that we already know a great deal about consciousness "as such." Further, what we know places boundaries on any complete theory of conscious events, and the variety of these facts is such as to suggest that what we know is not at all arbitrary -- the evidence is not somehow manufactured by our cultural expectations (viz., Nisbett and Wilson, 1977; Wilkes, 1988).

Table 1: Stimulus analysis.

Conscious Cases

1. Percepts: Conscious input representations

Comparable Unconscious Cases

1. Stimuli lacking in intensity or duration, and masked stimuli.
2. Pre-perceptual stimulus processes.
3. Post-perceptual (habituated) processing.
4. Unaccessed interpretations of ambiguous stimuli.
5. Contextual constraints on the interpretation of conscious percepts.
6. Expectations of specific stimuli

---

3. Contrastive Set 1: Stimulus analysis<3>

3.1 Conscious aspects of input processes

There is little disagreement that "perception" is conscious stimulus representation. Even a radical Behaviorist like B.F. Skinner suggests that consciousness is associated with stimulus control (1974). Since the advent of cognitive psychology, perception has come to be widely viewed as the process of representing the stimulus world, which, in recent years and somewhat reluctantly, is coming to be acknowledged to be conscious! (Neisser, 1967; Rock, 1983). But the converse does not necessarily hold -- not all kinds of stimulus-representation are conscious (e.g. Libet, 1977; Marcel, 1983ab; Baars, 1988).

3.2 Unconscious stimulus processing

There are at least five categories of unconscious phenomena that involve stimulus-representation. We will take them in turn, and consider arguments for their representational nature.

3.2.1 Below-threshold or masked stimulation

One obvious case where we experience a stimulus "fading out of consciousness" involves a decrease in the intensity or duration of a stimulus, or masking of one by another. There is a vast literature on subliminal effects produced in this way, which was revived in 1983 by Marcel's classic experiment showing that a masked unconscious word, which could not be reported, would still "prime" the processing of a related word. This persuaded many psychologists that unconscious words were still semantically represented in some sense. (Marcel, 1983; Dixon, 1981; Erdelyi, 1974; Shevrin & Dickman, 1980).

3.2.2 Pre-perceptual processes

There are several sources of evidence for the view that pre-perceptual processes are representational:

a) Studies of cortical evoked potentials during brain surgery suggest that activity in the primary sensory projection area occurs several hundred milliseconds before the stimulation becomes conscious (Libet, Alberts, Wright, & Feinstein, 1967; Libet, 1977).

b) Local perceptual information, such as may be acquired in a single foveal fixation, can be highly ambiguous (Gregory, 1966; Rock, 1983). Likewise, most words in the language have multiple meanings that do not become conscious in the course of normal language comprehension, so that any ambiguities must be resolved prior to the moment of conscious comprehension. Plentiful evidence suggests that local ambiguities must be resolved with reference to a larger framework representing the visual scene or the linguistic message as a whole. This, too, suggests the need for unconscious pre-perceptual hypothesis testing before one can arrive at a coherent, conscious representation of the world (Neisser, 1967; Baars, 1983; see 3.24).

c) When a stimulus is degraded so that automatic pre-perceptual processing is blocked subjects often begin to perform conscious hypothesis-testing. A good example of this occurs in reading upside-down, where the letter features and possible meanings of words

begin to be tested quite consciously. This appears to be a conscious analogue of a process that normally takes place quickly, automatically, and unconsciously.

3.2.3 Post-perceptual representations

a) Is an habituated stimulus still represented? The feeling of the chair once we have been sitting for a while, the ambient light and noise-level, one's orientation to gravity, and indeed all the multifarious sources of consistent and predictable stimulation in the environment tend to be unconscious. Nevertheless, many students of the subject believe, at least since the work of E.N. Sokolov (1963), that the nervous system continues to represent habituated stimulus events even after they have become unconscious. Sokolov's well-known arguments are based upon the occurrence of an Orienting Response (OR) whenever people or animals are confronted with a novel stimulus. (The OR consists of a large set of central and peripheral physiological events, from receptor orientation, to blocking of alpha waves, elicitation of the P300 wave in the evoked potential, and changes in a wide variety of autonomic responses such as increases in sweating, heart rate, capillary expansion or contraction, etc. In humans, we can be pretty sure that an OR-triggering stimulus is conscious.)

b) Suppose we present people with a train of white noise bursts of a certain duration, spectral distribution, onset and offset slopes, location in space, inter-stimulus interval, and so on. If the stimulus is not painfully loud, people will lose awareness of it rather quickly, but they will tend to become conscious of the noise again as soon as any stimulus-parameter changes: The noise can become louder or softer, the time between the noise-bursts can change, the intensity envelope or the frequency distribution can change -- and any of these changes will trigger a new OR. To explain this, Sokolov argues, we can only assume that there is a "model of the stimulus" against which the unconscious stimulus is matched; as long as the match fits reasonably well, one does not become conscious of the noise; only when there is a mismatch in any parameter of the stimulus do humans and animals produce another OR. This suggests that all constant or predictable sources of stimulation continue to be represented in the nervous system, even though they are unconscious.

