

Our Overt Behavior Makes Us Human

Howard Rachlin
Stony Brook University

The four commentaries all make excellent points; they are all fair and serve to complement the target article. Because they are also quite diverse, it makes more sense to respond to them individually rather than topically.

McDowell

Before discussing McDowell's (2012) thoughtful comments, I need to clarify his categorization of my position on consciousness as "eliminative materialism." He is correct that I would eliminate the *phenomenology* of consciousness from scientific discourse. However, I also claim that the concept of consciousness itself is extremely useful and has an important place in behavior analysis. So I would not eliminate the *concept* of consciousness from scientific discourse. The theory of consciousness implied by Watson II is a physical theory, like the neural identity theories to which McDowell refers. However, neural identity theorists believe that consciousness occurs within the organism and is identical to some pattern of nervous behavior. I claim that consciousness occurs in the world outside the organism and is identical to abstract patterns of overt behavior. The difference between my identity theory and theirs is not one of physical versus mental; we agree that the mental is real, and it is identical to an abstract pattern of activity of the organism. The difference is that, for them, the pattern occurs (wholly or mostly) over some spatial extent in the brain, whereas for me, the pattern occurs over time in the organism's overt behavior. It is not

the word *consciousness* that I would eliminate from scientific discourse and still less from everyday speech. Contrary to what McDowell says, I *do* "acknowledge the existence and reality of consciousness" (p. 25). Abstract entities, such as behavioral patterns, are as real as or more real than their components.¹ It is rather phenomenological introspection or internal "reflection" as a means of psychological investigation that I would eliminate. I recognize the importance of a kind of reflection (contingencies of reinforcement are essentially reflections from overt behavior to the world and back) but not a reflection that takes place wholly within the organism. Introspection, as a psychological technique, has been tried for at least a century and has produced little of value.²

One argument I take very seriously is that my view of the mind is bad for behavior analysis. But I cannot abandon that view because nonbehaviorists or antibehaviorists like John Searle are not able to understand why I have it. The history of science is full of *prima facie* facts that have been proven to be less useful than their contraries. Especially suspicious are those facts that put humans at the center of the universe (physical or spiritual). The sorts of existence postulated by the phenomenologists arguably come under this heading. From a pragmatic viewpoint (my

Address correspondence to Howard Rachlin, Psychology Department, Stony Brook University, Stony Brook, New York 11794 (e-mail: howard.rachlin@sunysb.edu).

¹Both Plato and Aristotle believed that abstract entities may be in a sense more real (because they are directly connected to their function) than their components. For Aristotle, a chair is more real than the parts that make it up, and for Plato, the user of the chair knows the chair better than does its maker, again because the user is directly involved in its function as a chair (Rachlin, 1994).

viewpoint), something is true because it is useful in the long run to behave as if it were true. The burden is on us behaviorists to show that our account is more useful than others. Once that happens, what seems obvious will change accordingly. Searle's objection, quoted by McDowell, rests on the implicit premise that what Searle cannot imagine or understand must be false. If the research based on teleological behaviorism by me and others turns out to be unfruitful or useless, then such objections will have weight. It is perhaps fair to say that there has not yet been enough research on behavioral patterns, or acceptance and understanding even within behavior analysis, to give teleological behaviorism a fair test. One purpose of the target article is to correct this lack. Meanwhile, I will have to take my chances with Searle. He may be beyond convincing, but hopefully not every philosopher is that closed minded. McDowell and others have reached across disciplines to make contact with philosophers and neuroscientists, and that gives one hope. If teleological behaviorism does not result in an improved behavioral technology then that is why it will fail; not because it contradicts a philosopher's entirely subjective certitudes.

² Nevertheless, introspection may be useful in everyday life. I may say, "I am angry," or "I love you," but not merely to report an internal state any more than I would say, "The grass is green," or "The sky is blue," merely to report an external state. Any statement must be made for a reason. The reason, in the case of "I am angry," and so on, is to predict one's own future behavior on the basis of one's own past behavior in similar circumstances. Such a prediction enables the hearer (it could be just one's own self) to react appropriately. A person (who is less observant of his or her own behavior than someone close to him or her) may be wrong about an introspective statement. I might say, "I am angry," and truly believe it, and my wife may say "No you're not," and she may be right. It is introspection as a scientific method, not introspection as a useful kind of everyday behavior, to which I object.

