The Journal of Mind and Behavior
Autumn 1999, Volume 20, Number 4
Pages 391–402
ISSN 0271–0137

Why Isn't Consciousness Empirically Observable? Emotion, Self-Organization, and Nonreductive Physicalism

Ralph D. Ellis

Clark Atlanta University

Most versions of the knowledge argument say that, since scientists observing my brain wouldn't know what my consciousness "is like," consciousness isn't describable as a physical process. Although this argument unwarrantedly equates the physical with the empirically observable, we *can* conclude, not that consciousness is nonphysical (some physical processes *might* be observationally inaccessible) but that consciousness isn't identical with anything empirically observable. But what kind of mind–body relation would render possible this empirical inaccessibility of consciousness? Even if multiple realizability may allow a distinction between consciousness and its physical substrata, why does this distinction make consciousness *empirically unobservable*? The reason must be that the emotions motivating attention direction, partly constitutive of phenomenal states, are *executed*, not *undergone* by self-organizing processes actively appropriating and replacing needed physical substrata; we feel motivations by *generating* them. But all consciousness is motivated; visual cortex activation is *unconscious* of red unless the emotional limbic system and anterior cingulate motivatedly "look *for*" red. Experiencing entails executing motivations. Experimenters do know what subjects' brain events "are like" — but from the standpoint of *the experimenter's* motivational processes.

That there is something empirically unobservable about phenomenal consciousness follows from a modified version of the knowledge argument. Traditional versions (Jackson, 1986; Noren, 1979; Robinson, 1982) run essentially as follows: if experiencing were equivalent with physical brain states, then complete empirical knowledge of brain states should constitute knowledge of everything about my experiencing; but complete empirical knowledge of brain states would *not* constitute knowledge of everything about experiencing (those alone wouldn't reveal what it's like to have that

Requests for reprints should be sent to Ralph D. Ellis, Ph.D., Department of Religion and Philosophy, Clark Atlanta University, Atlanta, Georgia 30314. E-mail: ralphellis @mindspring.com

experience); therefore, experiencing is *not* equivalent with physical brain states. This argument can be criticized for unwarrantedly assuming that everything physical is *empirically observable* (from an experimenter's standpoint). For example, Jackson (1986, p. 291) assumes that the what it's like aspect isn't "expressible in physical language," but the reason for granting this assumption is that what it's like is inexpressible in terms of possible empirical observations. Without the assumption that everything physical is empirically observable, we can conclude, not that consciousness is non-*physical* (since there *might* be physical processes that are observationally inaccessible), but simply that consciousness isn't identical with anything empirically observable. Still, given the intimate connection between each conscious event (C) and a corresponding *empirically observable* physiological event (P), what P–C relationship could render C empirically *unobservable*? If identical, they should be equally observable. But if C were *non*-physical, it would then seem very mysterious how the systematic P–C correlations might be explained. Can C be *physically* instantiated, yet distinguishable from anything empirically observable?

One possibility is that C is a *relation* among Ps which is not identical to the Ps because it is multiply realizable (Bechtel and Mundale, 1999; Putnam, 1993). If C hadn't been realized by P1, it could have been realized by P2, yet still have been the *same* relation. C might be a self-organizing pattern of activity which *appropriates* and *replaces* physical substrata needed for its maintenance or evolution in self-motivated directions. Thus the P–C relation wouldn't be an identity, epiphenomenalism, or dualism, but a biological relation between an organismic *process* and its actively appropriated material components, continually replaced and reproduced through the organism's self-maintaining patterns of organization (Kauffman, 1993; Monod, 1971).

But how could this account explain the empirical *unobservability* of the relation C? As Bickle (1992) notes, something as simple as the temperature of a gas is multiply realizable by various movements of particles, yet temperature is not empirically unobservable!

