

Clancey, W. J. 1999. *Studying the varieties of consciousness: Stories about zombies or the life about us?* *Journal of the Learning Sciences*, 8 (3–4), 525–40.

Review of Chalmers *The Conscious Mind*

The Conscious Mind: In Search Of a Fundamental Theory. David Chalmers. New York, NY: Oxford University Press, Inc., 1996, \$29.95 hardcover. ISBN No. 0195105532. 1997, 432 pp., \$14.95 paper. ISBN No. 0195117891.

William J. Clancey¹
Institute for Research on Learning

Introduction

How difficult it must be to be a philosopher! Talk of God, Zombies, representational content, awareness that isn't experienced—it boggles the mind. All these hypotheticals and definitions—can consciousness be treated like a form of mathematics? Isn't consciousness a phenomenon that can be scientifically described and modeled? What is the role of philosophy? Do Chalmers' analyses, however careful and scholarly, help us build better robots? These are the impressions and questions that swirled about my own conscious mind as I searched for value in Chalmers' book, *The Conscious Mind: Toward a Fundamental Theory*.

My conclusions are mixed: For the scholarship, this account is a useful compendium of alternative theories, with a well-balanced, thoughtful framework for relating and building on different points of view. But as an appraisal of what the collective scientific community now knows about consciousness and for determining what's worth *building* next, this book is not very useful for either the robot designer or neurobiologist. Those interested in the scientific study of consciousness will probably prefer to read (Damasio, 1994; Edelman, 1992; Rosenfield, 1992; Sacks, 1987).

How a non-zombie writes a book review

Before reviewing the contributions and shortcomings of this work, I'd like to frame the discussion by explaining how I think books like Chalmers' should be approached. By the end of my review, this side note will turn out to be a pivotal

¹ Correspondence to: Institute for Research on Learning, 66 Willow Place, Menlo Park, CA 94028. E-mail: bill_clancey@irl.org.

example for understanding why consciousness exists, the central question Chalmers explores.

Recently I received a review of a book manuscript I'd submitted for publication. The reviewer started with a negative attitude, which only got worse as he got caught up in a depressing funk. He found every possible opportunity to haughtily mock the ideas, as if he were slaying some threatening dragon. I learned from this unhappy experience that, first, we must be careful when reading another's work through the filter of our own preferred research approach; second, we must be careful of being dismissive of alternative approaches; and third, we must be generous to allow others to have their say.

In reading Chalmers I sometimes found myself writing "WRONG!" or "WHAT?" in the margins. I looked disparagingly at all his arguments without much data: Why talk about God and Zombies when you can study people or vervets? Why wonder about hypothetical possible worlds when you can wonder instead about the Neanderthal? Why argue about logical supervience when you can model how neurons self-organize? I discovered that my way of reading a large text—like the reviewer who found evidence of theoretical flaws in every typo of my own manuscript—was to interpret like a lawyer-investigator, looking for text strings that could be used to implicate the author for stupidity. I would then present these quotes to my audience, the jury.

But on my second reading of Chalmers' book, after getting back that harsh manuscript review from my own publisher, I realized that Chalmers has another point of view. Unlike me, he's not interested in studying autism, dreams, language, or visual imagination. He's interested in establishing the principles, the natural laws, that hold in a world in which consciousness exists. Yes, he's a philosopher and not a scientist; but we should look for the value in this difference rather than disparaging it. We should acknowledge the obvious, positive character of the book: He's trying to get a difficult, controversial idea across, yet his presentation is not contentious, but unusually broad and considerate.

Starting in a different place than my own investigations, working with different kinds of patterns (Zombies and God), and wrestling with different demons (Materialists and Dualists), Chalmers' story sometimes appears unnecessary and even—harshes attack of all!—unscientific. It may be easy to set

this book aside if you are a neurobiologist or robot builder. But closing down in this way is deciding to take a narrow approach in your own research, missing the opportunity to understand another person's considerate approach to a problem that has not been solved by anyone.

So I tried to calm myself, and reviewed the book, not as a contradiction of all that I knew to be true and right, but as if it were a book from an alien planet. If indeed it had come from outer space, surely I would try to understand it and would be entertained at least by the strange world of this fellow traveler. I would attempt to recognize the value of the perspective, rather than find a hundred ways to mock it.