McDowell's summary of the views of Brentano, Husserl, and Sartre is interesting and enlightening. There is certainly a commonality between behaviorism and their philosophy, perhaps coming to a head in the later Wittgenstein (1958) who said, "If one sees the behavior of a living thing, one sees its soul" (p. 357). More relevant to the current topic is McDowell's discussion of the modern philosophers, John Searle, Thomas Nagel, and Colin McGinn. It seems to me that, at least as McDowell presents their views, all three are dancing around the mind-body problem and coming no closer to solving it than did the European philosophers of the 18th or 19th centuries. But modern philosophy is not as negative about behavioristic thought (or, more aptly, not as positive about phenomenology) as McDowell implies. According to Noë (2009),

After decades of concerted effort on the part of neuroscientists, psychologists, and philosophers, only one proposition about how the brain makes us conscious—how it gives rise to sensation, feeling, subjectivity—has emerged unchallenged: we don't have a clue. (p. xi)
 Consciousness is not something that happens inside us. It is something we do or make. Better: it is something we achieve. Consciousness is more like dancing [overt behavior] than it is like digestion [covert behavior]. ... The idea that the only genuinely scientific study of consciousness would be one that identifies consciousness with events in the nervous system is a bit of outdated reductionism. (p. xii)

Searle, as quoted by McDowell (p. 21), claims that "neural activity and conscious experience are different aspects, or levels of description, of the same thing, in the same way that, say, the molecular structure of a piston and the solidity of the piston are different aspects, or levels of description, of a piston." Amazingly, Searle has it almost right. Substitute *behavioral activity* (overt) for *neural activity* (covert), and I would completely agree. But Searle, despite his intention to rid philosophy of Cartesian remnants, has not completely

eliminated Cartesian dualism from his own philosophy. If mental (or conscious) activity is an abstract version of physical activity, what is that physical activity? Why is it any more plausible for Searle, and the many philosophers who have considered this question, that conscious physical activity has to occur inside the head than that it occur in overt behavior? I understand why Descartes saw things this way. Because Descartes believed that the soul was located deep in the brain and the physical motions had to directly influence the soul, and vice versa, the physical motions also had to be in the brain. But Searle presumably does not believe that there is a nonphysical soul located deep within the brain that interacts with our nerves. Nor, as McDowell points out, is this inherently obvious. Some societies and some ancient philosophers believed that our minds as well as our souls were in our hearts. I would guess that if you name a vital organ, there will be or have been some society that believed it to be the seat of the soul; there may even have been some who identified the soul with the whole organism. So if the mind is a molar or abstract conception of some physical activity (as Searle and I seem to agree), and there is no *a priori* reason (such as connectivity with an internal, non-physical soul) to assume that the physical activity occurs in the brain, where does it occur?

In answering this question, usefulness is paramount, especially as consciousness, and talk of consciousness, must have evolved along with the rest of our human qualities. Organisms may die without reproducing because their behavior is maladaptive not directly because their nerves are maladaptive. Our nerves would be in direct contact with our souls if our souls, as the sources of consciousness, were inside of us. But if our environment is seen as the source of our consciousness (as

it would have to be if consciousness were a product of biological evolution), then it would be our overt behavior, not neural behavior, that is in direct contact with the source. Group selection (selection at the level of classes or patterns) may act at the level of nervous function, as Edelman and colleagues (e.g., Tononi & Edelman, 1998) have shown. It may act as well at the level of innate behavioral patterns across generations (Wilson & Wilson, 2008). And it may act as well at the level of learned patterns within the lifetime of a single organism (Rachlin & Locey, 2011).

Consciousness is therefore not an epiphenomenon or a faint halo that wafts up from a certain degree of complexity in our nervous systems, but is a vital property of our overt behavior with a vital function in our complex world. Our long-term patterns of behavior, including sobriety, moderation, cooperation with others, morality, rationality, as well as the language that reflects (and at the same time imposes) their organization, all evolved. These patterns are what we would have to create in Watson II for him to leap over those eons of biological evolution and be human. The mechanism that could create those patterns may very well turn out to resemble our actual nervous mechanism. Or it may not. But it is behavioral evolution, not neural evolution, that counts for Watson II's consciousness.

