

Searle on the Brink

Selmer Bringsjord

Department of Philosophy, Psychology and Cognitive Science
Rensselaer Polytechnic Institute
Troy, NY 12180. USA

selmer@rpi.edu

Copyright (c) Selmer Bringsjord 1994

PSYCHE, 1(5), August 1994

<http://psyche.cs.monash.edu.au/v2/psyche-1-5-bringsjord.html>

Keywords: cognitive science, consciousness, mind, property dualism, reductionism

Abstract: In his recent *The Rediscovery of the Mind* John Searle tries to destroy cognitive science *and* preserve a future in which a ``perfect science of the brain" (1992, p. 235) arrives. I show that Searle can't accomplish both objectives. The ammunition he uses to realise the first stirs up a maelstrom of consciousness so wild it precludes securing the second.

1.

1.1 What Dennett (1993) considers quixotic I regard accomplished: he thinks Searle's (1992) attempt to overthrow orthodox cognitive science [<1>](#) is at best a pipe dream, and at worst the near ravings of a very forgetful man [<2>](#) -- but I think orthodox cognitive science is *already* overthrown, for Searlean reasons, and others. [<3>](#) I worry, though, about Searle's desire, expressed in his recent *The Rediscovery of the Mind* (= RM), to both destroy cognitive science *and* preserve a future in which a ``perfect science of the brain" (1992, p. 235) arrives. [<4>](#) It seems to me that Searle stands at the brink of a maelstrom of consciousness so wild it won't *ever* be tamed by science. In what follows, I focus this metaphor by analyzing a pivotal section in RM.

1.2 The section I have in mind begins with Searle's serviceable synopsis of the well-known Jackson (1982)-Kripke (1971)-Nagel (1974) argument for the irreducibility of consciousness:

Consider what facts in the world make it the case that you are now in a certain conscious state such as pain. What fact in the world corresponds to your true statement, ``I am now in pain"? Naively, there seem to be at least two sorts of facts. First and more important, there is the fact that you are now having certain unpleasant conscious sensations, and you are

experiencing these sensations from your subjective, first- person point of view. It is these sensations that are constitutive of your present pain. But the pain is also caused by certain underlying neurophysiological processes consisting in large part of patterns of neuron firing in your thalamus and other regions of your brain. Now suppose we tried to reduce the subjective, conscious, first-person sensation of pain to the objective, third-person patterns of neuron firings. Suppose we tried to say the pain is really ``nothing but" the patterns of neuron firings. Well, if we tried such an ontological reduction, the essential features of the pain would be left out. No description of the third-person, objective physiological facts would convey the subjective, first-person character of the pain, simply because the first-person features are different from the third-person features. Nagel states this point by contrasting the objectivity of the third-person features with the what-it-is- like features of the subjective states of consciousness. Jackson states the same point by calling attention to the fact that someone who had a complete knowledge of the neurophysiology of a mental phenomenon such as pain would still not know what a pain was if he or she did not know what it felt like. Kripke makes the same point when he says that pains could not be identical with neurophysiological states such as neuron firings in the thalamus and elsewhere, because any such identity would have to be necessary, because both sides of the identity statement are rigid designators, and yet we know that the identity could not be necessary. (pp. 117-118)

As I say, this is an adequate synopsis of the argument. (In saying this I tacitly agree to conflate Jackson with Kripke despite the fact that Jackson's argument presupposes nothing about rigid designators, whereas Kripke's does. Jackson, it ought to be noted, didn't simply repeat Kripke's argument. At any rate, in a moment I set out the argument more carefully, in *Jacksonian* style.) But it will be helpful to both anchor things with a thought-experiment and to set out a version of the argument *with premises and inferences explicit*.

