Self-Deception and Confabulation

William Hirstein†
Elmhurst College

Cases in which people are self-deceived seem to require that the person hold two contradictory beliefs, something which appears to be impossible or implausible. A phenomenon seen in some brain-damaged patients known as confabulation (roughly, an ongoing tendency to make false utterances without intent to deceive) can shed light on the problem of self-deception. The conflict is not actually between two beliefs, but between two representations, a 'conceptual' one and an 'analog' one. In addition, confabulation yields valuable clues about the structure of normal human knowledge-gathering processes.

1. Introduction. Compare the following two cases:

In the first case, a doctor is applying for a health insurance policy. When the agent asks him if he has any preexisting conditions, he answers "No, and I would know." The agent accepts this and issues him the policy. But the doctor has deceived the agent; he has a form of cancer which is terminal and is fully aware of that fact.

In the second case, the doctor notices certain symptoms in himself which, had he observed them in another person, would immediately cause him to conclude that that person has a certain type of cancer which is invariably terminal. But he explains away each symptom as probably due to some minor ailment, and refuses to consider the cumulative probability that the conjunction of the symptoms is due to that form of cancer. The doctor believes that he does not have cancer, but he is self-deceived.

In the first case, the doctor has deceived the insurance agent, and this means that the doctor knew something to be the case, that he did have a preexisting condition, but withheld that from the agent, who believes the opposite to be the case. Applied to the case of *self*-deception, however,

†Send requests for reprints to the author, Department of Philosophy, Elmhurst College, 190 Prospect Ave., Elmhurst, IL 60126.

Philosophy of Science, 67 (Proceedings) pp. S418–S429. 0031-8248/2000/67supp-0032$0.00

this seems to require that the doctor believe both that he does have terminal cancer and that he is of sound body. How is this possible?

The philosophical work on this problem consists primarily of working out accounts of the mind (including, sometimes, accounts of the self) or of belief which are able to resolve the problem. These accounts generally fall under one of two types: separation theories and nonseparation theories. Separation accounts of irrational behavior trace back to Plato, through Freud, to Davidson, but Pears's (1984) theory is the most recent and thoroughly developed such approach. In the case of the self-deceiving doctor, Pears would say that there is in the doctor's mind a cautionary belief, that he is irrational to believe that he does not have cancer. But the belief does not interact with the doctor's belief that he does not have cancer, because it is contained within a subsystem, "a separate center of agency within the whole person." One objection to Pears's theory is that it describes not self-deception, but the deception of one entity by another. It also allows several entities in the mind to have their own intentional states, a move which worries some: where will the regress of subagents end? Is the existence of subagents plausible at all?

I deliberately named the second type of theory negatively, because they typically are constructed in response to what are seen to be unnecessary extravagances in the separation theories. Mele's (1987) theory is perhaps currently the best known example of a nonseparation theory. Mele avoids the implausible situation of an agent believing both p and not p by having the agent believe that p, then employ different techniques to hold the contradictory belief at bay. For instance, a man has a long-standing belief that his wife is faithful. But when signs begin to accumulate that she is having an affair, he prevents them from causing him to form the belief that she is having an affair by misinterpreting them, or simply ignoring them.

From a methodological point of view, the accounts are interesting, in that they are examples of a sort of a priori psychological theorizing. They both make detailed claims about the nature of the human psyche, of belief formation processes, and so on. As it stands, this is fine, but self-deception is a phenomenon with parts lying in the purview of several different disciplines. What I want to do in this paper is provide a firmer empirical grounding to the inquiry by connecting normal self-deception with what seems to be an extreme form of self-deception which occurs in certain neurological patients.

2. Three Types of Confabulating Patient. Certain types of brain lesion can produce a curious phenomenon known as *confabulation*: when asked a question which touches on deficits caused by the injury, rather than simply acknowledging his problems, the patient will give a false or irrelevant

answer, as if he were attempting to cover up his deficit. The three syndromes which most often produce confabulation are split-brain, anosognosia, and Korsakoff's syndrome.