Searle, Nagel, and McGinn, as presented by McDowell, all have double-aspect theories of mind: Body and mind are two aspects of the same thing. The traditional question to ask two-aspect theorists is: Two aspects of what? Searle gives the correct answer: The body is to the mind as the molecular ("molecular structure of a piston") is to the molar ("solidity of a piston"). This is a spatial analogy but it could just as well be a temporal one: ... as the notes are to the melody; ... as the steps are to the dance. But Nagel and McGinn both

posit a third entity that the two aspects are aspects *of*. For Nagel it is Factor X and for McGinn it is “unknowable.” Are these answers to the traditional question any more enlightening than the traditional answer to that question: two aspects of God? I do not believe so.

A view of consciousness proposed by Noë (2009) holds (as I do) that the mind cannot be understood except in terms of the interaction of a whole organism with the external environment. Nevertheless, for Noë, the brain remains an important component of mental activity. He retains a neurocognitive view of the mind while he expands its reach, beyond the brain, into the peripheral nervous system and the external environment. According to Noë, “My consciousness now—with all its particular quality for me now—depends not *only* on what is happening in my brain but *also* on my history and my current position and interaction with the wider world” (p. 4, italics added).

I believe that this is a step in the right direction, but its problem is that it mixes levels of explanation. Consider the following transcription of Searle’s distinction between physical activity and conscious experience: “[Behavioral] activity and conscious experience are different aspects, or levels of description, of the same thing, in the same way that, say, the molecular structure of a piston and the solidity of the piston are different aspects, or levels of description, of a piston.” If conscious experience is analogous to the solidity of the piston, then it cannot *also* be analogous to its molecular structure. Noë’s conception of conscious activity blurs the distinction between conscious and nonconscious activity. Extended cognition theory extends the domain of consciousness spatially beyond the brain, into the peripheral nervous system and out into the world. But it does not consider a temporally extended view of cognition that

extends behavior beyond the present moment into the past and future. It is this temporal extension, I believe, that gives Watson II his humanity.

Finally, McDowell proposes a mental rotation test and a visual oddity test as possible alternatives to the tough Turing test I proposed in the target article. The problem with these alternatives is that it would be extremely easy to build a machine that would pass these tests with flying colors. I believe the current Watson, with a little tweaking, could easily do it. Suppose Watson did pass these tests but failed the tough Turing test. Would anyone believe that it was human? Suppose Watson passed the tough Turing test (for sensation, perception, imagination, cognition, as well as the emotions of love, anger, hope, fear, etc.), but failed the mental rotation and visual oddity tests. Would it not be a violation of our common morality not to consider it human?

Schlinger

Schlinger (2012) claims that Watson “would be handicapped in that he would have no private world to experience and, thus, to describe” (p. 43). But he also agrees with me that “consciousness is in the behavior, not the mechanism” (p. 42). The question I would like to address in this reply is: Do covert talking and covert picturing properly belong to the class of movements we call *behavior* or are they themselves, like the chemical and electrical events involved in neural transmission, part of a *mechanism* that underlies behavior? If the latter, then by Schlinger’s own reasoning, Watson’s private world would be irrelevant to whether or not he could be conscious; we would then have to look, as I do in the target article, for Watson’s and our own consciousness in our overt rather than covert behavior.

The nub of Schlinger’s views is best captured by the following passage:

A ... circumstance that probably evokes the term *conscious* most often, and the one that is of most interest to consciousness scholars and laypeople alike, is the tendency to talk ... to ourselves about both our external and internal environments, and our own public and private behavior. ... It is these behaviors that give rise to what consciousness scholars refer to as *qualia* or subjective experience and consist of what I believe a conscious person behaves like. That is, a conscious person is taught by his or her verbal community to answer questions about his or her own behavior, such as "What are you doing?" "Why did you do that?" and "What, or how, are you feeling?" ... As a result, we are constantly describing our behavior and private events. (p. 42)

Let us start from the back of this statement. Why does our verbal community want to know what we are doing, how we are feeling, why we do this or that? What's in it for them to know? Or more precisely, what reinforces these requests of theirs? The answer is that we are interacting with our verbal community in a social system, our future behavior affects their welfare, and they would like to be able to predict better than they currently can what our future behavior will be.³ So when we answer their questions, we are essentially making predictions about our future behavior. Now let us consider the reverse question: Why should we bother to answer these questions? Why should we bother to make such predictions? The answer again must be that the questioners are interacting with us in a social system; their future behavior affects our welfare, and we are trying as best we can to maximize the value to us of their behavior both now and in the future. In other words, we are engaged with them in a joint venture, and it is to our interests to refine the flow of discriminative stimuli back

³There may of course be other reasons. It may be idle curiosity. Or the questioner might be a neighbor in the elevator saying "How are you?" and I answer "fine," even if I happen to be rushing to the doctor. But I think that the reasons for such interchanges, like the reasons for those Schlinger cites, are reducible to a mutual interest in greasing the wheels of our current and future interactions.

and forth between us and them. Schlinger may agree so far.