1.3 The protagonist in our thought-experiment is a Jacksonian character created elsewhere for my version of the argument in question, which appears in (Bringsjord, 1992). The character is Alvin, a cognitive scientist who lives and works in an isolated laboratory. Suppose, indeed, that Alvin, for the past five years, has been an absolute recluse, that during this time he has hardly had any contact with other humans, that what contact he has had has all been of a professional, scientific nature, and so on. Alvin, during this time, has mastered the purportedly complete reductive, cognitive scientific specification of human mentation. Suppose, in addition, that he has never encountered a long-lost friend -- Alvin has never even had an experience remotely like this.

1.4 Now, one day Alvin leaves his drab lab and encounters a long-lost friend, and thereby learns what it feels like ``on-the-inside" to meet a long-lost friend. ``So *this*," he says to himself, ``is what it feels like to meet a long-lost friend in the flesh, to see once again that gleam in her eyes, the light her hair seems to catch and trap...." Etc., etc.

1.5 The corresponding argument in (Bringsjord, 1992) is too complicated to reproduce here. So I'll adapt to Alvin's situation a more informal but elegant and powerful statement of the Jacksonian argument given by Jacqueline (1994):

Argument 1 (A1)

(1) To know everything knowable about a psychological state is to have complete first- and third-person knowledge of it.

(2) Alvin, prior to his initial first-person long-lost-friend experience, knows everything knowable about meeting long-lost friends from a third-person scientific perspective.

(3) To know everything knowable about meeting long-lost friends from a first-person perspective implies knowing what it's like to meet a long-lost friend in the flesh.

(4) Alvin, prior to his first-person long-lost-friend experience, doesn't know what it's like to meet a long-lost friend in the flesh.

(5) If reductivist cognitive science is true, then if Alvin, prior to his first-person long-lost-friend experience, knows everything knowable about meeting long-lost friends from a third-person scientific perspective, then, prior to his initial first-person long-lost-friend experience, he knows everything knowable about meeting long-lost friends.

Therefore (from 1, 3 & 4):

(6) Alvin, prior to his first-person long-lost-friend experience, doesn't know everything knowable about meeting long-lost friends.

Therefore (from 2, 5 & 6):

(7) Reductivist cognitive science is false.

This argument is obviously formally valid (as can be shown when it's symbolized in the propositional calculus).

1.6 Now that the argument is before us, let's be clear about Searle's verdict on the reasoning in question -- a verdict which is quite unmistakable: This argument, Searle says, "is ludicrously simple and quite decisive" (p. 118).

2.

2.1 This verdict would seem to imply that Searle's view of consciousness is non-scientific (or ascientific), since science at least *appears* to progress by reductively formalizing mysterious laic concepts as transparent ones susceptible of third-person analysis. But


Searle's position on consciousness, by RM itself, *can't* be non-scientific. For RM is, according to Searle himself, a specification and defence of what he calls "naive mentalism" (p. 54). Here's that view in a nutshell:

The brain causes certain "mental" phenomena, such as conscious mental states, and these conscious states are simply higher-level features of the brain. Consciousness is a higher-level emergent property of the brain in the utterly harmless sense of 'higher-level' or 'emergent' in which solidity is a higher-level emergent property of H₂O molecules when they are in a lattice structure... (p. 14)

2.2 Searle is well-aware of the tension between his affirmation of the irreducibility of consciousness on the one hand, and his championing of science (not *cognitive* science) on the other:


But to many people it seems that [the Jackson-Kripke-Nagel argument] paints us into a corner. To them it seems that if we accept that argument, we have abandoned our scientific world view and adopted property dualism. Indeed, they would ask, what is property dualism but the view that there are irreducible mental properties? In fact, doesn't Nagel accept property dualism and Jackson reject physicalism precisely because of this argument? (p. 118)

