« appendix »

MAKE YOUR OWN THEORY: A SUMMARY OF THE EVIDENCE

This section provides a more technical summary of the major bodies of evidence about conscious experience. By comparing contrasting pairs of phenomena that are very similar except that one is conscious while the other is not, we can hone in on just those elements that are uniquely associated with consciousness. Here are pairs of "conscious versus unconscious" facts that any complete theory must explain. Try to see if you can account for them—with a theater metaphor or in any other way!

Conscious and Unconscious Aspects of Input Processes

There is little disagreement that "perception" is conscious input representation. Even a radical behaviorist like B. F. Skinner suggested that consciousness is associated with "stimulus control." But not all kinds of stimulus representation are conscious. There are at least six categories of comparable *un*conscious stimulus-representation. We will take them in turn.

Below-Threshold or Masked Stimulation

One obvious case where we lose a stimulus from 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, revived in 1983 by Anthony Marcel's classic experiments showing that a masked unconscious word, which could not be reported, would still "prime" the processing of a semantically related word. This persuaded many psychologists that unconscious words were still semantically represented in some sense.

Preperceptual Processes

There are several sources of evidence for the view that preperceptual processes are representational. Perhaps the most persuasive involve brain electrical activity evoked by stimuli during coma or deep sleep. 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 preperceptual hypothesis testing before one can arrive at a coherent, conscious representation of the world.

When a stimulus is degraded so that automatic prepercepual 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.

Postperceptual Representations

Is a 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, 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, characteristic parts of the evoked potential, and changes in a wide variety of autonomic responses such as increases in sweating, heart rate, capillary expansion or contraction, et cetera. An animal or human being looks for the source of stimulation once a stimulus elicits an OR, so we can be pretty sure that those stimuli are conscious.

Suppose we present people with a train of white noise bursts of a certain duration, spectral distribution, onset and offset slopes, location in space, interstimulus 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—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. Habituated material includes all the things we take for granted, everything that has become predictable.

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

Unaccessed Interpretations of Ambiguous Stimuli

What happens to the "unaccessed" interpretation of an ambiguous visual figure like a Necker Cube when it is not conscious? Does it disappear? Or is it still represented in some sense in the nervous system? In a classic two-channel selective listening experiment, Donald G. MacKay showed in 1973 that an ambiguous word such as "bank" in the sentence "They walked to the bank" could be influenced by unconscious words bearing on one meaning of the ambiguity. Subjects "shad-

owed" one ear with the conscious sentence while words such as "river" or "money" were presented to the unconscious ear, simultaneous with "bank." When the subjects were given the unconscious word "money," the probability of *bank* being interpreted as a financial institution increased significantly, and vice versa.

In related research, MacKay showed in 1966 that preconscious processing of ambiguous words and phrases slows down when two alternative interpretations are balanced in likelihood, and 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 myself, Jonathan Cohen, Gordon Bower, and Jack Berry, who demonstrated that subjects who were hypnotically primed to feel anger were more likely to pick words with an unconscious angry meaning. Given incomplete sentences such as "At the end of the day, I still had customers to———," angry subjects more often unconsciously chose to fill in the blank with the rather violent "finish off", rather than to more peaceful choices like "help," "attend to," or "handle."

In the case of binocular rivalry, two streams of information are presented to a person's two eyes. Nikos Logothetis and collaborators have found clear electrophysiological evidence that the unaccessed interpretation continues to be processed

David Swinney and others have shown other distinctive aspects of ambiguity-resolution in language perception. A typical 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 such as "stood" or a noun-related one such as "flower" is presented immediately after "rose." The subject's task is to decide whether this word is a real word or not. 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 a context free automatism that takes over for as long as 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.

Contextual Constraints on Perception

Perceptual experiences are constrained by numerous factors that are not themselves conscious. Perhaps the most famous demonstrations of such unconscious constraints were devised by Adelbert Ames in the 1950s, 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 mostly 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. 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 size dramatically shaped by unconscious "presuppositions" about the space in which we live. Numerous other examples can be cited. Contextual constraints on perception and comprehension are the rule, not the exception.

Expectations of Specific Stimuli

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 the expected event. Yet we are not conscious of our expectations regarding the next word in this sentence, though those expectations clearly exist: try substituting a "glorb " in a sentence, for example. Thus, we can apply the Sokolov argument 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. Unlike percepts and images, expectations are not objectlike representations—they do not have figure-ground properties, qualitative perceptual dimensions, or a discrete moment of onset and offset in the flow of events.