Here are the questions I then considered with my conclusions summarized:

- What is the approach? What kind of theory is it (An architecture? A mathematical view of system levels in a feedback system? A framework for analyzing facts?)

The book argues that consciousness is a special kind of phenomenon requiring new kinds of physical laws to understand and model.

- What does it contribute? Does it uniquely explain any facts? Does it help us understand the limits of existing models?

Chalmers proposes that we must not only explain how consciousness is possible in physiological terms, but explain why it exists and some sense why it is necessary for this sort of world ("functional criteria for when consciousness arises" p. 229).

- Relative to the alternative approaches of situated robotics or synthetic biology (Beer, 1995; Effken and Shaw, 1992; Steels and Brooks, 1995; Turvey and Shaw, 1995) is the book an important contribution to the scientific community?

Yes, Chalmers provides a theoretical foundation for a functional analysis (a way of modeling complex systems); in particular, his approach is consistent with theories of social-interactional dynamics, as in the work of (Bickhard and Terveen, 1995; Gould, 1987) and especially recent work on the evolution of intentionality (Barresi and Moore, 1996).

What are the really hard problems?

Referring to the human experience of reporting experiences of awe, angst, or exhilaration, of relating what we perceive to future action, and of introspecting, Chalmers asks, “Why should there be conscious experience at all?” (p. 4) and “Why do individual experiences have their particular nature?” (p. 5). Various theories explain “reportability” or attention or other cognitive functions, but not the “really hard questions”: “The problem of consciousness goes beyond any problem about the explanation of structure and function, so a new sort of explanation is needed” (p. 121). That is, we must explain not only how consciousness arises as physical processes, but why does the phenomenon have the characteristic “qualitative feel” that is experienced?

Chalmers rejects reductive explanation: “The existence of consciousness will always be a further fact relative to structural and dynamic facts, and so will always be unexplained by a physical account” (p. 122). He explains that consciousness is not just a kind of emergent phenomenon (for physical arguments will explain that, too) but a kind of *feature* that requires additional laws to describe it. By analogy, modeling electromagnetism requires the physicist to include new features of physical materials and laws for relating them. The same is true of chemical properties (e.g., pH, temperature, reaction coefficients) relative to atomic facts and of sensory experience (e.g., bitter taste) relative to chemical facts. Thus, a neurological or even psychological theory is not enough, it doesn’t cover the experiences and their relations and why—the really hard problem—a cognitive agent need know that such detections are even occurring.

The zombie argument

The story of Zombies perhaps best reveals Chalmers’ point of view. Zombies are hypothetical cognitive agents without conscious experience. Chalmers observes that “a physically identical zombie world is logically possible” (p. 123) and hence “the presence of consciousness [in our world] is an extra fact about our world” (p. 123). Crucially, consciousness is assumed to play a role, it provides additional information to the agents: “There is extra character due to the presence of consciousness”, namely “phenomenal information” which constrains the way the world is in a way different than physical facts alone. In particular, consciousness changes consciousness itself, so explanation of

consciousness in physical terms alone would not explain either the experience or its effects. Chalmers calls this a “non-Materialist” theory.

Put the other way, the fact that consciousness exists but the physical facts do not alone account for the *experience* of consciousness, means that “consciousness is not logically supervient on the physical” (p. 124). A special kind of force is at work: “When God created the world, after ensuring that the physical facts held, he had more work to do.... To ensure that the facts about consciousness are the way they are [that we are not all zombies], further features had to be included in the world” (p. 124). Chalmers elevates these “fundamental new features” to emphasize that they require “new kinds of natural laws,” just as the addition of electromagnetic charge and forces were added to physics (p. 127).