Consciousness, unlike many self-organizing processes, is empirically unobservable because the emotions motivating the direction of conscious attention, partly constitutive *of* what it's like to experience anything, are *executed*, not *undergone* by organisms (Newton, 1996). Organisms feel motivations by *generating* them. But I shall suggest here that *all* phenomenal consciousness must be emotionally motivated; thus the what it's like aspect of a phenomenal experience is inseparable from the emotions that permeate it. So, visual cortex activation is *unconscious* of red unless the emotional limbic system and anterior cingulate motivatedly "look *for*" a red object (Aurell, 1989; Damasio, 1994; Posner and Rothbart, 1992). Directly experiencing someone's subjective consciousness entails executing her emotional motivations,

thus *being* the person whose organism produces them. An experimenter observing my brain events does indeed know what those events are like, but only for *her* organism, as motivated by *her* self-maintaining processes. Understanding or empathizing with another's emotions requires motivation by our *own* emotional processes, perhaps similar, but not *numerically identical* with the other's.

As discussed below, a self-organizing process is a particular *type* of multiply realizable one. In a self-organizing process, the internal *structure* of the organism is the reason for its tendency to appropriate and replace its own substratum as needed to preserve functional continuity. My suggestion is not that the motivations stemming from biological self-organization are a sufficient condition for consciousness to occur, but only a *necessary* condition. Even this much can resolve a good bit of the mystery as to how consciousness can be in principle empirically unobservable without thus becoming a metaphysical substance split off from the physical realm.

How Can an Empirically Unobservable Process Have Physical Substrata?

If the above modification of the knowledge argument entails that phenomenal consciousness is not identical with anything empirically observable, it also entails that consciousness must be a process or aspect of a process that appropriates the empirically observable physiological substrata needed for its maintenance, growth, and reproduction. This follows from the process of elimination; every other theory of the mind–body relation would entail untenable conclusions in light of the in-principle empirical unobservability of consciousness. That is, psychophysical identity and epiphenomenalism imply that consciousness should be directly empirically observable — which it is not — whereas dualism and interactionism of the Popper and Eccles (1977) variety (the claim that conscious and physical events causally interact, but that neither necessarily requires the other as its underpinnings) would imply that consciousness is a non-physical entity or event, requiring no physical substratum. Each of these implications can be shown to be untenable.

We have already seen that, if consciousness were *identical* with empirically observable events, consciousness itself would be empirically observable. But it isn't. As a general principle, no amount of empirical knowledge of physiological events, by itself, would entail knowledge of what it's like to experience the consciousness in question. The experimenter may *infer* what it might be like by comparing the empirically observable events to what *her own* experience would be like when in an analogous neurophysiological state, but that knowledge in turn would not result from *empirical* observations of her own brain, but rather from past subjective introspection. Thus it cannot

be inferred from any set of purely empirical observations alone that phenomenal consciousness would be like anything at all (Chalmers, 1995). Thus consciousness in principle cannot be identical with anything that is either directly or indirectly empirically observable.

Even granted that most empirical observations are at least partly indirect, consciousness is not even *indirectly* empirically observable in this sense. It is true, of course, that scientists empirically observe temperature only indirectly, by first observing a thermometer reading, and then *inferring* the temperature from this reading, along with theories that presuppose inductive and deductive inferences from previous observations. But the point is that temperature is the kind of phenomenon that can be inferred at least indirectly from some set of empirical observations alone, whereas what a state of consciousness is like cannot, even indirectly.

If the empirical unobservability of consciousness conflicts with psychophysical identity, it also rules out a purely epiphenomenalist approach. If consciousness were simply *caused by* its empirically observable physiological correlates (as opposed to being identical with them), as Jackendoff (1996) and Searle (1984) propose, then consciousness itself would have to be either physical or non-physical. If physical, it ought to be capable of causing other physical phenomena, just as any other physical phenomenon is. Epiphenomenalism must reject this possibility because its central thesis is that consciousness is caused *by* physical processes, rather than causing them. On the other hand, if the conscious events caused by physical processes were *non*-physical, epiphenomenalism would have to posit the existence of non-physical entities, thus becoming a metaphysical dualism. It would be very mysterious what the non-physical entities might consist of, and how physical entities could make causal contact with nonphysical ones.