Split-brain syndrome is due to commissurotomy, a surgically-induced lesion in which the corpus callosum, the large bundle of fibers interconnecting the two hemispheres, is severed in order to prevent the spread of epileptic seizures from one side of the brain to the other. The first split-brain patients initially appeared normal following the surgery, but researchers began to find that by employing techniques to lateralize stimuli to just one of the two hemispheres, such patients could be made to seem like two people in one body (Gazzaniga 1995, Sperry 1985). Only the left hemisphere could give verbal responses, but it was found that the right hemisphere could understand simple linguistic input, and could respond by pointing to pictures with the left hand (each hemisphere has control over the arm on the opposite side). In the course of this testing, however, a curious phenomenon occurred: when the patient was asked about the activities of the left hand, the left hemisphere would answer as if *it* had been controlling the left hand, whereas, due to the commissurotomy, it had no idea why the left hand was doing what it was (Gazzaniga 1985, 1995; Sperry 1985; Sperry et al. 1979). For instance, in one study, a picture of a snow scene was lateralized to the right hemisphere of a split-brain patient, while a picture of a chicken claw was lateralized to his left hemisphere. Then an array of associated pictures was shown to each hemisphere, who responded correctly by pointing at a chicken with his right hand and at a snow shovel with his left hand. But when the patient was asked why he had chosen these items, his left hemisphere said, "Oh, that's simple. The chicken claw goes with the chicken, and you need a shovel to clean out the chicken shed" (Gazzaniga 1995). In another study, a picture of a naked woman was shown only to the right hemisphere, using a tachistoscope. When the patient was asked why he was laughing, the left hemisphere said, "That's a funny machine you've got there."

'Anosognosia' means unawareness of illness. It is exhibited by many types of neurological patient, but it occurs most frequently following stroke damage to the inferior parietal cortex of the right hemisphere, located just behind and above the right ear. Damage here can produce paralysis of the left arm or of the entire left side of the body. This paralysis can be accompanied by *neglect*, a condition in which the patient ignores the left side of her body and its nearby surrounding space. A patient with neglect typically will not eat food on the left side of her plate or wash the left side of her body, and will not notice people standing quietly on her left. Some patients with this left-side paralysis will also exhibit anosognosia for several days following their stroke. Approached on her right side as she lies in bed and asked whether she can use her left arm, such a patient

will answer matter of factly that she can. When the neurologist tells the patient to touch his nose with her left arm, the patient will typically try in vain to reach. But often, the patient may produce a confabulation, saying something like "I could if this old arthritis weren't bothering me," or "I'm not feeling very ambidextrous this morning" (Ramachandran 1995, Joseph 1996). When asked whether she reached successfully, the patient who tried to reach will often say that she did, and a large percentage of these patients will also confabulate that they saw themselves reaching.

Korsakoff's syndrome is a form of amnesia, most often caused by a lifetime of heavy drinking. The memory deficit affects episodic memory, a system which stores information about autobiographical episodes, but not semantic memory, our knowledge of concepts, including word meanings (Tulving 1979). The locus of lesion is not as clear in the case of Korsakoff's amnesia as in our other two cases, but the most frequent sites of damage are the mammilary bodies and the dorsomedial nuclei of the thalamus. Korsakoff's amnesia is severe enough that the patient will typically have no memory at all of the events of the preceding day. But when asked what he did yesterday, the Korsakoff's patient will often produce a detailed description of plausible (or not so plausible) sounding events, all of it either entirely made up on the spot, or traceable to some veridical but much older memory.

The patients give no sign that they are aware of what they are doing; apparently they are not lying, and genuinely believe their confabulations. They do not give any outward signs of lying, and their demeanor while confabulating has been described as "rocklike certitude." In one experiment which affirms the sincerity of confabulators, anosognosics were given the choice of performing a two-handed task (tying a shoe) for a reward of $10 or performing a one-handed task (screwing a bulb into a socket) for $5. The patients uniformly selected, then failed, at the two-handed task. In contrast, control patients who had left-side paralysis caused by right hemisphere stroke but no anosognosia systematically chose the one-handed task (Ramachandran 1995).