c) While Sokolov's arguments have been widely accepted, some critics object that specific neural networks might produce the same results, thereby avoiding the notion that the nervous system "represents" the stimulus1>. From a cognitive perspective, this implies a misunderstanding of the argument. The claim that there continues to be a mental representation of all stimulus parameters after habituation does not imply that there is no specific neural substrate able to trigger another OR; rather, these are two different levels of theory, which are sometimes useful to keep separate. One does not need to believe in "strong functionalism" (e.g. Dennett, 1991) to take the position that we can usefully specify a psychological level of analysis even in ignorance of the neural substrate, though no doubt such an understanding of the neural substrate would deepen our knowledge in important ways. In this "weak functional" sense, *everybody*, including neurobiologists interested in mental processes, begins by postulating some reasonable

level of psychological analysis that can be rigorously defined even without a complete neural reduction (e.g. Crick, 1994; Edelman, 1989).

d) In sum, it seems safe to conclude that post-perceptual (habituated) stimulus events are representational, though unconscious. This contrast provides us one more empirical boundary that any adequate theory of consciousness must explain.

3.2.4 Unaccessed interpretations of ambiguous stimuli

a) What happens to the "unaccessed" interpretation of a Necker cube, or any other ambiguous word or figure, when it is not conscious? Does it disappear? Or is it still represented in some sense in the nervous system? There are in fact numerous ways to prime the unconscious interpretation of ambiguous stimulus is processed. In a classic two-channel selective listening experiment, MacKay (1973) showed that an ambiguous word like *bank* in the sentence *They walked to the bank* could be influenced by unconscious words bearing on one meaning or another of the ambiguity. Subjects 'shadowed' one ear with the conscious sentence, while words like *river* or *money* were presented to the unconscious ear, simultaneous with *bank*. Given the unconscious word *money* the probability of *bank* being interpreted as a financial institution increased significantly, and vice versa.

b) In related research, MacKay (1966) showed that preconscious processing of ambiguous words and phrases slows down when two alternative interpretations are more balanced in likelihood, as the ambiguity becomes harder to resolve. This kind of result also indicates that the brain is doing something different in response to unconscious ambiguity. Comparable results were obtained by Baars, Cohen, Bower & Berry (1992), who demonstrated that subjects who were hypnotically primed for to feel *anger* were more likely after amnesia suggestions were given, to pick words with an unconscious angry meaning. Given sentences like *At the end of the day, I still had customers to* [blank], the rather violent *finish off* was chosen more often, compared to more peaceful choices like *help*, *attend to*, or *handle.*

c) Continued brain activity relative to the unconscious interpretation of an ambiguous stimulus. In the case of binocular rivalry, there is electrophysiological evidence that the unaccessed interpretation continues to be processed (Logothetis & Schall, 1989; Shevrin \& Dickman, 1980).

d) Swinney (1979) has shown some distinctive aspects of experiment in this literature has subjects listen to a sentence fragment ending in an ambiguous word, such as *They all rose*. *Rose* can be either a verb or a noun, but in this sentence frame it must be a verb. How long will it take for this syntactic fact to influence sentence processing? To test this, a verb-related word like *stood* or a noun-related one like *flower* is presented immediately afterward. The task is to decide whether this word is a real word or not. (Non-words are presented on catch trials.) If the knowledge about the syntactic category of *rose* is available immediately, it should facilitate response time to *stood*; otherwise, there should be no difference in lexical decision time. A number of investigators have found that for

several hundred milliseconds, there is no context effect at all, though the standard priming effect occurs after that. The dead period suggests that there may be an autonomous and context-free processing module for a few hundred milliseconds after *rose* is presented. This surprising result takes for granted, of course, that there is rapid, sophisticated unconscious processing going on of both the accessed and the non-accessed meaning.

3.2.5 Contextual constraints on perception

Perceptual experiences are constrained by numerous factors which are not themselves conscious. Perhaps the most famous demonstrations of such unconscious constraints were devised by Adelbert Ames (1953), who noted, for example, that the rectangular walls, floor and ceiling of a normal "carpentered" room actually project trapezoids, not rectangles, onto the retina. Any single retinal projection can be interpreted as the result of an infinite set of trapezoids placed at different angles to the eye. But in Western culture we are exposed almost exclusively to rectangular walls, floors, and ceilings, and we interpret any consistent set of joined trapezoids to be box-shaped with rectangular sides. Hence the "Ames distorted room," which actually consists of joined trapezoidal surfaces, but is perceived as an ordinary rectangular room. And because we assess height in a carpentered environment by implicit comparison to the presumably constant height of the walls, people in an Ames room will appear to grow and shrink dramatically as they walk from one end of the trapezoidal wall to the other. In this way is our conscious experience of is size dramatically shaped by unconscious "presuppositions" about the space in which we live. Numerous other examples can be cited (Rock, 1983). Contextual constraints on perception and comprehension are the rule, not the exception (Baars, 1988, Ch. 4 and 5).