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.

Images and Inner Speech

Consider now another example of contrastive analysis: the comparison of conscious images with comparable unconscious representations. "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.

The Conscious Side of Imagery

We are conscious of more than external events. We can reexperience today's breakfast, important events from the past, and hear 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. 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 book we 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, 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.

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.

The Unconscious Side of Mental Imagery

Memory Images Before Retrieval

Where is our mental 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.

Currently Unrehearsed Items in Working Memory

Comparatively little work has been done on "inner speech" but the vast literature on Working Memory bears on this topic rather closely. In a typical Working Memory experiment, subjects are given a string of unrelated words, letters, or numbers, and requested to retrieve them shortly afterward. 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, Working Memory is closely associated with conscious experience, though not identical to it.

Automatic Mental Images

One of the most intriguing ideas about mental images is that they may fade from consciousness and yet continue to function. Hard evidence to that effect has been developed by John Pani among others.

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 that might be quite conscious, but which is excluded from consciousness because of a competing stream of input. 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.

Attended versus Unattended Messages

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 thirty-five times without subjects being able to report them, according to Donald Norman.

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. (see above) Thus unconscious events can influence the interpretation of simultaneous conscious events.

Voluntary versus 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 that 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.

Dishabituation of Orienting

On the basis of research with the Orienting Response (OR), 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 pattern tend to become conscious.

Thinking: Spontaneous Problem-Solving

Most Thinking Is Inexplicit

Entirely conscious problem-solving, such as working out an arithmetic problem on paper, is quite 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.

These are the famous phenomena of problem incubation and the "Aha!" experience, frst discussed by Gestalt psychologists like Köhler in the 1920s. 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. George Miller and others have pointed out 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.

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 creative mathematical problem-solving, the overall pattern seems the same.

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 whereby we do so is utterly opaque. Free association and numerous other memory tasks have the same character.

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), 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.

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.

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, such as 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 does not 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.

One intriguing possibility in this regard is that James's "stream of consciousness," which appears as a series of "flights and perches" of the mind on different topics, could actually consist 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.

Consciousness and Some Learning Phenomena

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, Ellen Langer and Nancy Imberhave 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 the table at the end of this appendix). 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).

Loss of Conscious Access to Visual Information that Nonetheless Continues to Inform Problem-Solving

A particularly interesting case of this pattern exists in skilled use of imagery. Lynn Cooper and Roger Shepard 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, John Pani has shown that mental images required to solve a problem become less consciously available with practice but can reemerge when the subject encounters unexpected difficulty.

Implicit Learning of 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. 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.

Capability Contrasts

Whatever we do really well, we tend to do unconsciously, from speaking to seeing to playing piano. This observation has led some psychologists to wonder why consciousness is needed at all. To get at this question of the role of consciousness, we can conduct another contrastive analysis, focused on the capabilities of comparable conscious and unconscious processes. The table at the end of this appendix presents a basic set of such Capability Contrasts. Notice that purely conscious processes are

handicapped by low computational efficiency: they are rather inefficient (slow, prone to error, and vulnerable to interference), serial, and limited in capacity. Consider mentally multiplying 23 × 79. For most of us this is not very easy to do, and the more the steps are conscious, the more difficult it is. Yet mental multiplication is trivial in complexity compared to the vast amount of processing that is needed to analyze the syntax of this sentence. But syntactic analysis is of course entirely unconscious. Mental arithmetic can become more efficient with practice by letting highly predictable steps become automatic and unconscious, but that illustrates the same point, that efficiency in computational processes is achieved when some algorithm becomes unconscious.

Purely conscious mental manipulations have a high rate of errors, are slow, and interfere with each other, suggesting that performing efficient symbolic computation is not the primary function of consciousness.

But the computational drawbacks of experience are balanced by clear advantages: Consciousness has a *vast range* of possible contents, it enables novel *access* to an astonishing number of skills and knowledge sources, and it shows exquisite context-sensitivity.

As an example of its vast range, consider all the possible percepts, images, memories, concepts, intentions, fringe experiences, and the like of which we can be conscious. As one poetic student put it, we can be conscious of everything from the "rumbling of our stomach to the return of a theme in a Bach cantata."