3. The Isolation Theory of Confabulation. In order to explain the confabulation which split-brain patients engage in, Gazzaniga (1995) hypothesized that the left hemisphere contains an *interpreter*, a module whose function is to produce a verbal explanation of the agent's activities based on the sources of data available to it, in order to respond to a question. When the interpreter is isolated from a source of data, it simply makes do with what it has, and creates an explanation from that. The existence of such a module explains how the confabulation is created and offered, but it does not explain why the patient does not recognize that the confabulation is false or irrelevant. This suggests that in addition to having certain

abilities, the isolated left hemisphere has a disability, specifically an inability to assess the plausibility of the confabulations it creates.

In order for such an isolation theory of confabulation to be of interest, we need to understand the two other syndromes which lead to confabulation as also cases of separation between the two hemispheres (or at least as cases in which the left hemisphere is isolated from sources of information). In anosognosia, the left hemisphere is unable to receive information about the left side of the body and its nearby space, because the part of the right hemisphere which normally represents those areas is destroyed.

The damage in Korsakoff's syndrome is bilateral, opening up the possibility that it may involve a left-right disconnection, just at a lower level (anatomically and functionally) than the disconnection present in the split-brain. Several recent imaging studies show high activity levels in the right prefrontal cortex during tasks involving recall of episodic memories (e.g., Cabeza et al. 1997). It may be, then, that Korsakoff's isolates the left hemisphere from the activities necessary for recall occurring in the right hemisphere, just as the lesion in neglect deprives the left hemisphere of its source of information about the body, located in the right hemisphere.

4. Representational Differences Between Left and Right. In trying to understand why confabulation seems to be associated with isolation of the left hemisphere, one place to begin is to examine differences in representational style between the left and right. The fact that linguistic abilities are located in the left hemisphere (in right-handed people and in most left-handed people) has been one of the best-confirmed findings in neurology ever since its discovery by Broca in 1863. Difficulties in perceiving and comprehending language occur with damage to the posterior parts of the left hemisphere's cortex, while difficulties with forming utterances and speaking occur with damage to prefrontal areas. Linguistic concepts would seem to be the preferred mode of representation of the left hemisphere.

What about the right hemisphere? One of the most important and well-documented findings of recent years has been the presence of two separate routes of visual processing leaving the occipital lobe, where visual information first enters the cortex. These routes have been called the 'what' stream and the 'where' stream; the 'what' stream seems to primarily serve the function of object identification, while the function of the 'where' stream seems to be to represent the agent's nearby visual space, for purposes of navigation, reaching, and so on. The 'what' route runs ventrally from the rearmost portion of the cortex into the temporal lobes. Damage here can result in inability to visually recognize familiar objects, including people. The 'where' route leaves the occipital lobe and runs in a dorsal

direction, toward the parietal lobe, significantly, to the very area that is injured in neglect with anosognosia.

The finding that these two different routes exist is best documented in the brains of monkeys. When the hypothesis is tested on humans through the use of imaging studies, however, only the right hemisphere shows signs of a clear what/where separation (Ungerleider and Mishkin 1982). These studies support a hypothesis of Bear (1979) that the human left and right hemispheres, unlike those of the monkey, have fundamentally different structures. According to Bear, the human left hemisphere's processing is biased toward a sort of super 'what' stream, dedicated to quickly attaching linguistic tags to incoming information. This involves a sort of digitizing of information which is initially in analog form: As visual information moves from the retina back to the occipital lobe, it retains its topographic form, so that the shapes on the back of the eyeball are roughly maintained on the sheets of cortex in the occipital lobe, a type of representation referred to as a retinotopic map. Ultimately this analog information must be parsed into information about separate objects, each corresponding to some concept, with its accompanying linguistic tag.