2.3 Searle tries to dissolve the tension by showing "that the irreducibility of consciousness is a trivial consequence of the pragmatics of our definitional practices" (p. 122). He begins by directing our attention to cases of successful reduction. For example, consider the story of HEAT, a story which, in broad strokes, begins with the nebulous, unanalyzed concept of HEAT, but ends with a gloriously precise account of the kinetic energy of molecular movement. The story skeleton, Searle tells us, is straightforward: HEAT is first divided into two sorts of properties, the subjective and objective. On the subjective side we might have a first-person property like *my feeling hot*; on the objective side we might have the third-person *the thermometer's registering m when immersed*. We then proceed to reduce the third-person property to concepts found in kinetic theory. A little picture sums up our storyline:



2.4 So far, Searle is sailing on clear water. The target for reduction, in this case, *is* some cluster of third-person properties, not private, subjective feelings. But what does this


unexceptionable analysis have to do with consciousness? Searle asks us to consider the storyline produced by applying this reduction protocol to conscious mental states. In order to fix the situation a bit, suppose that on the subjective side the property in question is *my fearing rabid bats*, whereas the objective side is encapsulated by the property *area 13 of the cortex registering PET-measured cerebral blood flow of X*. Suppose, in addition, that cerebral blood flow readings have been thoroughly cashed out in a robust theory ('N theory,' say) whose primitives are terms referring to low-level constituents of the brain -- neurons, dendrites, etc. Our storyline then looks like this:



2.5 You can doubtless see where Searle wants this storyline to take him. He says the moral of the story is that consciousness is irreducible for the simple and undeniable reason that the reduction of properties like *my fearing rabid bats* to N theory is what it would take to achieve the reduction -- but, so the story goes, this reduction isn't accomplished.

3.

3.1 There are some bad ways to reply to Searle's reasoning. One would be to claim, a la Paul Churchland (1984), that Searle suffers from a lack of imagination. For, so the rebuttal goes, he need only conceive of a futuristic, but not improbable, scenario in which PET measurements (or something similar) *supplant* talk of such things as fearing rabid bats. That is, suppose that by 2027 humans are rigged up to self-monitoring machines [I called such devices 'brain boxes' in (Bringsjord, 1986)] which measure cerebral blow flow (or something similar); and that by this time, instead of exclaiming, upon finding a brown bat clinging to the wall above one's bed, ``Argh!! I'm petrified!``, one screams ``Ahh! Layer IV has just maxed out at 80 Hz!`` This scenario even comes with its own storyline, modified so as to lack the dead ends Searle has pressed against reductionists:



|
v
N theory

3.2 This scenario may be imaginative, but it doesn't produce a sound argument. In fact, the "zombie" thought-experiments Searle presents elsewhere in RM shoot down the rebuttal. In one such experiment a fellow whose brain is deteriorating receives a silicon-based replacement, and finds, to his horror, that, on the inside, he is dying away to non-consciousness, while, on the outside, his behavior continues smoothly -- so smoothly that the neurosurgeons consider their procedure to be a smashing success. This scenario shows that a brain box future is merely one in which conscious states are *correlated with* brain states. There's no conceptual connection; there needn't even be a causal connection. And for that reason it's rather a stretch to declare, as in Figure 3, that consciousness is *reduced* to N theory.

3.3 Other thought-experiments bring home the same point. Suppose that the highly complex psychological profiles of nations (which certain governmental organizations have traditionally been in the business of charting) can be supplanted with pictures taken by satellite from outerspace of the distribution of bodily heat across the countries in question. So instead of an analyst for the CIA giving the U.S. Joint Chiefs of Staff a disquisition on the mindset of business people and politicians, etc. in Iraq, she merely presents snapshots that look something like random dot stereograms, and offers her predictions accordingly. Would the Chiefs be entitled to hold that the folk psychological states which formed the core of such briefings in the past had been *reduced*, ontologically speaking, to abstract technicolor snapshots? Hardly.

4.