Now let us consider to what we may refer when we answer their questions. We could be referring, as Descartes believed, to a spiritual state, a state in a nonphysical world with its own rules, located somewhere inside of us (perhaps in our pineal glands), to which our questioners have no access but to which we have direct and unimpeachable access through introspection. Or we could be referring to a state of our nervous systems (the chemicals and electrons running through our nerves) or to a kind of organization of those chemicals and electrons in which they mimic the executive function of a computer program. I assume that Schlinger agrees with me that such neurocognitive events are interesting and valuable objects of study but are mechanisms rather than behaviors and are not what we refer to when we answer questions such as "How are you feeling?" (Moreover, why, unless they are neurologists, should other people be interested in the state of our nervous systems?)

Or, when we answer such questions, we could be referring to what we say to ourselves. According to this scenario, if my wife asks me, "What did you think of those people we met for dinner last night?" and I say, "I think they were a pair of creeps," I must actually be referring not to the people themselves, nor to their actual behavior, nor to my interaction with them, but to some sentences I was saying to myself or some image of them undetectable (to my wife) that I created in my muscles between her question and my answer. But even that implausible scenario would not be getting at my consciousness. According to Schlinger, it is not the covert words or images that constitute my consciousness but my proprioceptive feedback from these words and images. Schlinger claims that covert behavior "give[s] rise" (p. 42) to consciousness and "Without a

sensory system that, in addition to exteroception, also includes interoception or proprioception, Watson II would not be able to describe private stimulation or, in other words, how he feels" (p. 42). But, aren't interoception and proprioception chemical and electrical events in our nerves? You can't have it both ways. Covert movements cannot just "give rise" to consciousness; if they are to explain consciousness, they must be consciousness itself. And, if covert behavior is consciousness itself, consciousness cannot also be the *perception* of covert behavior. But let us suppose for a moment that consciousness is perception of internal speech by our proprioceptive nervous system. What exactly would that perception be? Is it identical to the entirely physical activity in our proprioceptive nerves? Or, do we need a still more covert activity (the perception of the perception) to explain the perception? And so on until we get to the center of the brain where the only remaining possibility is a nonphysical soul, and we are back to Descartes' model. Moreover, what a waste it seems for such an important functional property as consciousness to have evolved to rely on the relatively impoverished proprioceptive system when our exteroceptive system is so exquisitely accurate. It is our past behavior (our reinforcement history) that best predicts our future behavior. If, as I claim, the purpose of answering Schlinger's questions is to predict our overt behavior, the part of our behavior that will affect them, why would our answer refer to our unreliable inner speech? There is no question that we do talk and picture things to ourselves. I believe that these covert acts, when they occur, are part of the *mechanism* by which our overt behavior is sometimes organized. But I do not believe that they can be usefully identified as thinking, perceiving, sensing, imagining, and so on. There is insufficient room between our central and peripheral nervous systems on the one

hand and our overt behavior on the other for a massive covert behavioral system, a system that, if the covert-behavior view of consciousness is right, would have to be the referent for our entire mental vocabulary.

In the face of this unlikelihood, bordering on impossibility, what is a behaviorist to do? One tactic would be for behaviorists to join many philosophers and declare that the mind is simply inaccessible to scientific study. Such an attitude is understandable coming from philosophers, because by implication *they* would be the experts on mental life. But, for a psychologist, to give up on the scientific study of the mind and consciousness is to give up on what psychology is supposed, by the people who support our research, to be all about. Such a tactic, if adopted, would marginalize behaviorism still further within psychology. But these are just extrinsic reasons. The intrinsic reason for a behavioral science of mind, the reason I wrote the target article, is that a view of the mind as overt behavior is the best, the most logically consistent, the most satisfying (try it and see) view of the mind that one can take.