Starting with the data

At this point, it should be apparent that Chalmers’ approach is scientific after all. He simply wants to emphasize what kind of theory we should be looking for and how we should be talking about consciousness. He wants us to avoid being mired in discussions of neurons when we need to be talking about experience. So why should we talk about God and Zombies when we can study animals and neuropsychology? The role of philosophy is to help define the scope and nature of the problem, what really needs to be explained, what kinds of explanations are possible, and how these relate to other kinds of phenomena that have been previously explained. Specifically, avoiding reductionistic functionality means not defining consciousness away as if it were a byproduct of something else, but giving it operable presence and a role that transcends what the neurons are doing.

In this respect, it is surprising that Chalmers makes no attempt to taxonomize, describe, and lawfully relate the phenomena of consciousness. What could be a better argument than to give a list of these features, which he claims should be the proper concern of a scientific approach? By my reading, the lack of empirical analysis reflects Chalmers’ preferred approach (theoretical argumentation), rather than the impossibility of engaging in such research.

Chalmers argues that “consciousness is not directly observable in experimental contexts” (p. 215). But indirect observation (studying effects to infer causes) is common in science, including especially subatomic physics and

astronomy. Furthermore, research over the past twenty years reveals a wide variety of progress in empirically studying consciousness: vervet monkeys are experimentally manipulated in the wild to study their calls and hence to investigate their intentionality (Cheney and Seyfarth, 1992), the capabilities of children to form referential relations are studied experimentally (notably including the consciousness dysfunctions of autism (Baron-Cohen, 1995), and many patients with aberrant experience are studied to understand the varieties of consciousness and especially the nature of emotion, temporal sense, and (Cytowic, 1993; Damasio, 1994; Rosenfield, 1992; Sacks, 1987). Studies of prehistoric humans and our evolutionary ancestors are revealing how representational capability evolved, in a field sometimes called *mimetics* (Calvin, 1994; Donald, 1991). Cognitive capabilities of ravens, great apes, and beavers are compared to humans, suggesting that other forms of consciousness exist in the animal kingdom (Griffin, 1992). We can't observe feelings or beliefs, but we can observe people who are conscious or animals conscious in different ways. From this, we are starting to develop taxonomies that can be related to features of intentionality, categorization, and reference. Table 1 shows the kind of synthesis that is now possible.

Table 1. Levels of categorizing required for inference

Technical term	Internal relations	Example/paraphrase
5. symbolic reference	categorizing object as having referential interpretation	viewing a photograph as representing an animal
4. referential categorization	categorizing relation of belief to other categorizations, i.e., conception of having an idea about something	viewing own beliefs as being attributions
3. second-order intentionality	categorizing about categorizing, i.e., conception of beliefs; depends on concept of transactions	attributing beliefs or desires to self or another animal
2. conceptualization, i.e., second-order categorization = first-order intentionality (having beliefs and expectations)	categorizing about relations, i.e., conception of object-properties; depends on concept of activities (“instrumental actions and their expected consequences” (Barresi and Moore, 1996, p. 112))	“mother-offspring” relation between female vervet and juvenile
1. perceptual categorization	sensorimotor relations	a fish or insect being attracted to colored spots

Chalmers suggests that a theory of consciousness should explain principles that are independent of a particular evolutionary process (p. 121). I believe that relevant principles include mathematical relations of recursion and orders, process theories of self-organization (such as *selectionism* (Edelman, 1992)), and a theory of categorization. In Table 1 summarize how these ideas might be expressed together. These are levels of categorization, starting with perceptual categorization and then categorizing relations between categories (aka conceptualization). Higher orders compose categories to develop what we

typically call the “content” of knowledge: To conceive of beliefs (as “things” that agents “possess”), to conceive of a relation between a belief and a thing it is “about”, and to conceive of an object as denoting. Studies of autistics indicate that the higher orders of conceptualization require capability to coordinate multiple categorization in time—to hold multiple categorizations active and compose them, a basic form of abstract thinking that I call “conceptual coordination” (Clancey, in press). Crucially, forming such categorizations is inherent to conscious experience of different sorts, especially a sense of time and identity (Rosenfield, 1992; Sacks, 1987). Viewed another way, this ability to conceptually coordinate (juggling multiple ideas in the head) requires forms of neural activation (e.g., “global maps” and “reentrant links” (Edelman, 1992)), value correlation, sequencing, persistence, dynamic feedback, and composition that are not found in all animals and that appear to have different manifestation in different birds and mammals.