On the other side, processing in the right hemisphere seems to be biased toward the 'where' stream. Retinotopically structured information moves forward from the occipital lobe, while topographic representations of the agent's body—somatotopic maps—located near the central sulcus, which divides the posterior parts of the cortex from the anterior part, send information toward the back of the brain. These two streams of information meet in the inferior parietal lobe, where a representation (or a set of representations) is constructed of the body, situated in its nearby personal space. Hence, in the right hemisphere's 'where' stream, information is kept in an analog form, rather than being translated into a conceptual/linguistic form as in the left hemisphere.

The idea that information in the right hemisphere is kept in analog form provides another way to link clearly right hemisphere problems such as neglect with Korsakoff's syndrome. What fails in Korsakoff's is autobiographical memory, a form of memory which, introspectively at least, has the form of a sort of videotape replay: I remember the events of yesterday from my point of view, as I experienced them. This is an analog form of representation, just as the representation of the body and its space which is damaged in neglect. In general then, confabulation seems to occur when analog representation systems are disturbed or isolated from the conceptual representation system in the left hemisphere.

If we think of the mind/brain as modularized, there is no reason to think that each module has the sort of meta-knowledge about inputs which we expect the entire normal person to have. In order to know that a normal source of information to me is missing, I need to represent past instances

in which I engaged in causal contact with that source. A module which monitors its own input is a fancy, cognitively expensive module. When the left hemisphere is disconnected from the source by which it obtains information about the body, or recent autobiographical memories, it may have no way to 'realize' that is no longer receiving information from these sources. In general a representational system which trades in concepts may be entirely blind to information not in a conceptual form.

Analog representation systems, on the other hand, can still embody information about an unidentifiable something; they have representational space for something like 'grey blob in this spot in the visual field'. The brain's interfaces with the world—which all trade in analog representations (be they retinotopic, tonotopic, or somatotopic)—need this sort of approach in order to take in as broad a sample of the stimulus array as possible. Not so with the conceptual system, which is fed with highly processed, prepared data by analog systems. One hallmark of analog representations approaching the level of the conceptual system is that any gaps or holes in the representation have been *filled in*, by processes specialized in completing and filling in analog representations (Ramachandran and Gregory 1991, Ramachandran 1993, Churchland and Ramachandran 1993, Ramachandran and Hirstein 1997).

There may be another factor leading to the typical confidence which confabulators have in their creations. In addition to sending information to the left hemisphere, it may be that the right hemisphere sends a signal which tells the left hemisphere that it has the answer. Or, if the right hemisphere does not have the appropriate information, it may send a signal indicating so. Phenomenologically, this signal may take the form of a feeling of confidence in the case in which the information is there, or a feeling of anxiety, lack of confidence, or even alarm or danger in the case in which the information is not there.

There is a positive and a negative way in which disruption of such a signal might lead to confabulation. Positively, it may be that the channel through which the information itself is transmitted is destroyed, while the channel carrying the emotional 'confidence signal' is intact. It has been shown with split-brain patients, for instance, that emotional information may still be able to cross from the right side of the brain to the left via intact lower-level limbic pathways. For instance, one patient had a picture of Hitler lateralized to his right hemisphere (Sperry 1985). The speaking left hemisphere was unable to name the person. Astonishingly however, it did report that the person had done something negative.

In a recent series of studies on confabulation in two of the original series of split-brain patients (AA and LB), we set up an experiment designed to test for the presence of a non-informational confidence signal. The experimental set-up was as follows: The subject was seated and fixed

his eyes on the nose of an assistant sitting directly across from him, about four feet away. The experimenter sat on the subject's left. For each trial, the experimenter selected an object from a set of about ten and held it up, either in the patient's left visual field, so that the right hemisphere could see it, or behind the patient's head, so that neither hemisphere could see it. The subject was then asked to name the object being held up. Both the subjects were far more likely to confabulate when the right hemisphere was able to see the object, perhaps due to the fact that it was able to send a confidence signal to the left.