4.1 So where does this leave us? We have the result that, given *today's scheme* for reduction, consciousness can't be *reduced* -- a result Searle produces (via the reasoning we've just unpacked on his behalf) and cheerfully embraces (p. 124). But Searle then tells us that consciousness' "irreducibility has no untoward scientific consequences whatever" (p. 124). It's this comment, alas, which signifies, it seems to me, the placing of his hands in front of his eyes so as to stay blind to the crevasse which yawns beneath him. There's a mighty big difference between *irreducible on today's scheme for reduction* and *irreducible, period*. The former -- which, we can agree, doesn't doom science of the brain -- is what he has just established. But the latter is supposed to be the conclusion of the Jackson-Kripke-Nagel argument! And if consciousness is irreducible come what may, it's exceedingly hard to see how Searle's sanguinity is more than blind faith. For recall Searle's encapsulation of naive mentalism:

The brain causes certain "mental" phenomena, such as conscious mental states, and these conscious states are simply higher-level features of the brain. Consciousness is a higher-level emergent property of the brain in the utterly harmless sense of 'higher-level' or 'emergent' in which solidity is a higher-level emergent property of H₂O molecules when they are in a lattice structure... (p. 14)

If Jackson and Co. are right, consciousness isn't harmless in the least, for it would not be *possible* to cash it out in terms of lower-level properties of the brain and brain parts.

4.2 This is why I say Searle stands on the brink of a maelstrom of consciousness so wild it won't *ever* be tamed by science.

5.

5.1 Of course, my analysis isn't entirely unproblematic. For one thing, someone might object on behalf of Searle as follows: ``You've told us that 'science at least *appears* to progress by reductively formalizing mysterious laic concepts as transparent ones susceptible of third-person analysis.' It's true that this appearance is taken as gospel by *non*-philosophers of science, Searle included. But appearances can be deceiving -- and this one is precisely that. For without being either hypercritical or radical, it's fair to say that most of what we call scientific progress has *nothing to do* with laic concepts and therefore cannot be interpreted as reducing the latter to anything at all! So, if we're charitable, when Searle says that the irreducibility of consciousness is no impediment to the scientific study of consciousness, we can read him as holding that such study need not involve the sort of reduction you feature."

5.2 Two types of reply to this objection come immediately to mind: We could attack the latitudinarian view of scientific progress presupposed by it. Or, we could retreat to the claim that on *Searle's* view of scientific progress, his vision of a perfect science of the brain is forever unattainable.

5.3 Replying in either of these ways is problematic. The first route would simply take us too far afield. The second has the present essay produce, at most, the result that Searle is inconsistent. Fortunately, there's a third route: viz., focus on what brain scientists working on the problem of consciousness think about the situation. Here, for example, is part of what Patricia Churchland (1994) has to say about Searle's view that consciousness is irreducible:

Synoptically, here is why Searle's manoeuvre is unconvincing: he fails to appreciate why scientists opt for identifications when they do. Depending on the data, cross-level identifications to the effect that *a* is *b* may be less troublesome and more comprehensible scientifically than supposing thing *a* causes separate thing *b*. This is best seen by example.... Science as we know it says electrical current in a wire is not caused by moving electrons; it *is* moving electrons. Genes are not caused by chunks of base pairs in DNA; they *are* chunks of base pairs (albeit sometimes distributed chunks). Temperature is not caused by mean molecular kinetic energy; it *is* mean molecular kinetic energy. Reflect for a moment on the inventiveness required to generate explanations that maintain the *nonidentity* and causal dependency of (a) electric current and moving electrons, (b) genes and chunks of DNA, and (c) heat and molecular motion. Unacquainted with

the relevant convergent data and explanatory successes, one may suppose this is not so difficult. Enter Betty Crocker. (Churchland, P.S., 1994, p. 30)

Churchland goes on to claim that Betty Crocker's view of microwave cooking, which is predicated upon the nonidentity of heat and molecular motion, is unworkable scientifically. The Betty Crocker approach, in science, Churchland (1994, p. 30) tells us, would be "like trying to nail jelly to the wall."