Metaphysical dualism and interactionism are untenable for similar familiar reasons. Dualism is unable to account for the extensively documented regular correlations between conscious and physical events; and interactionism, because it requires that some physical events be caused by mental events, would imply that the chain of physical events should not be causally sufficient without input from the non-physical conscious events; yet we do not observe the frequent violations of chemical and physical patterns here entailed.

That consciousness is not empirically observable does not entail that it is a "non-physical" entity. Of course, the term "physical" here is not meant narrowly. Consciousness is not something that can be kicked, like a stone, but neither is temperature, a sound wave or a photon. None of these is non-physical, since each has relations to observable physical events which ultimately can be described in terms of the behavior of things that *can* be kicked.

If we eliminate psychophysical identity, epiphenomenalism, dualism, and interactionism, then the relation between consciousness and its physiological

substratum can be neither a causal relation between separate entities, nor a reducibility of one entity to another. What remains is a relation in which, as James (1908/1968) said, consciousness is not an entity but a function; moreover, unlike most functions it is in principle empirically unobservable. That consciousness could be a function, different enough from its physical substratum to be empirically unobservable, is possible only if consciousness is related to underlying physiological events as a process relates to the elements of the physical substratum for that process. For example, a transverse wave takes physical particles and discrete movements of these particles as its substratum when the wave passes through that particular material medium (as when a sound wave passes through a wooden door), but the wave is not identical with the door, nor is it caused by the door. The wave could have been the same wave even if some other material medium (a different door, or a volume of air) had been in a position to serve as its substratum.

Consciousness, like a sound wave, is multiply realizable in this sense; but unlike things like sound waves, it is not empirically observable. Sound waves are completely describable in terms of the observable movements of their substrata, whereas not everything about consciousness is. This feature of consciousness cannot be accounted for sheerly by its multiple realizability. Consciousness, if it is to be empirically unobservable without being nonphysical, must therefore be a special *type* of multiply realizable process.

I shall argue that, to account for the way consciousness differs from other multiply realizable processes, we must grant that it is a self-organizing process, or an aspect of such a self-organizing process. But this will require a somewhat precise understanding of what it means for a process to be self-organizing. A self-organizing process is one that, by virtue of the pattern of its organization, appropriates and replaces physical components needed for the continuation of that pattern of organization into the future. Current theories of self-organizing processes treat them as open thermodynamic systems, called "dynamical systems" (for example, see Kauffman, 1993). The system is "open" because it continually exchanges energy and materials with its environment; otherwise it would be a dead system. And it is "dynamical" because, although it continually replaces its own material substrata, it maintains the continuity of its pattern of organization, which is definitive of its form of being.

I shall further argue that only certain kinds of self-organizing processes are capable of consciousness. In order to be capable of consciousness, a process must not only appropriate and replace its own substrata, but it must do so in ways that are motivated by emotional interests in interaction with the environment. The reason for this is that the mere impression of sensory data into the material of an organism does not by itself result in consciousness of those data. Only when we actively look for information in a motivated way does the information enter into consciousness — even if the only motivation for

the conscious attention is a generalized feeling of curiosity. According to Panksepp (1998), a generalized curiosity is one of the most important emotional systems motivating animals to search their environments. Without actively turning our attention toward a stimulus and looking for it, we would not see it. That is why, when tracking the movement of a soccer ball, we can be tricked by a deceptive leg movement. We look for the ball in one place, and if it goes in an unexpected direction, we must first reorient ourselves to look for it in that direction before it can enter our consciousness, *even if it was in our visual field the entire time*. Perceptual experiments by Mack and Rock (1998) have now demonstrated systematically that even when a datum appears at or near the focal point in a subject's perceptual field, the subject is not conscious of that datum if already in the process of attending to a different visual task either there or elsewhere in the visual field. Add to this fact Panksepp's thesis that all attention must be motivated (even if only by general curiosity), and it follows that only emotionally motivated organisms, structured by the requirements of self-organization, are capable of consciousness.

What Is Emotion?