In order to quantify the experimental results, we devised an objective measure of amount of confabulation, based on the idea mentioned above that failure to say "I don't know" can in certain situations be a confabulation (Mercer et al. 1977, Dalla Barba 1993). Suppose you are asked ten rather difficult questions, to which you must answer either 'yes', 'no', or 'I don't know'. The ideal here is to get all ten questions right of course. But given that you are fallible, there are going to be some things that you don't know, so the next best thing to giving a correct answer is to avoid giving wrong answers, something you can achieve by being aware of your epistemic limitations and answering "I don't know" when you don't have clear, warranted knowledge of the answer. Given your fallibility, your best strategy is to answer the ones you know, and abstain from answering the others, by simply saying "I don't know." Given this, it is better to say "I don't know" than to give a wrong answer. An index of the amount of confabulation then can be devised as follows: The ratio of the number of incorrect answers to the don't know answers added to the incorrect answers: wrong/(don't know + wrong). We found that the patients were roughly twice as likely to confabulate when the right hemisphere had access to the appropriate information. We used a second measure in order to ensure that the patients were genuinely confabulating and not merely throwing out guesses because that was what the experimenters seemed to be calling for. After each response, the patient was instructed to indicate the level of his confidence in his answer, on a scale from one to five. Average levels of confidence were significantly lower in the trials in which neither hemisphere had access to the answer.

There is also a 'negative' way in which the phenomenon of a confidence signal may bring about confabulation. It may be that the left hemisphere will confabulate in the absence of 'lack of confidence' signal from the right hemisphere. A number of different lines of inquiry support the idea that something like a lack of confidence or an alarm signal may be detectable simply by measuring the person's skin conductance response (formerly known as 'galvanic skin response'). The difference, for instance, between lying and confabulation is that the liar is fully aware that what she is saying is false. The point of using skin conductance response to detect lying (it is

the primary measure of the polygraph test) is that in a normal person, lying is accompanied by a skin conductance response. Sociopaths, on the other hand, are famous for being very convincing speakers, for being able to sell just about anything to just about anyone. It has been shown that such people do not get the normal skin conductance response when they lie. They lie convincingly because, while they are aware that they are lying, the awareness lacks the emotional sting which such an awareness would bring in a normal person. Another line of evidence in support of the idea that skin conductance response may be a symptom of the activation of a 'lack of confidence' signal which the right hemisphere sends is the finding that skin conductance response is initiated primarily by the right hemisphere. It is well-confirmed that damage to the right hemisphere can seriously diminish or abolish the skin conductance response. In addition, there is some evidence that damage to the *left* hemisphere has a sort of releasing effect on skin conductance response, causing an increase in it. More pertinent to our concerns here, often the sort of right hemisphere damage patient which has an abolished skin conductance response (even to surprise or strenuous exertion) is a neglect patient.

5. Confabulation Succeeds in the Social Milieu. Confabulation is driven by a kind of epistemic overconfidence, the opposite of which we might call epistemic overprudence. One can be too prudent epistemically, abstaining completely when there is any doubt at all, as Descartes did at the beginning of his famous Meditations, where he refuses to hold any beliefs gained via his senses. This "strategy" is just as bad as being confabulatory, however, since the overly prudent person is completely paralyzed by his lack of certainty. An epistemically overprudent weatherman, for example, who when asked about tomorrow's weather always replies, "I don't know," or "Can't be sure," will soon be an unemployed weatherman. Excessive epistemic prudence is especially counterproductive when it occurs in someone who has a position of power or responsibility. Imagine a general who was never sure what to do, and as a consequence never did anything, never instructed his soldiers because he always felt that he didn't know what the right strategy was. Armies which don't move forward aggressively are soon overtaken by their enemies. Sometimes, then, an answer which is possibly wrong is better than none at all.