3.2.6 Expectations of specific stimuli

a) Expectations about the stimulus world are clearly representations of some sort. For instance, we can immediately detect a violation of an expectation, in any dimension of an expected stimulus or action. Yet we are not conscious of our expectations regarding the next word in this sentence, though it those expectations clearly exist: try substituting a non-word in a "glorb," for example. Thus, we can apply the Sokolov argument from section 3.23 here as well -- if the nervous system can detect a change in any dimension of some event, it has a representation of all the mismatchable parameters <4>.

Unlike percepts and images, expectations are not object-like representations -- they do not have figure-ground properties, qualitative perceptual dimensions, or a discrete moment of onset and offset in the flow of events.

b) Stimulus expectations behave in other ways like percepts and images: they are representational, they represent the environment, they sometimes involve ambiguities, they are internally consistent, yet they are not objects of conscious experience. This point is rarely made, but it seems indisputable, and theoretically significant.

3.3 Some patterns emerging from the contrastive analysis

We can now briefly look at some generalizations across the stimulus-related contrasts described above.

a) First, there is a great deal of evidence to suggest that subthreshold stimuli continue to be processed. What is the sub- and suprathreshold stimulus processing? One can argue for critical differences in terms of stimulus intensity, degree of activation of corresponding mental representations, lack of time to make unconscious events available to consciousness, and so on (see Libet, 1977; also Shevrin & Dickman, 1980). One popular move currently is to propose an activation construct, which is itself influenced by stimulus duration, intensity, and the availability of top-down influences. (MacKay, 1992). Stimuli fail to reach consciousness when they have insufficient activation (Herbart, 1824/1961).

c) If that is so, then the next contrast in Table 1 suggests that there can be *too much* activation! This is the fact of *habituation* of awareness, which occurs in all sense modalities after a stimulus has been presented to the point of predictability. Baars (1988) suggests that the critical variable may be informativeness rather than activation. A novel stimulus is high on informativeness, but a repeated one, if it has been completely represented, may drop out of awareness simply because it is redundant.

3.4 To sum up what we have found so far: It seems as if, to be conscious, an event must be presented long enough, intensely enough, identified by top-down processing enough, or, in summary, have accumulated enough activation (created by these three factors) -- but not to the point of complete predictability, or habituation will occur. The conscious representation always involves only one interpretation at any moment, even if the stimulus is really ambiguous.

Table 2. Imagistic contrasts.

Conscious Cases

1a. Retrieved images.
2a. Currently rehearsed item in Short-Term Memory.
3a. Sensory memory (iconic and echoic) and after-images.

Comparable Unconscious Cases

1b. Unretrieved images.
2b. Currently unrehearsed items in Short-Term Memory.
3b. Long-Term Memory representations needed to process sensory memory.

---

4. Images and inner speech

Consider now another example of contrastive analysis: the comparison of conscious images with comparable unconscious representations <5>. "Images" are broadly defined to include quasi-perceptual events that occur in the absence of external stimulation in any sensory modality, including inner speech and emotional feelings.

4.1 The conscious side of Table 2

a) We are conscious of more than external events. We can re-experience today's breakfast, salient events in the past, and our own inner speech. People sometimes experience hallucinations, and we all have dreams. Over the past few decades a large and reliable research literature has emerged, especially in the area of visual imagery, so that now a great deal is known about this phenomenon (Paivio, 1971; Cooper & Shepard, 1973; Kosslyn & Schwartz, 1981). Imagery in other sense modalities, inner speech, and feelings associated with emotion have seen much less research, but it is hard to see any principled reason why one could not investigate these domains with the same kind of reliability. Images are conscious representations, experienced in the absence of the imagined object. In this paper we will use the word "imagery" very broadly, to mean all of those quasi-perceptual conscious experiences we can have in the absence of an external stimulus. Visual images resemble visual percepts in a number of respects (Finke, 1980), and in fact many of the same means of assessment can be used for both perception and imagery -- notably, we can use verbal report, which is in practice our primary means for deciding whether people are conscious of something. We can hope that the kind of reliable evidence that has been found in the past few decades regarding visual images can be extended to the study of imagery in the broad sense we are using here, including inner speech and emotional feelings.

b) Clearly any adequate theory of conscious experience should be able to explain why images are conscious, while the following "stimulus-representations in the absence of the stimulus" are not.

4.2 The Unconscious side of Table 2: four cases

4.2.1 Images prior to retrieval

Where is our image of yesterday's breakfast before we bring it to mind? If it is accurate, it must in some sense be represented in long term memory. And after such images are lost from consciousness, some representation must continue to exist, since we can retrieve it again, the first retrieval primes the second one, and so on.