To argue that consciousness requires a particular kind of self-organization that includes emotional motivations requires that these terms be defined with a certain degree of precision. On the one hand, we need to distinguish emotional and motivational phenomena from the mere *tendencies* that nonconscious and even sometimes inorganic self-organizing systems exhibit. But on the other hand, we don't want to rule out a priori the possibility of unconscious emotions. Also, it is useful to distinguish "motivations" from even *non-conscious* "emotions," since we sometimes want to say that the organism "is motivated" to maintain this or that equilibrium, but without identifying these motivations a priori with any specific "emotions."

These terms should therefore be defined so as to leave open four different possibilities: (1) unified *tendencies* in systems that have neither emotions nor motivations; (2) motivations that are not associated with any emotions; (3) motivations that are associated with unconscious emotions; and (4) motivations that are associated with conscious emotions.

Even self-organizing systems that do not act as agents of their own actions can have *tendencies* to sustain themselves by appropriating their own needed material substrata, but these tendencies do not constitute *intentions* — intentional motivations and emotions. This lack of "intentionality" correlates with the organism's lack of representations (perceptual and imaginative) that could be associated with any *aims* or *objects* to which its behavioral tendencies could relate. For example, plants seem to fall into this category (see Faw, in press).

By contrast, there are other organisms that do act as agents of unified actions in the sense that their systems can execute plans of action in relation to perceived objects, for example, to move from point A to point B if there is more food at B; but this ability does not necessarily entail that such organisms are conscious of intending to perform the actions for these reasons. Nonetheless, at this level, some sort of rudimentary representation or perception is associated with the motivated action, even if not a conscious representation (Newton, 1996). For example, insects, and perhaps even frogs, might fall into this category (Faw, in press; Merleau–Ponty, 1942/1963). A colorful example is the snake, which is cold-blooded both literally and metaphorically: it perceives, but between the perception and the action there is no need for evaluation; the snake is instinctually predetermined to respond in a stereotyped way when the object is presented.

If there is a difference between motivation and emotion, it must be that the emotion has the *capability* of being experienced with certain feeling tones or qualitative conscious properties, even in those cases where the emotion does not *actually* enter into consciousness. We should allow at least for the possibility of unconscious emotion. Humans, for example, seem sometimes to have emotions and yet be completely unconscious of what the emotions are really "about." We may be unaware of the aims and objects of an emotion, or even mistaken as to what the quality of the emotion really is. I may think I am "angry" about my son's not taking out the garbage, but it is not the case that if he had done it then I would feel fine; what may really be happening is that I am "frustrated," not angry — and not because of what my son did, but about the way my career or love life is going. So we want to leave room for *unconscious* emotions in this sense, and these are very important in understanding motivation, since more often than not our first kneejerk reaction about what our emotions really want us to *do* (their aims in relation to their objects) are very inaccurate, and to get more accurate requires getting more clear on what we really feel and what the feeling really wants us to do about it.

We then need to clarify why there is a difference between an unconscious *motivation* and an unconscious *emotion*, in some such way as this: a motivation that is incapable of becoming conscious in the particular organism in which it occurs is not an "unconscious emotion." A motivation can be an emotion only in an organism that, in general, is capable of being conscious that it feels so-and-so. This requires (a) that the organism be able to form proprioceptive and exteroceptive imagery; and (b) that the particular motivation is of a kind that can be proprioceptively sensed as related to exteroceptive and/or other proprioceptive imagery for that organism. We would not say that a plant or even an insect has even "unconscious" *emotions*, although the insect can have *motivations* (whereas the plant cannot). And we would not say that I have an "unconscious emotion" that wants to regulate my

blood sugar level, even though I am *motivated* to do so (because my system acts as a whole to regulate it, and will readjust its different systems and its behavioral activity as a whole in relation to perceived environmental conditions such as food, if that is needed to facilitate the process). But, if I am frustrated and do not know it or cannot feel it, we *do* want to say that I have an "unconscious emotion" of frustration. Those who are averse to speaking of emotions as unconscious might prefer the term "preconscious," since the latter term implies that the emotion *could* become conscious if we were to engage in some kind of conscious process aimed to bring it into awareness. But it is still important here not to say that the conscious emotion is just the unconscious emotion *plus* a direction of conscious attention *to* it, because when we become conscious of something, we execute a vast complexity of neurophysiological processes, and this will undoubtedly *change* the physiological structure of the emotion itself. This is important for avoiding what Natsoulas (1993) calls an "appendage theory" of consciousness. Consciousness is not just an extra layer that is tacked on to processes that could have occurred on an unconscious basis without their basic structure's being affected by whether they are conscious or unconscious.