5.4 But what about *consciousness* and reduction? Is it really true that brain scientists are already thinking in terms of reduction when they dream of scientific progress on this front? Absolutely. In fact, Churchland (1994) introduces two programs heading in this direction: one piloted by Crick (1990, 1994) and one by Llinas (1991, 1993), two brain scientists who presumably share Churchland's views on Searle and reduction.

5.5 It would seem, then, that Searle's "consciousness is irreducible come what may" position, by the lights of contemporary brain science, *has* pushed him to the brink of the maelstrom Churchland and like-minded neuroscientists take pains to avoid.

6.

6.1 Now, what might *Searle* say about my diagnosis of RM? What if he removes his hands from in front of his eyes and takes an unblinking look at the storm of consciousness brewing below? I think he would probably say something he said nowhere in RM: that the Jackson-Kripke-Nagel argument doesn't establish irreducibility, period; it establishes only what Searle has established by taking a different route (the one involving an analysis of reduction as practised, which we traversed above): viz., that consciousness can't, *today*, be reduced.

6.2 But how, exactly, is the Jackson-Kripke-Nagel argument to be modified? What modification of (A1) is entailed by the move we're attributing to Searle? Unfortunately, if the modification is to index the reasoning in question to *current knowledge* and preserve formal validity, as it must, then Searle is in some trouble. This is so because the resultant argument can't be sound. Consider, for example,

Argument 1' (A1')

(1') To know everything *currently* knowable about a psychological state is to have complete first- and third-person knowledge of it.

(2') Alvin, prior to his initial first-person long-lost-friend experience, knows everything *currently* knowable about meeting long lost friends from a third-person scientific perspective.

(3') To know everything *currently* knowable about meeting long lost friends from a first-person perspective implies knowing what it's like to meet a long-lost friend in the flesh.

(4') Alvin, prior to his first-person long-lost-friend experience, doesn't know what it's like to meet a long-lost friend in the flesh.

(5') If reductionist cognitive science is true, then if Alvin, prior to his first-person long-lost-friend experience, knows everything *currently* knowable about meeting long-lost friends from a third-person scientific perspective, then, prior to his initial first-person long-lost-friend experience, he knows everything *currently* knowable about meeting long-lost friends.

Therefore (from 1', 3' & 4):

(6') Alvin, prior to his first-person long-lost-friend experience, doesn't know everything *currently* knowable about meeting long-lost friends.

Therefore (from 2', 5' & 6):

(7') Reductionist cognitive science is false.

6.3 (A1'), like its predecessor, (A1), is formally valid. But of course the argument is unsound: (2') and (4) are true [more precisely: (2') reflects the implied adjusted thought-experiment, and (4) is just as plausible as its counterpart in (A1), which is itself (4)], but (1') is rather obviously false, since current knowledge may be impoverished relative to all that is knowable.

6.4 So Searle can't back away from the brink so easily.

6.5 Ironically, Dennett's own position is that (A1') is the best that can be mustered. (Dennett is of course *happy* to find that (A1') is unsound!) He thinks that (A1)'s

(2) Alvin, prior to his initial first-person long-lost-friend experience, knows everything knowable about meeting long-lost friends from a third-person scientific perspective.

is unimaginable, and that once (2') -- which, he concedes, corresponds to an imaginable scenario -- supplants this premise, we get at best (A1'), which, as we've seen, *is* unsound. Here's what Dennett has to say about (2):

The image [of Alvin] is wrong; if [seeing Alvin make a discovery] is the way you imagine the case, you are simply not following directions! The reason no one follows directions is because what they ask you to imagine is so preposterously immense, you can't even try. The crucial premise is that "[Alvin] has *all* the physical [computational] information." This is not readily imaginable, so no one bothers. They just imagine that [he] knows lots and lots -- perhaps they imagine that [he] knows everything that anyone knows *today* about the neurophysiology [etc.] of [such psychological states]. But that's just a drop in the bucket, and it's not surprising that [Alvin] would learn something if *that* were all [he] knew. (Dennett, 1991, p. 399)

6.6 Dennett has been conveniently struck by an uncharacteristic failure of imagination. The fact of the matter is that the thought-experiment and corresponding argument can be put in austere terms wholly divorced from the particular sciences of today. Jackson and Co. needn't refer to existing scientific theories in the least; they can put their point exclusively in terms of the very framework which allows us to consider the question in the first place. The generalized Jacksonian thought-experiment, a parallel of the one I offered in (Bringsjord, 1992), would be formulated as follows.