The high levels of success which sociopaths often attain in society bears witness to the idea that a confabulatory strategy works. But it is important to notice that this strategy works only on other people, who are subject to being fooled by fast talk coming from a person who appears very confident. In a nonsocial environment, someone who can lie convincingly has no advantage over any other person. Presenting oneself in a favorable way

to conspecifics is a well-known evolutionary strategy. The conceptual system in the left hemisphere, then, can be seen as a social part of the brain, one of the main goals of which is to win over other people, using language. The analog system based in the right hemisphere, on the other hand, is a pure representation system, devoted to accurately capturing the world, on this hypothesis.

Given a brain structured as I have described it, it is not clear why there should be any problem: the agent simply uses the conceptual system when dealing with other people, but uses the accurate analog system for the planning and execution of his own behavior. A person like this would be too inconsistent, however, and would quickly be found out as a liar. Some sort of more complicated method of reconciling the two systems is needed. One sort of solution would be to control the activity of the two different representation and computation systems with some sort of executive process. One of the more interesting consequences of such a view in general is that it implies that the human intellect, instead of being a seamless, unified process is actually a combination of at least two systems, one of which is rather brittle and ignorant without its analog partner.

6. Back to Normal Self-Deception. Mele expresses the philosopher's exasperation with the problem of self-deception: If I deceive you into believing that p, (normally) contained in this claim is the idea that I believe that not p. Applied to self-deception, then, when I deceive myself, I cause myself to believe that p, while at the same time believing that not-p. "But, how is this possible?" he asks (1987, 121).

One source of difficulty may be that we have trouble thinking of the information in the analog system as consisting of *beliefs*. Searle (1983) gives an example of a problematic epistemic state, which we are not sure is a belief or not. Suppose I were to enter my office one day and find a huge chasm in the floor. I would be surprised, now, does this mean that I believe that my office floor is solid, or worse, that I have a belief that there is no chasm in the middle of my office floor? We get a robust intuition that there is something wrong in ascribing such beliefs. I suggest that this is because, while I do represent my office floor as solid, I represent this in analog form, rather than in the form of explicit belief. I have been in my office hundreds of times, and my right hemisphere body-in-space representation system has represented the room each time.

Self-deception may occur, then, when a belief that not p is contained in the conceptual system, while the analog system faithfully continues to represent that p. This points a way to resolving a problem with Mele's theory. How can a person selectively avoid evidence for the belief that p without in some way believing that p? How else does he know what evi-

dence to avoid? One way to deal with this is as follows: the information that p is already represented in the analog system. There is a sort of conflict in the person's mind, but the conflicting information is represented in two different forms, conceptual and analog. This is different from holding two contradictory beliefs in full conceptual form. What is happening is that the brain has a way of preventing certain types of analog information from being represented in conceptual form, from being explicitly thought and believed. Self-deception is traceable to this phenomenon.

Another connection to Mele's account is as follows. One of the techniques which Mele claimed people use to keep an undesirable belief at bay is what he calls selective focusing and attending. One difference between the 'what'-oriented conceptual system in the left hemisphere, and a 'where'-oriented analog representation system in the right is that the 'what' system is fed primarily by focal vision, whereas the 'where' system is fed primarily by peripheral vision. The confabulatory conceptual system can selectively ignore information in the periphery by keeping the eyes focused elsewhere. More drastically, by simply refusing to receive any information at all from the analog system, the conceptual system can prolong its illusions.

In general, the difference between a confabulatory neurological patient and a normal self-deceived person is that the normal person experiences a sense of conflict because of the conflicting information contained in the two systems. The neurological patient, on the other hand, is completely at ease, and confabulates with a felicity even a sociopath would admire, because the analog system which would normally represent the conflicting information has been destroyed (or disconnected). At different points in our testing of split-brain patients, after they had given a confabulatory answer, we would ask them "Are you sure?" They would immediately answer "yep" or "positive." A self-deceived person, on the other hand would either pause to think or answer defensively if asked whether he was sure his wife wasn't having an affair, or his son wasn't dealing drugs.