4.2.2 Currently unrehearsed items in Short Term Memory

Comparatively little work has been done on "inner speech" but the vast literature on Short Term Memory bears on this topic rather closely (e.g. Baddeley, 1976; Norman, 1976). In a typical Short-Term Memory experiment, subjects are given a string of unrelated words,

letters, or numbers, and requested to retrieve them shortly afterwards. A great deal is known of the resulting memory patterns, but relatively little attention has been given to the fact that during the retention interval, only the currently rehearsed item is conscious at any single moment. Thus Short Term Memory is closely associated with conscious experience, though not identical to it. (Baddeley, 1993)

4.2.3 The contents of sensory memory vs. the long-term memory information that is needed to process those contents

a) Immediately after a bright visual stimulus we tend to see after-images. Such phenomena have been studied for more than a century, and a great deal is known about them. They are typically interpreted as a by-product of neuronal activity in the visual system. In this interpretation, it is not clear whether they always function as representations. However, sensory memory, a closely related phenomenon, is generally interpreted as a momentary representation which is operated upon and further processed by other parts of the cognitive apparatus. In a classic experiment, Sperling (1960) showed that immediately after exposure to a matrix of visual symbols we can access much more information than we can report. Neisser (1967) and Atkinson and Shiffrin (1968) suggest that information from the sensory information stores is transferred to short-term memory, and may in turn be coded into long-term memory.

b) Yet the very definition of an "item" in sensory memory is profoundly affected by long-term mechanisms. To give an example: Consider the number sequence 4, 9, 9, 1, 0, 0, 9, 1, 6, 7, 7, 1, 6, 6, 0, 1. Being more than 7 plus or minus 2, it exceeds short-term memory. But reading the sequence backwards helps dramatically, because it contains only four well-known "chunks": the years 1066, 1776, 1900, and 1994. Suddenly the number series fits well within short term memory limits. Our "long-term knowledge" about important years certainly affects the immediate information store. In this example the famous years can be brought to consciousness, but there are many cases where unconscious knowledge shapes the identity of items in sensory and short term memory. Perceptual learning is another excellent example of unconscious factors shaping conscious events in sensory memory.

Table 3. Attentional contrasts.

Conscious Cases

1a. Attended streams of stimulation.
2a. Previously unattended events interrupting the attended stream.
3a. Conscious events involved in directing attention voluntarily.
4a. Attention drawn to new or changed stimuli.

Comparable Unconscious Cases

1b. Wholly unattended streams of events
2b. Unattended events affecting the interpretation of simultaneously attended events.
3b. Events controlling involuntary attention.
4b. Predictable repetitions of stimuli (see Table 1).

---

5. Contrasts involving attention

This class of events also overlaps with previous categories, but it emphasizes the selective and directive aspects. That is, there is always perceptual information which might be quite conscious, but which is excluded from consciousness because of a competing stream of input (Broadbent, 1958; Norman, 1976; Neisser, 1976). Further, we make the traditional distinction between voluntary attention, which is itself preceded by a conscious decision to pay attention to something, and involuntary attention, in which an unexpected stimulus disrupts the attended stream (James, 1890/1984).

5.1 Attended versus unattended channels

There is obviously a difference in consciousness of an attended and an unattended stream of information. However, aspects of the unattended stream frequently can become conscious. In the standard "shadowing" paradigm where subjects repeat a continuous stream of speech in one ear while another one is presented to the other ear, subjects can typically identify voice quality in the "unattended" channel, though single words have been repeated as much as 35 times without subjects being able to report them (Norman, 1976). It is possible that under conditions of extreme workload in the attended channel, even the unattended voice quality would fail to filter through. Because of these considerations it may be good to be cautious about claiming that the unattended stream is utterly unconscious, though words, sentences and meanings seem truly unconscious.

5.2 Interruption of, and influence on, the attended stream

A loud noise in the unattended channel can interrupt clearly conscious information in the attended channel. Further "significant" stimuli in the unattended channel (such as one's name) can disrupt the conscious stream, even when they are not particularly loud. This can be contrasted with events in the unattended channel that change the interpretation of the conscious stream of information, but without disrupting it overtly. Thus MacKay (1973) reports that the interpretation of an ambiguous word in the conscious stream will be changed when a disambiguating word is presented in the unconscious channel. For example, the word "bank" in the sentence. "They were walking near the bank" can be interpreted either as the edge of a river or as a financial institution. If simultaneously with "bank" one presents the word "money" vs. "water" in the unattended channel, the likelihood of one interpretation over the other is changed. Thus unconscious events can influence the interpretation of simultaneous conscious events.

5.3 Voluntary vs. involuntary attention

We can ask someone to pay attention to something in the unattended channel voluntarily. In this case, conscious information (our request) precedes the shift in attention. Alternatively we can make someone pay attention to something in the unattended channel by presenting a loud noise, the name of the subject, and perhaps a variety of other "significant" stimuli which will disrupt the conscious stream of information without voluntary involvement by the subject. In that sense, events preceding voluntary attention are conscious, while those preceding involuntary attention are not.

5.4 Dishabituation of orienting

On the basis of research with the Orienting Response (OR) (Sokolov, 1963) we know that a change in any parameter of any habituated stimulus may elicit a new OR. Since the OR is clearly associated with consciousness (at least in humans), we can claim that, while predictable repetitions in stimulation remain unconscious, *changes* in this predictable patterns tend to become conscious.

Table 4. Spontaneous problem solving.