The ability to verbally model the world and behavior qualitatively changes the nature of conscious experience (Edelman, 1992). Indeed, much of cognitive science has viewed intelligence as equivalent to forming and manipulating descriptions (i.e., the inference and planning emphasis of “symbolic AI”). Just as cognitive science has mostly identified cognition with the verbal phenomena of higher-order consciousness, Chalmers’ approach to consciousness leaves out the different forms of the phenomena in the animal kingdom. With a taxonomy in front of us, such as is sketched in Table 1, we don’t start by wondering about “what special feature” or “what new natural laws” are required.

Consciousness is not a unitary phenomenon

Chalmers generally views judgment, beliefs, sensations, and perceptions as if they are all linguistic phenomena. Referring to color sensations he says, “A system judges that a proposition is true if it tends to respond affirmatively when queried about the proposition, to behave in an appropriate manner given its other beliefs and desires, and so on” (p. 174). But what does this have to do with ravens, monkeys, and proto-humans? What is the nature of beliefs in an animal without language, such as a cat?

A cat is not “queried about propositions”; it cannot participate in dialogues at all. A cat does not have desires like a person, for it cannot describe what it wants. But a cat does apparently *conceive* desires, such as to be affectionate with a

person or to catch a bird etc. Cats, like other animals such as ravens and monkeys, apparently categorizes relations between objects and their actions. Animal behavior that is playful, as in mock fighting of dogs, shows an ability to *conceive modalities of interaction*, ways of relating to the other animals. Yet all this occurs without describing what is happening. By not clearly distinguishing judgment from belief, perceptual categorization, and sensations, and by equating them all with *propositional description*, Chalmers treats conscious experience uniformly, in a highly anthropomorphic way.

The danger of a purely philosophical-analytic approach is that it tends to lump together a variety of data. Contrast Table 1 with the statement that the “basic processing correlate of consciousness is awareness or global availability” (p. 239). Is Chalmers referring to the primary consciousness found in chimpanzees (Edelman, 1992) or the higher-order consciousness in humans?

Consciousness is not a unitary, all or nothing phenomenon. Having lumped a wide variety of phenomena together by this term (not viewing it as a spectrum of relations and levels of categorization), Chalmers views variations as forms of experience or awareness—different kinds of “access.” Similarly, Chalmer’s idea of “representational content” appears as a flat analysis, mixing together emotion, judgment, and perception.

Notice how not distinguishing between levels of consciousness itself contributes to the view that consciousness should be treated as “a special kind of feature,” something fundamental requiring a new kind of natural law. Obviously consciousness fundamentally changes how a psychological system remembers, controls, plans, and indeed experiences its own behavior and state. But isn’t understanding the variety of such experience, ranging from perhaps bats to primates, the necessary starting place before expressing laws about it or even arguing that such laws are necessary?

Dogs and cats (being imaginative, playful, and social) are conscious in a different way than cows. Chimpanzees (conceiving of themselves as individuals in a group, but lacking a descriptive modeling language) are conscious in a different way than people. It’s not so much a law or new kind of “fundamental feature” that we need to understand as perhaps levels of categorization and self-organization—with experience itself reified, related over time, imagined, in different modalities, interactively shaped, and described. This view of a wider set

of natural phenomena—of data to be organized and explained—suggests a complementary approach that would help Chalmers' project by separating and relating reasoning, reference, intentionality, belief, conception, and perceptual awareness.

Relating a foundational theory to biological facts

Possibly Chalmers' insistence on non-reductionist theories is preventing him from treating consciousness as a scientific problem at all. He says that Edelman's "discussion is often vague," but Edelman lays out a theory based on principles of 1) selectionism, 2) orders of categorization, and 3) a reentrant-coupling (versus linear-causal) mechanism. Chalmers says that "Edelman gives no account of how all this processing should give rise to conscious experience" (p. 117). But in three dense books and one popularization, Edelman provides a new theory of categorization supported by computer models of neural systems, which significantly differs from the stored-description theories of memory, conceptualization, learning, and coordination. Chalmers doesn't appear to appreciate how much is new in this theory or perhaps he takes it for granted. (Similarly, his reformulation of Searle misses how the Chinese Room was a critique of programs like Mycin, which identify memory and knowledge with stored descriptive models of the world and behavior, rather than a way of understanding the dynamics of symbol systems in the brain.) Saying that Edelman is "only explaining the processes underlying conscious experience" (p. 118) is a bit odd, given that this fundamental goal of AI is far from being attained.