To take this view, however, we need to give up on the strict efficient-cause, mechanical, interacting billiard-ball view of causation in which each cause must lie temporally as well as spatially up against its effect, and to adopt a teleological view of causation. From a teleological viewpoint, abstract patterns of movements are final causes of the particular acts that make them up. Instead of efficient causes prior to their effects, final causes are more abstract and extended in time than their effects. For example, fastening a board is a final cause of hammering a nail, building a floor is a final cause of fastening a board, building a house is a final cause of building a floor, sheltering a family is a final cause of building a house, and so on. Each final cause is an answer to the question WHY? Efficient causes are

answers to the question HOW? Thus, final causes are more appropriate than are efficient causes for Skinnerian behaviorists who are focused on explaining behavior in terms of reinforcement. Skinner's notion, that a contingency of reinforcement (that takes time to occur) can be a cause, and that a response rate (that takes time to occur) can be an effect, is an example of departure from efficient causation. We do not need to justify the effect of contingencies by imagining miniature contingencies represented in the brain efficiently causing behavior. Physics long ago gave up the billiard-ball view of the connection between cause and effect (gravity, magnetism, electric fields), not to mention all of quantum physics. In economics, utility functions are viewed as causes of the particular economic exchanges that make them up. A utility function need not be represented in the brain or anywhere except in the economist's observations. Aristotle believed that final causes are actually more scientific than efficient causes because they are more abstract (Randall, 1960). In the target article I tried to demonstrate that our mental vocabulary fits like a glove on patterns of overt behavior over time. It is in that teleological sense and in that sense only that, as Aristotle claimed, the mind can cause behavior (Rachlin, 1992, 1994).

Hutchison

As Plato claimed, the knowledge of the user of a chair is prior to that of the builder of a chair. The *function* of the chair, the reason for the chair's existence, is part of the user's knowledge (or that of the builder as user). Before the builder can even begin to work, he or she must first know the chair as a user knows it. The same is true for the builder and user of an automobile and the builder and user of a robot. But the reverse is not the case. To use a chair properly it is not necessary to know how to operate a

lathe; to drive a car properly it is not necessary to know how to forge the steel of the chassis; to use a humanoid robot properly it is not necessary to know how to build its program. We behavior analysts are like the users of human beings. Our knowledge leads to prediction and control of the behavior of whole organisms (including our own behavior). The builder of a humanoid robot needs to know what behavior analysts know, but the behavior analyst does not need to know what the roboticist knows. The behavior analyst can go about his or her business without knowing anything about how to build a human robot (as interesting and valuable as that knowledge certainly is), but the reverse is not the case. I do not believe that anything in Hutchison's (2012) highly interesting and informative comments contradicts this assertion, but it is worth making it explicit.

Consistent with the above, if learning works like evolution, it is not because minievolutionary systems have been found in our brains (again, as interesting and valuable as that finding may be). Rather, it is because our *behavior* as whole human beings evolves over our lifetimes (behavioral evolution) according to the same selective principles as individuals evolve across generations (biological evolution). As the target article indicates, group selection (in which the replacement of unfit individuals within a group is relatively slower than replacement of unfit groups within a population) is much more likely on the behavioral than on the biological level. This group selection process (in which whole behavioral patterns evolve as such) is responsible for the learning of self-control, social cooperation, and altruism (Rachlin & Locey, 2011), which are behavioral attributes necessary for Watson's behavior to approach that of an evolved human. (All biological selection is in a sense group selection. Just as tribes are groups of families, and

families are groups of individuals, individual organisms are groups of organs, etc.; Wilson & Wilson, 2008.)

A cautionary note: As useful as delay-discounting functions are to identify addicts and potential addicts, as central to our personalities as are their slopes relative to the slopes of other individuals (Odum, 2011), I do not believe that discounting functions are encoded anywhere in our nervous systems, and I believe that it would be a mistake to program them as such in robots. Discounting functions vary too much from time to time and from commodity to commodity (e.g., automobiles are discounted differently from money, and money is discounted differently from food). Delay-discounting functions are measures of self-control, not determiners of self-control. What are they measures of? I believe that they are measures of the breadth of overt behavioral patterns. It is the capacity for such breadth to evolve (by group selection of behaviors over a lifetime) that should be programmed into robots if we want to make them behave more like humans. Even if there were delay-discounting functions inside our heads, we would then have to ask how they were formed and how the results were converted into actual behavior.