One further terminological problem is that some people reserve "emotion" to refer only to *instinctual and hardwired* motivational affective feelings, and not to motivational affective feelings in which the aims and objects are not hardwired; whereas other people just lump them all under the term "emotion." Musicians routinely talk about "emotional expression" in music, but the "emotions" in question are hardly ever, and maybe never, the instinctual, hardwired ones. Philosophers also are prone to use "emotion" in the broader sense, as in A.J. Ayer's "emotivist theory of ethics." In this paper, I am using "emotion" to refer to the broader class — all kinds of affective and motivational feelings, both conscious and unconscious. Within this category, of course, different kinds of affective and motivational feelings can be identified, and if someone wants to reserve "emotion" for the more stereotyped ones, and use some other term for the more complex ones, that is merely a terminological convention; I have argued elsewhere (Ellis, 1995, 1996, in press) that such a rigid distinction cannot be maintained, and that all emotion involves complex intentionality, but there is no need to resolve that for the purposes of this paper.

An emotionally motivated self-organizing system, then, is a particular type of process–substratum relation. In any process–substratum relation, the process is not reducible to the sum of the behaviors of its substrata, because the same process could have occurred using different substrata, as long as those different substrata had interrelated with the same overall pattern; and the process could have failed to occur, given those same substrata, if the substrata had related to each other in different patterns. If the process could have

occurred without those substrata and vice versa, then the process is not equivalent with the substrata. For this reason, a process–substratum relation is the type of relation that supervenience theorists have sought (Kim, 1992, 1993; Searle, 1992; van Gulick, 1992). Those who object to supervenience (Bickle, 1992; Newton, 1996) insist that, if consciousness is to be a property that supervenes on physical things, then it ought to be a physical property as amenable to physical description and observation as any other. But consciousness does not seem describable in terms commensurate with physical parameters, and thus with the idea that it supervenes on physical things.

The process–substratum model provides a way of conceiving of consciousness that is commensurate with the physical realm, but without identifying it with anything that is objectively observable. The process of consciousness, though it takes empirically observable events as substrata, may not be accessible to someone not living *in* the location of the organism that works as substratum for that process. The reason, as suggested above, is that all consciousness is permeated by emotional feelings; we have these feelings only by actively *generating* them out of the total motivational structure of our own organisms.

Emotions, Organismic Purposes, and Self-Organizing Processes

We must now address two important questions: How do self-organizing processes differ from phenomena like temperatures and sound waves, which are indeed multiply realizable but not self-organizing? And how can the fact that consciousness is self-organizing account for its empirical unobservability?

Self-organizing processes differ from other multiply realizable phenomena in this sense: temperature is multiply realizable in that the same temperature could have been realized by an infinite number of different combinations of movements of particles in the substance realizing that temperature. Similarly, a sound wave could have been the same sound wave even if it had travelled through a different medium. How do self-organizing processes differ from these? A self-organizing process, as Monod (1971) and Kauffman (1993) suggest, is one whose *internal structure* is what makes it especially prone to replace its own substrata if needed to maintain that structure. Typical examples of such an internal structure can be found at the cellular level. One example is a structure in which "The enzyme which catalyzes the first reaction of a sequence . . . is inhibited by the final product of the sequence. The intercellular concentration of this metabolite therefore governs its own rate of synthesis" (Monod, 1971, p. 64) Another example is where "The enzyme is activated by a product of degradation of the terminal metabolite." In each case, the internal structure of the process guarantees its strong tendency to be maintained, even across multiply realizable replacements of its own substratum elements.