Chalmers' discussion of Rosenthal suggests a further lack of consideration of the variety of consciousness. Rosenthal makes a distinction about higher-order consciousness (HOC), in which a conscious state is "the object of a higher-order thought." Chalmers appears to mock this by saying "there is little reason to believe that we form second-order judgments about all of our experiences, including experiences of every detail of the visual field, of background noises, and so on" (p. 230).

First, notice that Chalmers here again characterizes perceptual experience as "judgments," a term Rosenthal uses to refer to a higher-order thought. They are talking past each other. Rosenthal has a notion of orders of categorization (as in Table 1). He would strongly agree with Chalmers that not every experience is an

object of thought and further correct Chalmers to say that not every experience is a judgment. But because Chalmers has a relatively flat view in which *all experiences are judgments*, he interprets Rosenthal as saying that “for every detail of experience, [there is] a first-order and a second-order judgment” (p. 231). Here Chalmers acknowledges a first and second order distinction, but he refers to second-order thought as reflective consciousness (making judgments about judgments) and first-order as “judgments available for global control” —which leaves us with some judgments that aren’t conscious at all!

Nevertheless, true to his thorough approach, Chalmers considers that he might be misusing the term “judgment.” Perceptual categorizations might be “registrations” and not beliefs or judgments “endorsed by the subject” (p. 232). But here his language is still entangling: Chalmers describes an illusion where “I mistakenly take there to be eight fingers before me... I judge that there are eight fingers... but my phenomenal experience is of seven fingers” (p. 233). How can you “experience” something of which you are not aware? “Taking” (a term from ecological psychology (Turvey and Shaw, 1995)), phenomenal experience, and judgment appear to be conflated. A distinction needs to be made between *conceptualization of number* as a higher-order categorization (relation) and *perceptual categorization of objects* (one finger-thing, one finger-thing, one finger-thing...). Conceptualization of number is not a judgment, for no rationalization is required (cf. the parrot that counts (Vauclair, 1996)). Judgment is a process in which Chalmers engages when trying to *explain* why he experienced eight fingers.

In summary, as indicated by Table 1, consciousness exists as a spectrum involving levels of categorization, roughly like the following: perception — conception (primary consciousness) — higher-order consciousness — judgment. Using other terms, which don’t map onto these directly, we could also express the levels as: qualia — intentionality — identity — rationality. Issues of awareness and access are thus reformulated in terms of *categorizations of categorizations*. I stress that verbal conceptualization, as described by most cognitive models, is a special case; representing (conceiving) intentionality and social action does not require verbal models (linguistic expressions of rules, plans, and causal relations).

Again, Chalmers acknowledges that consciousness is not a unitary phenomenon, at least partially, when referring to “nonconceptual content” (p. 383). But his discussion gets lost in a footnote of groping remarks: “There seems to be a consensus in the literature that the contents of experience are nonconceptual,” which is contrasted with the opinion that “animals do not have experiences.” Concepts are still described as being “possessed” rather than constructed by a higher-order categorization (operating on perceptual categorizations).

Similarly, Chalmers’ analysis of forms of awareness such as a background noise (e.g., a drill) rightly suggests that perceptual categorization may occur without HOC. A detection may occur subconsciously, such that after the fact we may report that an event. But Chalmers prefers to say that “the subject was aware of the drill all along” (p. 228), but didn’t experience it. But experience is the very hallmark of consciousness; separating the two as distinct phenomena makes a mess of the analysis. It is better to classify *different kinds of consciousness*. Detecting the drill is not a matter of awareness without consciousness, but of *categorization without consciousness*. A level higher, dreaming is an excellent example of *experience without HOC*. Indeed, primary consciousness may be what some birds and other primates experience. Edelman calls this phenomenon “the remembered present,” a flow of recognition and adjustment, without the processes of holding active and comparing that allows objectifying events, objectifying their relations, and thus categorizing experience itself (cf. Table 1). (See also (Kihlstrom, 1984).)