One more note. Hutchison identifies talking to oneself with thinking (p. 32). As I argued in my responses to McDowell and Schlinger, although talking to oneself may be incidental to thinking, it is neither necessary nor sufficient to talk to oneself in order to think, be conscious, or engage in any mental activity.

Wojcik and Chemero

I believe that the experiments on Wojcik and Chemero's (2012) radical embodied cognitive science, as described in this comment, are good examples of molar extensions of behavior in time as well as in space. Not only are they not objectionable, they point in the direction that one

hopes cognitive psychology will take in the future. I was unaware of this work; I apologize to the authors for my ignorance.

With regard to that new direction I have one reservation; perhaps it is just a matter of semantics. The commenters, in an ongoing study of what they call *extended conscious experience*, measure changes in conscious experience in terms of physiological measures such as heart rate, respiration rate, and galvanic skin response during an extended cognitive game. These are important variables to measure and may yield interesting new information. But it is a testament to the degree with which the Cartesian model is ingrained in our thinking to suppose that by measuring these variables one is getting closer to the participants' actual conscious experience. In taking these physiological measurements, the authors exemplify the direction taken by much modern cognitive research: First find or explore some cognitive activity by means of purely behavioral measures; then explore the physiological correlates of that activity by means of an MRI machine or other physiological technique. As I have said in my previous responses in this exchange, I have no objection to such research. It is important, perhaps vitally important, to understand the physiological correlates of and the mechanism that underlies cognition. But I do object to the presumption that this line of research will lead to better understanding of the mind or consciousness or even such presumably basic processes as sensation.

Another direction in which one might go in studying such processes would be to expand rather than to narrow the scope of the behavior being measured. One might extend the research described in this comment, for example, by developing an economic model of attention in which the tradeoff between performance on the two simultaneous tasks (pushing the cursor and counting backward)

was studied as a function of the properties of those tasks individually (say, the programmed contingency between cursor and mouse) and the reward (implicit or explicit) for doing each of them well or punishment for doing them poorly. Then the model might be generalized to other competing tasks, and the effect of past experience (i.e., reinforcement history) on attention might be studied. I do not say that my direction is better than that taken by most cognitive psychologists, just that my suggested direction is towards consciousness and theirs is away from it.

Let us consider a person sitting quietly and listening to a Mozart quartet. You may study a person's behavior (or the person him- or herself may study it) as a function of or in the presence of musical notes, of melodies or passages of music, of movements, of quartets, of Mozart's works, of classical music, of music. Some of your descriptions may be made with precision (albeit probabilistic precision) on the basis of relatively brief observations; some will require many extended observations and highly abstract terms. As your descriptions progress from the particular to abstract, they will be getting closer and closer to that individual's consciousness. And this would be as true for the person's observations of his or her own

behavior (and these too are overt behavioral patterns) as it would be of yours.

REFERENCES

- Hutchison, W. R. (2012). The central role for behavior analysis in modern robotics, and vice versa. *The Behavior Analyst, 35*, 29–35.
- McDowell, J. J. (2012). Minding Rachlin's eliminative materialism. *The Behavior Analyst, 35*, 17–27.
- Noë, A. (2009). *Out of our heads: Why you are not your brain, and other lessons from the biology of consciousness*. New York: Hill and Wang.
- Odum, A. L. (2011). Delay discounting: I'm a k, you're a k. *Journal of the Experimental Analysis of Behavior, 96*, 427–439.
- Rachlin, H. (1992). Teleological behaviorism. *American Psychologist, 47*, 1371–1382.
- Rachlin, H. (1994). *Behavior and mind: The roots of modern psychology*. New York: Oxford University Press.
- Rachlin, H., & Locey, M. L. (2011). A behavioral analysis of altruism. *Behavioural Processes, 87*, 25–33.
- Randall, J. H., Jr. (1960). *Aristotle*. New York: Columbia University Press.
- Schlinger, H. D., Jr. (2012). What would it be like to be IBM's computer, Watson? *The Behavior Analyst, 35*, 37–44.
- Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science, 282*, 1846–1851.
- Wilson, D. S., & Wilson, E. O. (2008). Evolution "for the good of the group." *American Scientist, 96*, 380–389.
- Wittgenstein, L. (1958). *Philosophical investigations* (3rd ed.). (G. E. M. Anscombe, Trans.). New York: Macmillan.
- Wojcik, K., & Chemero, A. (2012). Nonneurocognitive extended consciousness. *The Behavior Analyst, 35*, 45–48.