Monod's concept of self-organization does not entail a causal interactionism: the process does not *cause* the behavior of its substratum elements. The behavior of each substratum element has causal antecedents at the substratum level that are both necessary and sufficient, under the given circumstances, to bring about that behavior. But the self-organization of the organism in which this behavior occurs is partly constitutive *of* the given circumstances under which those antecedents are necessary and sufficient for those consequences. This makes it possible that, if the needed antecedent for a behavior had not been available, the self-organizing organism is structured so that it could have changed some of its other functions in order to allow some *other* antecedent to be used as the necessary and sufficient antecedent of that same behavior. A typical example is the reorganization of brain function in mild stroke recovery. Even though the specific behavior of each substratum element has antecedents at the substratum level which are necessary and sufficient to produce it under the given circumstances, the structure of the self-organizing process as a whole is such that those given circumstances will tend to be changed when that is what is needed to maintain the general contour of the functioning of the overall process as such.

In other multiply realizable processes, such as temperatures or sound waves, it is true that the same temperature or sound wave could have been realized by different substrata if the given circumstances had been different; but only a *self-organizing* process is structured in such a way that it can change those given circumstances in order to ensure the continuity of the overall process. If the causal antecedents of the movement of a gas molecule are not such as to cause the temperature in the chamber to be 98 degrees, the chamber has no built-in structural tendency to change some other aspect of itself to ensure the maintenance of the 98 degree temperature; self-organizing organisms do this all the time.

Monod, like Kauffman, notes the similarity of living processes to crystal formations in this respect. Forming crystals are like growing biological organisms in that a property of their very *structure* as such is that additional materials coming into their vicinity have a tendency to enter into patterns that reflect that initial structure. But, unlike living organisms, crystals are not fluid or flexible enough to *change* aspects of themselves if needed to maintain continuity in some other aspect. If a particular element's causal antecedents had not been necessary and sufficient to produce the pattern under the given circumstances, the crystal as a whole would not have *changed* the given circumstances. This is just what *is* done by biological organisms.

A specific motivation, say the desire to raise my hand, results ultimately from the organism's self-organizing tendency. This self-organizing structure is therefore present, and embodies a tendency for me to want to raise the hand, even before the desire becomes pronounced enough to be a conscious aware-

ness. Thus the "expectancy wave" accompanying the decision to raise the hand is measurable *before* I am aware of a desire to raise it (Libet, Curtis, Wright, and Pearl, 1983; Young, 1988, pp. 164ff). Expectancy waves (also manifested in Libet's readiness potential) indicate that motivational feelings arise out of the organism's generally self-organizing nature.

Monod's concept of self-organization is essentially the same as Merleau–Ponty's (1942/1963) concept of a *purposeful* organism — in which a change in one part can be compensated for by changes in other parts where needed to maintain the continuity of the whole. Merleau–Ponty thus defines his notion of what it means to be a living being in terms of the idea that a part tends to be subordinated to the *purposes* of the whole.

If self-organization is the same as purposefulness, and if emotional motivations are purpose-directed activities, it follows that only self-organizing processes can be characterized by emotional motivations. (Not *all* self-organizing processes are characterized by emotional motivations, of course, but just the converse.) And if there can be no consciousness without some emotionally motivated direction of attention, then it follows that only self-organizing processes can be conscious. (Here again, I am not claiming that *all* self-organizing processes are conscious.) Moreover, we have seen that the empirical unobservability of consciousness stems from the fact that to experience a state of consciousness entails generating the emotional motivations that are a crucial part of the constitution of conscious states, and this in turn requires *being* the organism that generates those emotional motivations. Thus the reason why the phenomenal character of a state of consciousness cannot be inferred purely from observations made from an external or empirical standpoint is that emotions are motivated *actions* which an organism performs, and to experience them is to perform them. A scientist observing a subject's brain, in order to be conscious of the observation, must motivatedly perform the emotions *in herself* that direct *her* conscious attention; but by doing so she does not experience the *subject's* emotions, but rather her own. Therefore her empirical observations do not yield knowledge of what the subject's consciousness is like, although they *do* yield knowledge of what her *own* consciousness is like. The complex self-organizing process constitutive of the emotional motivations needed for the subject's phenomenal consciousness are experientially accessible only from the standpoint of the organism that executes them, because conscious experiencing per se entails executing rather than merely observing emotional processes.