A conscious event, like this sentence, can access new information in memory, combine knowledge from different sources in the brain in novel ways, and trigger unconscious rule systems that will pick up errors in any level of analysis—in the meaning, syntax, word level, sound, intonation, or printing of this "snetnecne."

The context sensitivity of consciousness can easily be shown to permeate whole domains, such as perception, thinking, or language. Take the predominance of lexical ambiguity, the fact that most words in natural languages have multiple meanings. The *Oxford English Dictionary*, for example, devotes seventy-five thousand words to the many different meanings of the little word "set." Thus, whatever we experience is shaped by unconscious processes, just like a theater in which we see only the actors onstage, but not all the people behind the scenes who make it all work. If the same actors were off stage (unconscious) their actions would not be contextualized by the entire supportive apparatus of the theater. The context sensitivity of conscious events extends far beyond language to perception, action control, memory, problem solving, et cetera. But there is no evidence that novel combinations of sub-

liminal words can be understood unconsciously. Likewise, when we make a navigational error simply because we are so used to turning right on the street going home, the less conscious we are, the more we are likely to make the error. The less conscious some event is, the less it is sensitive to context.

Unconscious processes have their own advantages, however. Unconscious automatisms, such as the ones that control nearly all aspects of the act of reading, show impressive speed and accuracy in routine matters, a tendency to perform parallel or concurrent processing whenever possible, and, when all unconscious resources are taken together, a vast capacity. But of course there is constant interaction between conscious and unconscious processes. In listening to a friend describe last night's party, we follow the conscious flow of sound, words, and meaning with no awareness of the complex acoustic, phonological, morphological, lexical, syntactic, semantic, intonational, and pragmatic processes happening at the same time, all of which are needed for us to become aware of the message.

Yet we can easily prove that these sophisticated unconscious processes are going on all the time. Should our friend commit any error, such as saying "entomology" instead of "etymology, we would immediately detect the error (if we were paying attention), even though it occurred at any of a dozen different levels of analysis. Further, concurrent with all this fast, complex, and unconscious linguistic activity, we also maintain balance and upright posture, represent predictable aspects of all incoming stimuli, and shape our actions in terms of the social and pragmatic demands of the situation.

In the laboratory, the limitations of purely unconscious language processing have been highlighted in selective attention studies. If we receive two dense flows of information, such as two simultaneous stories, one in each ear, or two different ball games shown on the same television set, we can follow only a single, consistent flow of the action. Under these conditions we can present information to the "unattended channel," the ear one is not listening to, for example. In general, it has been found that semantic priming from individual words in the unattended channel can influence the experience of the conscious, attended channel. Thus, the word "money" in the unattended message can bias understanding of the word "bank" toward "financial institution" instead of "shoreline of a river". However, the information in the unconscious channel does not extend to the meaning of longer passages.

Consistent patterns of evidence now begin to emerge from the contrastive analysis. We see the interplay of a serial, integrated, and very

limited stream of consciousness with an unconscious system that is distributed, composed of autonomous modules, and of enormous collective capacity.

Here is a summary table of the distinctive capabilities of both conscious and unconscious mental processes. The details are explained in this appendix, and in previous chapters. Here is your data. Can you design a theory that fits? Good luck!

Table 1
Capabilities of Conscious and Unconscious Processes

Conscious Processes	*Unconscious Processes*
1. Computationally inefficient: e.g., mental arithmetic.	1. Very efficient in routine tasks: e.g., syntax.
Many errors, relatively low speed, and mutual interference between conscious processes.	Few errors, high speed, and little mutual interference.
2. Great range of contents.	2. Each routine process has a limited range of contents.
Great ability to relate different conscious contents to each other.	Each routine process is relatively isolated and autonomous.
Great ability to relate conscious events to their unconscious contexts.	Each routine process is relatively contextfree
3. High internal consistency at any single moment, seriality over time, and limited processing capacity.	3. Routine, unconscious processes are diverse, can sometimes operate in parallel, and together have great processing capacity.
4. The clearest conscious contents are perceptual or quasi-perceptual (e.g. imagery, skill learning, problem-solving, action control, etc.).	4. Unconscious processes are involved in *all* mental tasks, from perception and imagery, inner speech, and internally into memory, knowledge representation and access, generated bodily feelings, etc.).