Conscious Cases

1a. Stage of problem definition (first stage)
1c. Aha! experience (final, or third stage)

Comparable Unconscious Cases

lb. Incubation (intermediate stage)

---

6. Thinking: Spontaneous problem-solving

6.1 Most thinking is inexplicit. Entirely conscious problem-solving, such as working out an arithmetic problem on paper, is extraordinarily rare. Rather, we tend to solve problems "spontaneously": to be conscious of the stage of problem-assignment, not conscious of some intermediate stage, and conscious again of the solution of the problem.

6.2 These are the famous phenomena of problem incubation and the Aha! experience (Kohler, 1929). In these cases one is conscious of the stage of problem-definition, but not of the incubation stage, in which the problem is presumably moving toward solution. Finally, the problem "comes to mind" again, and the solution is clear. Miller (1962) and others claim that we are typically conscious of the results of mental processes, but not of

the mental processes themselves. But very significantly, we are also conscious of the stage of problem definition. Further, in a reasonably complex problem we are usually conscious of intermediate steps on the road to a solution.

6.3 This pattern of conscious problem-definition, unconscious intermediate processes, and conscious solution of the goal or subgoal extends far beyond the classic cases of "creative problem-solving" in mathematics or other formal domains. It is typical of numerous commonplace mental processes.

6.4 Word-retrieval and question-answering

We may be conscious of an incomplete sentence, unconscious of the retrieval process and conscious again of the arrival of the proper word. Similarly, if someone asks a question, we are conscious of the question, usually not of the process of searching for an answer, and conscious again of the arrival of the answer. While the time intervals involved in these commonplace processes are much shorter than in the case of mathematical problem-solving, the overall pattern seems the same.

6.5 Recall from long-term memory

The same may be said of other recall processes. We can retrieve the image of the American flag, but the process thereby we do so is utterly opaque. Free association and numerous other memory tasks have the same character.

6.6 Action planning and control

We may have some conscious planning process about the next sentence we intend to say (though not all of an intention is conscious, as James pointed out above), and we have no access to the process whereby our conscious plans are converted into detailed movements; however, we can typically monitor conscious perceptual feedback from the results of an action. (Baars, 1988, 1994)

6.7 Perceptual reorganization

We can see two interpretations of a Necker cube, but we have little conscious insight into the process that brings us from one to the other. We may be solving a visual puzzle or trying to understand a sentence spoken in a very heavy dialect; in either case, we are conscious of some early information, often appearing to be very complex and difficult to organize, but this early organization is succeeded by a second, simpler conscious experience without any awareness of the details of intermediate processes.

6.8 Thus the conscious-unconscious-conscious pattern of problem solving processes is very general indeed. It can be found in explicit, deliberate problem solving in mathematics; in minor everyday problem-solving, like question-answering; in memory recall; in action planning and execution; and in perceptual organization. Especially in the last case it is clear that the problem-solving process need to be intentional in the usual

sense. All we need to do is be aware of the Necker cube, and suddenly we may see it become reorganized.

6.9 One intriguing possibility is that James' "stream of consciousness", which appears as a series of "flights and perches" of the mind on different topics, actually consists of an interwoven series of such conscious-unconscious-conscious triads. It may be that we are continuously engaged in a number of overlapping problem-solving processes, in which unconscious mechanisms attempt to resolve issues posed consciously, returning their answers to consciousness as well. These answers may, in turn, provide the conscious input for another unconscious problem-solving process.

Table 5. Learning and practice contrasts

Conscious Cases

1a. unskilled activities
1a'. de-automatized activities
2a. instances-to-learn generated by unconscious rule structure

Comparable Unconscious Cases

1b. "automatized" activities
2b. learning of complex rule structure

---

7. Consciousness and some learning phenomena

7.1 Developing automaticity with practice in predictable tasks

It is commonly observed that when we begin learning a difficult skill, we may be conscious of many details; after skill acquisition we are conscious of much less; and if the skill is disrupted in some way, we become conscious of some missing ingredient. Indeed, Langer and Imber (1979) have shown that subjects learning a simple coding task cannot retrieve the number of steps in the task once it has become automatic, although this is quite easy before automatization of the task. This pattern suggests that conscious involvement may help to integrate new information, but that it is not required for the smooth, routine execution of complex tasks. (There is obviously a close relationship between this pattern and the habituation phenomena discussed in Table 1). When automatic execution of a skilled task is disrupted, as in reading upside-down, the opposite occurs: we tend to become more conscious of the details of the task (Baars, 1988).

7.2 Loss of conscious access to visual representations that nonetheless continue to function.

A particularly interesting case of this pattern exists in skilled use of imagery. Cooper and Shepard (1973) already noted that subjects who are skilled in their classic mental rotation task often report losing awareness of their own processes. Nevertheless, the unconscious 'image' continues to rotate at the same rate, as shown by reaction time, matching to sample, and the like. Similarly, Pani (1982) has shown that mental images required to solve a problem become less conscious available with practice, but can re-emerge when the subject encounters unexpected difficulty.