6.7 An account of conscious mental states is still denoted in familiar fashion, i.e., by way of folk psychological expressions of a sort visited above. We still say such things as that "Alvin doesn't want to close his eyes and go to sleep, because he fears rabid bats, and fears one may at any moment swoop into his bedroom." But on the other side, when considering candidates for what such folk accounts are to be reduced to, we simply consider mathematical objects *as general as possible*. For example, the candidate comes in the context of a logical system [5](#) which can be as powerful as one likes. It could even be an infinitary logical system [6](#) quite beyond the capacity of a Turing machine to formalize (in which case we're immediately generalizing well beyond what cognitive science, bound as it traditionally is to the algorithmic, has to offer).

6.8 Accordingly: It's 2027, and Alvin, a reclusive, blind cognitive scientist, is given the purportedly complete reductive cognitive scientific formalization of the psychological state of meeting long-lost friends (and all other relevant compu-physical scientific information), expressed as formulas in some set D in an infinitary logical system L. Alvin assimilates the specification. [Picture him there in his laboratory "reading" (via Braille) formula after formula. If you like, imagine him "reading" infinitely many formulas by working faster and faster in the perfectly imaginable fashion of Zeus machines (Boolos & Jeffrey, 1980)]. After his sedulous stint is up, Alvin leaves his lab and meets a long-lost friend -- an experience he never had before. "So *this*," he says to himself, "is what it's like to meet a long lost friend!" [7](#)

6.9 This argument can be set out explicitly by simply inserting the parenthetical '(D in L)' at the appropriate place in (A1)'s (1), (2) and (5). The result is a formally valid argument the premises of which dodge Dennett's complaint.

6.10 Alvin's discovery is all entirely possible, no matter how wondrous D and L are; and D and L are wholly arbitrary over the domain of *all* logical systems, a class which in turn spans the space of substrates for scientific theories and the constituents thereof. Ergo, conscious mental states (here, those associated with meeting a long-lost friend), not only haven't been reduced in the parable to the L-based specification: such states are irreducible, period.

6.11 You may demand to see this argument worked out more formally [in which case you would want but an extension of the "Argument from Alvin" in (Bringsjord, 1992, Chapter I)], and you would of course be well within your rights in so demanding, but such a demand would almost certainly come from someone who was sceptical about this argument's ancestor, (A1). Searle, however, has, as we've seen, *affirmed* this reasoning.

This affirmation, despite his attempt to trivialize it via his discussion of current patterns of reduction, carries him to the brink of consciousness untamed and untameable -- a brink over which many thinkers have voluntarily plunged,<8> but only once they shed the chains of a hoped-for science of the mind to which Searle, by his own admission, is still bound.<9>

Notes

<1> The textual evidence supporting the view that Searle intends nothing less than to demolish cognitive science is quite unmistakable. As Dennett (1993, p. 195) puts it,

The central doctrine of cognitive science is that there is a level of analysis, the information-processing level, intermediate between the phenomenological level (the personal level, or the level of consciousness) and the neurophysiological level. Searle sees that his position requires that this central doctrine be entirely and hopelessly mistaken. [Searle says] ``There are brute, blind neurophysiological processes and there is consciousness, but there is nothing else" (Searle, 1992, p. 228).