Chalmers says that reductionist functionality “collapses the conceptual distinction between consciousness and awareness” (p. 165). But a different view is that reductionist functionality reformulates this distinction so there are *varieties of consciousness*, with awareness being a functional capacity involving *orders of categorization* (categorizing operating on itself, i.e., categorizations of categorizations). Thus “conscious experience” has different forms that can be understood in terms of what kinds of categorizations (and hence beliefs) are forming. By this approach, the word “experience” is not merely a label applied arbitrarily, but is itself broken apart to cover subconscious experience (not aware at the time of categorization), perceptual experience, conceptual experience, etc. All categorization is experience; it’s not a special kind of phenomenon needing

explanation. The experience of HOC in people then becomes different from the experience of perceptual categorization, sensation, or emotion in other animals. Furthermore, by relating emotion to categorization, in the way Edelman makes “value” inherent to categorization, we can begin to explain why the “qualitative feel” of consciousness exists and why it is necessary for higher-order forms of categorization (e.g., Damasio’s theory of the relation of emotion to decision making).

In the animal kingdom there are different forms of consciousness with distinctly different functional capabilities to project, reexperience, relate, model, rationalize, theorize about experience. Consciousness can therefore be productively related to physical facts about the brain and social relations, which scientifically speaking would be a necessary step before attempting to formulate new natural laws for characterizing it. In this respect, the study of consciousness would be like the early study of electricity—with consideration of lightning, “static” sparks, the Linden jar, and even light bulbs before new fundamental features (“charge” or “spin”) needed to be postulated. Indeed, some would say that such physical features will include features of neuronal groups and that’s where we should be looking now, so later we might create fully equivalent physical organizations in other, artificial “implementations” of life (e.g., see the architectural theories of (Calvin, 1994)).

Against reductive functionalism: Why not a social-interactional approach?

But have I fallen back into the explanatory enterprise that Chalmers says is the easy part? Chalmers doesn’t want consciousness to be reduced to a certain kind of physiological property, such as learning or categorization, that “play[s] a certain causal role in a system” (p. 164). He claims that “This view... misrepresents what it means to be a conscious experience.” Am I failing to appreciate the problem that interests Chalmers and substituting my own preferred approach? In fact, Table 1 is derived from a social-interactional approach, not a reductionist attempt to explain everything in terms of neural categorization. In this section, I will briefly discuss these considerations.

First, Chalmers is concerned that even after we have explained the functional capacities of consciousness the *experience* of being conscious will not be explained. But if we understand consciousness as internal categorizing, why wouldn’t it become evident that the feeling of being aware (by the feedback of

changing attention), as a feeling of being in control, of experiencing what we intend to experience, would be explained at least as well as the explanation of taste or bright light or harmonious sound or any emotion might be explained physiologically? Again, if we don't need a special theory for explaining the experience of sensation (and patently that is so, for not all creatures with emotion have HOC), then why would we need a special theory for explaining the experience of HOC.

Chalmer's second objection to reductive functionalism is that "we can imagine any functional role being played in the absence of conscious experience" (p. 165). I can imagine a poorly understood function being treated that way, but I doubt that the functional role of HOC relative to (social) *identity* could occur without conscious experience. Specifically, a society of Zombies would not be functionally equivalent to a society of humans. Without the ability to emphasize, to project, to experience anticipated effects in one's mind, morality would be impossible.

Furthermore, studies of cognition as a situated phenomenon reveal that rationality is inherently a conscious phenomenon. Representing is experienced, interactive construction and reinterpretation, whether the media consists of paper, verbal dialogues, or imagination (Clancey, 1993; Roschelle and Clancey, 1992; Schön, 1987). That is, causal models and plans are inherently *experienced* constructions. Without consciousness there can be no planning (and visual-gestural conceptualization may allow a form of planning in the imagination without verbal description (Heinrich, 1993). Hence, an "interactional" or "transactional" view is required for understanding the function of consciousness; this is not reductionistic, but an analysis that combines the development of neural and externalized representations.