7.3 Implicit learning in miniature grammars.

Subjects who are given a set of stimuli generated by a simple "grammar", unconsciously induce the underlying grammar as shown by successful recognition of novel cases generated by the same rule-systems (Franks & Bransford, 1971; Reber, 1993). Because humans routinely learn numerous rule-systems without being able ever to state the rules, this finding has implications for a great deal of actual learning. It is one among many indications that consciousness "focuses" many unconscious capabilities upon problems to be solved in the world (Baars, 1983, 1988).

8. Some common properties of these contrastive analyses

8.1 A main theme of this paper is that "contrastive analysis" of conscious vs. unconscious phenomena yields empirical constraints that any adequate theory of consciousness must explain. Theory is not our current objective. Nevertheless. we can suggest commonalities in the facts arrayed above.

8.2 The first generalization that seems to "leap out" is that conscious experiences are most clearly articulated in the case of perceptual or quasi-perceptual phenomena: perception in all modalities, imagery in all modalities, inner speech, feelings associated with emotions -- all are conscious in this narrow sense. When people speak of intentions, expectations, beliefs, abstract ideas, and meanings they seem to mean something different from conscious percepts and images. These two senses are not unrelated of course: they both involve readily reportable representations that take up limited capacity. Further, one can make a strong theoretical case for a deeper similarity (cf. Baars, 1988). Nevertheless, it is important to point out the distinction between conscious experience and "conscious" knowledge (see also Mangan, 1993).

8.3 The fact that clearly articulated conscious experiences are perceptual or imagistic object-like events does not mean that consciousness is just stimulus representation. Numerous kinds of stimulus representation have been found that are not conscious.

8.4 Another pattern that emerges from these facts is what the 19th century referred to as "the unity of consciousness": that is, conscious contents are always internally consistent (Baars, 1983; Mandler, 1984). Ambiguous stimuli always have only one conscious interpretation at a time, though one can make a strong case that the unaccessed interpretation is represented unconsciously.

8.5 There seems to be a clear non-monotonic time factor in whatever will be conscious. Pre-perceptual processes represent aspects of a stimulus but are not conscious. Percepts represent stimuli and are conscious, at least in significant degree. Habituated percepts again are lost from consciousness.

8.6"8.6 In sum: To be conscious, a mental representation must be have adequate duration, but not so long that it becomes completely predictable; it is most clearly conscious if it is perceptual or imagistic (including inner speech); it is internally consistent, even if the stimulus is really ambiguous; object-like rather than contextual. These points apply to stimulus analysis, imagery, attentional phenomena, and "spontaneous problem solving." Images are relatively evanescent, concrete rather than abstract, associated with limited capacity processes (i.e., they are vulnerable to competition from other potentially conscious events). They are often dependent upon voluntary control to retrieve and maintain them, and when competition occurs between different potentially conscious images and percepts, the more personally "significant" tend to win out. These are all points we would expect from processes associated with consciousness (e.g. Mandler, 1984; Baars, 1988): contrastive analyses permit us to be precise about them. Productive uses of contrastive analysis have focused especially on the functional capabilities of conscious vs. unconscious processes, the necessary neural conditions of conscious experience, the nature of voluntary control, and the concept of self (Baars, 1988; 1994).

8.7 No theory of conscious processes can be complete without working out the relationship between consciousness in the core sense discussed in this paper, and consciousness in several broader senses: at the very least, conscious experience should be explicitly connected to "conscious action" in the sense of "intentional, premeditated action", "conscious knowledge and belief", and "consciousness" in the sense of "being awake". This enterprise is beyond the aim of this paper, but has been attempted in Baars (1988).

9. Summary and Conclusions

What can we conclude overall?

9.1 That criterial facts about consciousness can be specified without either evading the issue, or demanding that we meet the impossibly demanding Bat Criterion (Nagel, 1974).

9.2 That a great deal of reliable empirical information on conscious experience is already at hand, based on decades of research under other names.

9.3 That the pattern of phenomena is sufficiently complex and non-obvious to make it extremely unlikely that it is merely dictated by social convention (see Nisbett and Wilson, 1977; Wilkes, 1988).

9.4 Finally, that even the atheoretical approach we have taken in this paper yields some generalizations about conscious experience, thereby placing strong constraints on any reasonable theory of consciousness.

Notes

An earlier version of this paper appeared in Baars (1986a). The current version has undergone significant revisions and additions to reflect recent findings and theoretical developments.

<1> Some connectionists argue that the notion of mental representation should be replaced by the a state of an "adaptive system". However, any representation already exists in an adaptive medium, such as neural tissue or computer chips. A usable adaptive medium maintains a flexible but robust isomorphism between a representation and the represented event. Potential media like stellar constellations and water cannot maintain such a robust and flexible isomorphism, and are therefore poor representational media for psychobiological purposes. Thus the notion of representation *already requires* adaptiveness. Adaptation and representation are two parts of the same elephant.

<2> Contrastive analysis has also been applied fruitfully to the comparison of conscious and unconscious functional capabilities, to the issue of "self" and voluntary control.