I will end by describing a couple of the advantages of this account of self-deception: First, it is grounded in experimental findings, rather than a priori psychological theorizing. And second, it makes sense of the idea that conflicting information in the mind of the self-deceived person is segregated in some way, without relying on the troublesome notion of subagents within a person with their own intentional states, as Pears's theory does.

REFERENCES

Bear, David M. (1979), "Personality Changes Associated with Neurologic Lesions", in A. Lazare (ed.), *Textbook of Outpatient Psychiatry.* Baltimore: Williams and Wilkins.

Cabeza, Roberto, Shitji Kapur, Fergus I. M. Craik, Anthony McIntosh, Sylvain Houle, and Endel Tulving (1997), "Functional Neuroanatomy of Recall and Recognition: A PET Study of Episodic Memory", *Journal of Cognitive Neuroscience* 9: 254–265.

Churchland, Patricia S. and V.S. Ramachandran (1993), "Filling In: Why Dennett is Wrong", in Bo Dahlbom (ed.), *Dennett and His Critics: Demystifying Mind.* Oxford: Blackwell Scientific Press, 28–52.

Dalla Barba, Gianfranco (1993), "Different Patterns of Confabulation", *Cortex* 29: 567–581.

Gazzaniga, Michael S. (1985), *The Social Brain.* New York: Basic Books.

———. (1995), "Principles of Human Brain Organization Derived From Split-Brain Studies", *Neuron* 14: 217–228.

Hirstein, William S. and V.S. Ramachandran (1997), "Capgras Syndrome: A Novel Probe for Understanding the Neural Representation of the Identity and Familiarity of Persons", *Proceedings of the Royal Society of London* 264: 437–444.

Joseph, Rhawn (1996), *Neuropsychiatry, Neuropsychology, and Clinical Neuroscience.* Baltimore: Williams and Wilkens.

Mele, Alfred R. (1987), *Irrationality: An Essay on Akrasia, Self-Deception, and Self-Control.* Oxford: Oxford University Press.

Mercer, Brian, W. Wapner, H. Gardner, and Frank Benson (1977), "A Study of Confabulation", *Archives of Neurology* 34: 429–433.

Pears, David (1984), *Motivated Irrationality.* Oxford: Oxford University Press.

Ramachandran, V.S. (1993), "Filling in Gaps in Logic: Some Comments on Dennett", *Consciousness and Cognition* 2: 165–168.

———. (1995), "Anosognosia in Parietal Lobe Syndrome", *Consciousness and Cognition* 4: 22–51.

Ramachandran, V.S. and Richard L. Gregory (1991), "Perceptual Filling In of Artificially Induced Scotomas in Human Vision", *Nature* 350: 699–702.

Ramachandran, V.S. and Willam Hirstein (1997), "Three Laws of Qualia: What Neurology Tells Us About the Biological Functions of Consciousness", *The Journal of Consciousness Studies* 4: 429–457.

Searle, John R. (1983), *Intentionality: An Essay in the Philosophy of Mind.* Cambridge: Cambridge University Press.

Sperry, Roger W. (1985), "Consciousness, Personal Identity, and the Divided Brain", in D. F. Benson and Eran Zaidel (eds.), *The Dual Brain: Hemispheric Specialization in Humans.* New York: Guilford, 11–25.

Sperry, Roger. W., Eran Zaidel, and Dahlia Zaidel (1979), "Self Recognition and Social Awareness in the Deconnected Minor Hemisphere", *Neuropsychologia* 17: 153–166.

Tulving, Endel. (1983), *Elements of Episodic Memory.* Oxford: Clarendon Press.

Ungerleider, Leslie G. and Mortimer Mishkin (1982), "Two Cortical Visual Systems", in D. J. Ingle, Melvin A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of Visual Behavior.* Cambridge, MA: MIT Press, 549–586.