Such analyses, as I have presented in Table 1, are now becoming relatively sophisticated in cognitive science, as in studies of autism, the calls of vervets, and neurological dysfunctions. Crucially, the conceptual content of the categorizations, the relations being represented by the subject, concern social action. Thus, this is sometimes called a social-interactional analysis. The zombie thought experiment presented by Chalmers does not incorporate the social-interactional aspects of a functional study of consciousness. Chalmer's apparent ability to imagine human culture without consciousness reflects a lack of

appreciation of the nature of identity, sense of time and place, and social relations. Experience is crucial because it involves emotional feedback; action is not automatic (no matter how complexly planned), but is conceived as participation and contribution. Action is social when the actor conceives of himself/herself as an actor, a participant in a social choreography (Clancey, 1997).

For example, consider again how I began this review by first considering what kind of review would be socially proper: I recalled my recent painful experience and consequently reformulated my behavior so it would avoid hurting someone the way I had been hurt. I thought, what do I want to accomplish? How do I want to relate to Chalmers? What is my identity? To be a scientist with an open mind who will learn something by reading this book and reviewing it? Or to be a warrior putting down a strange, foreign intrusion on my thoughts? To advertise my own ideas (yes, a little, but that's fair) or to relate Chalmers' contribution to others?

Without the emotional aspect of experiencing another person's actions, I would not have formulated my own behavior in terms of a social relation between me and Chalmers (whom I have not met). However rational, my behavior as a Zombie could not have been moral. A Zombie might learn the rules of etiquette, but the interpretations required and the ongoing learning must occur through individual experience, through *caring about what others feel* because you care about what you feel. Without conscious experience, there would be no experience of what it feels like to be mocked or to have one's work treated thoughtlessly and hence no basis for becoming more adept as an actor in a scientific community of practice. Without emotion zombies could not form communities at all, let alone organize to develop the cognitive capability of an intellectual community. Zombies could be cognitive, but their collaborative participation could only be like the repetitive work of ants, rather than the collaborative conceptual construction of science.

Chalmers has aptly chosen the Zombie as a counterpoint to humans. What he must do now is consider again what he has emotions about, what he experiences. He will observe that it is not all dry "representational content," but *a sense of himself as a feeling person*, who is hurt, elated, thoughtful, extemporaneous, and so on. This awareness of *being a person with feelings* is fundamental for being a

person at all, of being part of and helping to reshape a culture. So before we go off focusing on why a neurological explanation is not sufficient so new kinds of natural laws are required, we might start by looking for what the social-psychological sciences have to offer about the experience.

In summary, explaining why there is experience at all will have multiple forms: how it arises physiologically, how it develops interactively, why it is functional socially and psychologically, and how it evolved. We might not have predicted the existence of consciousness from physical facts alone, but we couldn't have predicted the existence of butterflies either. Indeed, to acknowledge the difficult problem Chalmers emphasizes, we have far to go in explaining why we experience grandeur looking at a landscape, find beauty in a flower, and enjoy the taste of bitter chocolate. But these explanations will have a different form than understanding *how consciousness evolved* (which I claim is the same as asking why it exists). Nevertheless, together the relation between consciousness and "qualitative feel" is clear: Consciousness is not just a capability to relate or reflect (higher-order categorization), but inherently involves feelings. By this theory, "value" is part of every categorization. Without value there would be no *interest* or exploration or investigation or inquiry. Without curiosity there would be no science. Having an experience necessarily entails conceiving, and all conceiving is evaluative.

Conclusion

Chalmers inquires about the nature of consciousness by which he means "the subjective quality of experience: what is it like to be a cognitive agent?" He wants to explain why a cognitive agent has conscious experience at all. I have advocated that we base a fundamental theory on what we know about the varieties of consciousness from studies of neural categorization, structural analysis of neurodysfunctional phenomena, emotional experience, and cross-species evolutionary data.