<3> The same analysis can be applied to abstract concepts. Conscious concepts are not the same as images, but often have typical images associated with them. This issue is discussed in Baars (1988) in considerable detail.

<4> Indeed, in a curious sense, the Sokolov model of the stimulus IS an expectation, though, unlike other expectations, it represents the unconscious, habituated stimulus.

<5> Formally, "imagistic" events can be defined as "internally-generated, quasi-perceptual events."

References

Ames, A.. Jr. (1953) Reconsideration of the origin and nature of perception. In S. Ratner (Ed.) *Vision and action.* New Brunswick: Rutgers University Press.

Atkinson, R.C. & Shiffrin, R.M. (1968) Human memory: A proposed system and its control processes. In K.W. Spence and J.T. Spence (Eds.) *Advances in the psychology of learning and motivation: Research and theory,* Vol. 2. NY: Academic Press.

Baars, B. J. (1983). Conscious contents provide the nervous system with coherent, global information. In R. Davidson, G. Schwartz, & D. Shapiro (Eds.) *Consciousness and self-regulation,* Vol. 3, 45-76. New York: Plenum Press.

Baars, B.J. (1986a) What is a theory of consciousness a theory of? The search for criterial constraints on theory. *Imagination, Cognition and Personality, 6(10),* 3-23.

Baars, B.J. (1986b) *The cognitive revolution in psychology.* NY: Guilford.

Baars, B.J. (1988) *A cognitive theory of consciousness.* London: Cambridge University Press.

Baars,. B.J. (1993) Why volition is a foundation issue for psychology. *Consciousness & Cognition, 2* 281-309.

Baars, B.J., Cohen, J., Bower, G.H. & Berry, J. W. (1992) Some caveats on testing the Freudian slip hypothesis: Problems in systematic replication. In B. J. Baars (Ed.) *Experimental slips and human error: exploring the architecture of volition,* 289-316. NY: Plenum Press.

Baddeley, A.D. (1976) *The psychology of memory.* NY: Basic Books.

Baddeley, A. (1993). Working memory and conscious awareness. In A..F. Collins & M.A. Conway (Eds.), *Theories of memory.* Hove, UK: Erlbaum.

Broadbent, D.E. (1958) *Perception and communication.* NY: Pergamon.

Cooper, L.A. and Shepard. R.N. (1973) Chronometric studies of the rotation of mental images. In W.G. Chase (Ed.) *Visual information processing.* New York: Academic Press.

Crick, F.C.H. (1994) *The astonishing hypothesis: The scientific search for the soul.* NY: Scribner.

Crick, F. C. H. & Koch, C. (1992). The problem of consciousness. *Scientific American, 267(3),* 152-160.

Dennett, D. (1991). *Consciousness explained.* Boston: Little, Brown.

Dixon, N.F. (1981) *Preconscious processes.* NY: Wiley.

Edelman, G. (1989). *The remembered present: A biological theory of consciousness.* NY: Basic Books.

Erdelyi, M. (1974) A new look at the New Look: Perceptual defense and vigilance. *Psychological Review, 81,* 1-25.

Ericsson, K.A. & Simon, H.A. (1984/ 1993) *Protocol analysis: Verbal reports as data.* Cambridge, MA: MIT Press.

Finke, R. (1980) Levels of equivalence in imagery and perception. *Psychological Review, 87(2),* 113-132.

Franks, J.J. & Bransford, J.D. (1971) Abstraction of visual images. *Journal of Experimental Psychology, 90(1),* 65-74.

Gibson, J.J. (1966) *The senses considered as perceptual systems.* Boston: Houghton.

Gregory, R.L. (1966) *Eye and brain: The psychology of seeing.* NY: Harper & Row.

Herbart, J.F. (1824/1961) Psychology as a science, newly founded on experience, metaphysics, and mathematics. Reprinted in T. Shipley (Ed.) *Classics in psychology.* NY: Philosophical Library. (pp. 22-50)

Hilgard, E. R. (1992) Divided consciousness and dissociation. *Consciousness and Cognition, 1(1),* 16-32.

Hochberg, J. (1964) *Perception.* Englewood Cliffs, N.J.: Prentice-Hall.

Hull, C.L. (1937) Mind, mechanism, and adaptive behavior. *Psychological Review, 44,* 1-32.

James, W. (1890/1984) *The principles of psychology.* NY :Holt. Reprinted, Harvard University Press.

Kohler, W. (1929) Gestalt psychology. NY: Liveright.

Kosslyn, S.M. and Schwartz, S.P. (1981) Empirical constraints on theories of visual mental imagery. In J.Long and A.D. Baddeley (Eds.) *Attention and performance, Vol. IX,* 214-260. Hillsdale, N.J.: Erlbaum.

Langer, E.J. and Imber, L.G. (1979) When practice makes imperfect: Debilitating effects of overlearning. *Journal of Personality and Social Psychology, 37(11),* 2014-2024.