<2> I'm afraid I must insist that I'm not guilty of hyperbole here. Consider these words of Dennett's (1993, p. 203):

Is it possible that although Searle has at one time or another read all the literature, and understood it at the time, he has actually forgotten the subtle details, and (given his supreme self-confidence) not bothered to check his memory? For instance, has he simply forgotten that what he calls his *reductio ad absurdum* of my position (81 [in (Searle, 1992)]) is a version of an argument I myself composed and rebutted a dozen years ago? There is evidence of extreme forgetfulness right within the book. For instance...

In the next paragraph, speaking about another of Searle's supposed lapses, Dennett says, ``But he forgets all this (apparently!) when forty pages later (107 [in (Searle, 1992)]) he sets out to explain the evolutionary advantage of consciousness...."

<3> Chapter V of (Bringsjord, 1992) is a victorious reformulation of Searle's Chinese Room Argument. This monograph contains a number of other arguments designed to overthrow ``strong" AI/cognitive science.

<4> I intimated these worries in the swift and impressionistic (Bringsjord & Patterson, forthcoming). From this point on, unadorned page numbers refer to RM.

<5> See, for example, Ebbinghaus, Flum & Thomas (1984, Chapter XII).

<6> For a brief look at an elementary system of this sort, see Ebbinghaus, Flum & Thomas (1984, Chapter IX). For a more mature study of such systems see Dickman, M. A. (1975).

<7> The point here can be brought home all the better, it would seem, if we throw in a Greek tragic twist: let the L-based specification of Smith's life be a specification of Alvin's, unbeknownst to Alvin! I flesh out this twist, in the context of L set to the logical systems equivalent in power to Turing machines, in Bringsjord (1992, Chapter I).

<8> James Ross (1993) has recently defended, with aplomb, his taking the plunge.

<9> I'm greatly indebted to Kevin Korb and two anonymous referees for presenting a number of trenchant objections which led to improvements in the paper.

References

Boolos, G. S. & R. C. Jeffrey (1980). *Computability and logic*. Cambridge, UK: Cambridge University Press.

Bringsjord, S. (1992). *What robots can & can't be*. Dordrecht, The Netherlands: Kluwer.

Bringsjord, S. (1986). Swinburne's argument from consciousness. *International Journal for the Philosophy of Religion*, 19: 127- 143.

Bringsjord, S. & Patterson, W. (forthcoming). A perfect science of the brain?, a review of John Searle's *The Rediscovery of the Brain*. *Minds & Machines*.

Churchland, P. M. (1984). *Matter and consciousness*. Cambridge, MA: MIT Press.

Churchland, P. S. (1994). Can neurobiology teach us anything about consciousness? *Proc. APA*, 67(4): 23-40.

Crick, F.H.C. (1994). *The astonishing hypothesis*. New York, NY: Scribner's & Sons.

Crick, F.H.C. & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 4: 263-276.

Dennett, D. (1993). Review of John Searle's *The Rediscovery of the Brain*. *The Journal of Philosophy*, 93: 193-205.

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

Dickman, M. A. (1975). *Large infinitary languages*. Amsterdam, The Netherlands: North-Holland.

Ebbinghaus, H. D., Flum, J., Thomas, W. (1984). *Mathematical logic*. New York, NY: Springer-Verlag.

Jackson, F. (1982). Epiphenomenal qualia, *Philosophical Quarterly*, 32: 127-136.

- Jacquette, D. (1994). *Philosophy of mind*. Englewood Cliffs, NJ: Prentice Hall.
- Kripke, S. (1971). Naming and necessity. In D. Davidson and G. Harman (Eds.), *Semantics of natural language* (pp. 253-355, 763-769). Dordrecht, The Netherlands: Reidel.
- Llinas, R.R. & Ribary, U. (1993). Coherent 40-Hz oscillation characterizes dream state in humans. *Proc. National Academy of Sciences*, 90: 2078-2081.
- Llinas R.R. & Pare, D. (1991). Of dreaming and wakefulness. *Neuroscience*, 44: 521-535.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83: 435-450.
- Ross, J. (1993). Immaterial aspects of thought. *The Journal of Philosophy*, 84: 136-150.
- Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.