From this starting point, cognition appears to be inherently emotional, interactive (behavioral-reflective), and socially conceived (about participation and contribution). A society of robots or Zombies without emotion would be more like horses standing in a field than people rushing down the street; a society of Zombies without a moral (social) conception, would consist of individuals without identity and hence groups without culture. Such a society

would be more like fish perhaps (and certainly not as complex as hyenas) and, according to current theories of cognitive evolution, would statically reproduce limited cultural forms (such as the apparently mindless patterning of Homo Erectus tools over tens of thousands of years).

Chalmers' book provides a good introduction to different ways of understanding consciousness. Its strength is in philosophically relating different kinds of theoretical approaches and especially explaining why a reductionist explanation is unsatisfactory. But by not working with the phenomenon itself and trying to relate its various forms to both biological and social facts, Chalmers has ironically missed the opportunity to elevate his study above the reductionist approach. What is needed now to complement Chalmer's book and to address his concerns is another book titled, *The Conscious Animal: In search of a fundamental theory*.

References

- Baron-Cohen, S. (1995). *Mindblindness*. Cambridge, MA: The MIT Press.
- Barresi, J., & Moore, C. (1996). Intentional relations and social understanding. *Behavioral and Brain Sciences*, 19(1), 17.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence Journal*, 72:, 173-175.
- Bickhard, M. H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Amsterdam: Elsevier Science Publishers.
- Calvin, W. H. (1994). The emergence of intelligence. *Scientific American*, 271(4), 100-108.
- Cheney, D. L., & Seyfarth, R. M. (1992). Precis of how monkeys see the world. *Behavioral and Brain Sciences*, 15(1), 135-182.

Clancey, W. J. (1993). Situated cognition: How representations are created and given meaning. In *International Federation for Information Processing, Working Conference "Lessons from Learning"*. (pp.). Archamps, France: Neuropes Lab.

Clancey, W. J. (1997). The conceptual nature of knowledge, situations, and activity. In P. Feltovich, K. Ford, & R. Hoffman (Eds.), *Expertise in Context*, (pp. 247-291). Cambridge, MA: The MIT Press.

Clancey, W. J. (in press). Conceptual coordination: Abstraction without description. *International Journal of Educational Research*.

Cytowic, R. E. (1993). *The Man Who Tasted Shapes*. New York: GP Putnam's Sons.

Damasio, A. R. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: G. P. Putnam's Sons.

Donald, M. (1991). *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, MA: Harvard University Press.

Edelman, G. (1992). *Bright Air, Brilliant Fire: On the Matter of the Mind*. New York: Basic Books.

Effken, J. A., & Shaw, R. E. (1992). Ecological perspectives on the new artificial intelligence. *Ecological Psychology*, 4(4), 247-270.

Gould, S. J. (1987). *An Urchin in the Storm*. New York: W.W. Norton & Company.

Griffin, D. R. (1992). *Animal Minds*. Chicago: University of Chicago Press.

Heinrich, B. (1993). A birdbrain nevermore: When put to the test, ravens display insight. *Natural History*, 102, 50-58.

Kihlstrom, J. F. (1984). Conscious, subconscious, unconscious: A cognitive perspective. In K. S. Bowers & D. Meichenbaum (Eds.), *The Unconscious Reconsidered* (pp. 149-211). New York, NY: John Wiley & Sons.

Roschelle, J., & Clancey, W. J. (1992). Learning as social and neural. *The Educational Psychologist*, 27(4), 435-453.

Rosenfield, I. (1992). *The Strange, Familiar, and Forgotten*. New York: Vintage Books.

Sacks, O. (1987). *The Man Who Mistook His Wife for a Hat*. New York: Harper & Row.

Schön, D. A. (1987). *Educating the Reflective Practitioner*. San Francisco: Jossey-Bass Publishers.

Steels, L., & Brooks, R. (1995). *The artificial life route to artificial intelligence: Building embodied, situated agents*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Turvey, M. T., & Shaw, R. E. (1995). Toward an ecological physics and a physical psychology. In R. L. Solso & D. W. Massaro (Eds.), *The Science of the Mind* (pp. 144-172). New York: Oxford University Press.

Vauclair, J. (1996). *Animal cognition*. Cambridge, MA: Harvard University Press.