Libet, B. (1977) Neuronal vs. subjective timing for a conscious sensory experience. In P.A. Buser & A. Rougeul-Buser (Eds.), (1978) *Cerebral correlates of conscious experiences: INSERM Symposium No. 6,* 69-82. Amsterdam: Elsevier.

Libet, B., Alberts, W.W., Wright, E.W. and Feinstein, B. (1967). Responses of human somatosensory cortext to stimuli below threshold for conscious sensation. *Science, 158,* 1597-1600.

Logothetis, N. D. and Schall, J. D. (1989). Neuronal correlates of subjective visual perception. *Science, 245,* 761-763.

MacKay, D.G. (1966) To end ambiguous sentences. *Perception and Psychophysics, 1,* 426-436.

MacKay, D.G. (1973) Aspects of a theory of comprehension, memory and attention. *Quarterly Journal of Experimental Psychology: General, 110,* 341-462.

MacKay, D.G. (1992) Errors, ambiguity, and awareness in language perception and production. In B.J. Baars (Ed.), *Experimental slips and human error: Exploring the architecture of volition,* 39-67. NY: Plenum Press.

Mandler, G.A. (1975). Consciousness: respectable, useful, and probably necessary. In R. Solso (Ed.) *Information processing and cognition: The Loyola Symposium.* Hillsdale, NJ: Erlbaum.

Mandler, G. (1984). *Mind and body: psychology of emotion and stress.* New York: Norton.

Mangan, B. (1993) Taking phenomenology seriously: The "fringe" and its implications for cognitive research. *Consciousness and Cognition, 2(2),* 89-108.

Marcel, A.J. (1983a) Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15,* 197-237.

Marcel, A.J. (1983b) Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology, 15,* 238-300.

Miller, G.A. (1962) *Psychology: the science of mental life.* Harper & Row.

Nagel, T. (1974) What is it like to be a bat? In N. Block (Ed.) (1980) *Readings in the philosphy of science, 1,* 159-70. Cambridge, MA: Harvard University Press.

Natsoulas, T. (1982) Conscious perception and the paradox of "blindsight". In G. Underwood, *Aspects of consciousness, 3,* 79-110. N.J.: Erlbaum.

Neisser, U. (1967) *Cognitive Psychology,* NY: Appleton- Century-Crofts.

Neisser, U. (1976) *Cognition and Reality.* San Francisco: W. H. Freeman.

Newman, J. & Baars, B.J. (1993) A neural attentional model for access to consciousness: A Global Workspace perspective. *Concepts in Neuroscience, 4(2),* 255-290.

Nisbett, R.E. & Wilson, T.D. (1977) Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84,* 231-59.

Norman, D.A. (1976) *Memory and attention.* NY: Wiley.

Norman, D.A. & Shallice, T. (1986) Attention and action: Willed and automatic control of behavior. In R. Davidson, G.E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation, Vol. 4.* New York: Plenum.

Paivio, A. (1971) *Imagery and verbal processes.* NY: Holt, Rinehart.

Pani, J. R. (1982) A functionalist approach to mental imagery. Paper presented at the 23rd Annual Psychonomic Society Meeting.

Posner, M. I., & Rothbart, M. K. (1991). Attentional mechanisms and conscious experience. In A. D. Milner & M. D. Rugg (Eds.), *The neuropsychology of consciousness,* 91-112. NY: Academic Press.

Reber, A.S., Kassin, S.M., Lewis, S., & Cantor, G. (1980) On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning and Memory, 6(5),* 492-502.

Rock, I. (1983) *The logic of perception.* Cambridge, MA: Bradford/MIT Press.

Ryle, G. (1949) *The concept of the mind.* London: Hutchinson.

Shallice, T. (1972) Dual functions of consciousness. *Psychological Review, 79(5),* 383-393.

Shallice, T. (1978) The Dominant Action System: An information-processing approach to consciousness. In K.S. Pope & J.L.Singer (Eds.) *The stream of consciousness: scientific investigations into the flow of experience.* NY: Plenum.

Shevrin, H. and Dickman, S. (1980) The psychological unconscious: A necessary assumption for all psychological theory. *American Psychologist, 35(5),* 421-434.

Skinner, B.F. (1974) *About behaviorism.* NY: Knopf.

Smith, K. (1969) *Behavior and conscious experience: a conceptual analysis.* Athens Ohio: Ohio University Press.

Snodgrass, M., Shevrin, H. & Kopka, M. The mediation of intentional judgments by unconscious perception. *Consciousness & Cognition, 2(3),* 169-193.

Sokolov, E.N. (1963) *Perception and the orienting reflex.* NY: MacMillan.

Sperling, G. (1960) The information available in brief conscious presentations. *Psychological Monographs, 74,* (Whole No. 11).

Swinney, D. (1979) Lexical access during sentence comprehension. (Re) consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18,* 645-60.

Weiskrantz, L. (1988). *Blindsight. A case study and its implications.* Oxford, UK: Oxford University Press.

Wilkes, K. V. (1988) ------, yishi, duh, um, and consciousness. In A.J. Marcel and E. Bisiach (Eds.) *Consciousness in Contemporary Science.* Oxford: Clarendon Press.