References

Aurell, C.G. (1989). Man's triune conscious mind. *Perceptual and Motor Skills, 68,* 747–754.

Bechtel, W., and Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science, 66,* 175–207.

Bickle, J. (1992). Multiple realizability and psychophysical reduction. *Behavior and Philosophy, 20,* 47–58.

Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2,* 200–219.

Damasio, A. (1994). *Descartes' error.* New York: Putnam.

Ellis, R.D. (1995). *Questioning consciousness: The interplay of imagery, cognition and emotion in the human brain.* Amsterdam: John Benjamins.

Ellis, R.D. (1996). *Eros in a narcissistic culture.* Dordrecht: Kluwer.

Ellis, R.D. (1999). The dance form of the eyes: What cognitive science can learn from art. *Journal of Consciousness Studies, 6,* 161–175.

Ellis, R.D., and Newton, N. (Eds.). (in press). *The caldron of consciousness: Affect, motivation, and self-organization.* Amsterdam: John Benjamins.

Faw, B. (in press). Consciousness, motivation, and emotion: Biopsychological reflections. In R.D. Ellis and N. Newton (Eds.), *The caldron of consciousness: Affect, motivation, and self-organization.* Amsterdam: John Benjamins.

Jackendoff, R. (1996). How language helps us think. *Pragmatics & Cognition, 4,* 1–34.

Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy, 83,* 291–295.

James, W. (1968). Does "consciousness" exist? In J. McDermott (Ed.), *The writings of William James* (pp. 169–184). New York: Random House. (Originally published 1908)

Kauffman, S. (1993). *The origins of order.* Oxford: Oxford University Press.

Kim, J. (1992). Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research, 52,* 1–26.

Kim, J. (1993). The nonreductivist's troubles with mental causation. In J. Heil and A. Mele (Eds.), *Mental causation* (pp. 189–210). Oxford: Oxford University Press.

Libet, B., Curtis, A.G., Wright, E.W., and Pearl, D.K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain, 106,* 640–642.

Mack, A., and Rock, I. (1998). *Inattentional blindness.* Cambridge, Massachusetts: MIT/Bradford.

Merleau–Ponty, M. (1963). *The structure of behavior* [A. Fischer, Trans.]. Boston: Beacon. (Originally published 1942)

Monod, J. (1971). *Chance and necessity.* New York: Random House.

Natsoulas, T. (1993). What is wrong with appendage theory of consciousness. *Philosophical Psychology, 6,* 137–154.

Newton, N. (1996). *Foundations of understanding.* Amsterdam: John Benjamins.

Noren, S. (1979). Anomalous monism, events and "the mental." *Philosophy and Phenomenological Research, 40,* 64–70.

Panksepp, J. (1998). *Affective neuroscience.* New York: Oxford University Press.

Popper, K., and Eccles, J. (1977). *The self and its brain.* Berlin: Springer Verlag.

Posner, M.I., and Rothbart, M.K. (1992). Attentional mechanisms and conscious experience. In A.D. Milner and M.D. Rugg (Eds.), *The neuropsychology of consciousness* (pp. 91–111). London: Academic Press.

Putnam, H. (1993, December 28). *Functionalism.* Paper presented at meeting of the American Philosophical Association, Atlanta, Georgia.

Robinson, H. (1982). *Matter and sense.* New York: Cambridge University Press.

Searle, J. (1984). *Minds, brains and science.* Cambridge: Harvard University Press.

Searle, J. (1992). *The rediscovery of the mind.* Cambridge: MIT Press.

van Gulick, R. (1992). Nonreductive materialism and the nature of intertheoretic constraint. In A. Beckermann, H. Flohr, and J. Kim (Eds.), *Emergence or reduction? Essays on the prospects of nonreductive materialism* (pp. 157–178). Berlin: Walter de Gruyter.

Young, J.Z. (1988). *Philosophy and the brain.* Oxford: Oxford University Press.