

Introduction

Michel Weber and Anderson Weekes

We introduce this volume with a brief preview of its overall concerns (Section I), followed by a précis of the individual contributions (Section II) and an overview of the most important literature with a similar or related focus (Section III). The preview will show that this is not a collection of specialty scholarship, but a volume rightly intended for the broadest possible learned readership. The uniqueness of its approach is tempered by the generality of its concerns. The précis then situate each contribution in the larger context of the book's philosophical and interdisciplinary ambitions, while the last section situates the book in the broader context of today's intellectual landscape, where a growing body of literature reinforces its cause without anticipating its results.

In this Introduction, we adopt the following conventions in referring to the chapters that follow. (1) Chapters are identified by authorship. Contributors' proper names, including those of the editors, refer to their respective contributions to Parts II–V. Proper names do not reference the contributions to Part I, which resulted from collaboration (chapters 1 and 2) or consensus (chapters 3 and 4) between the editors. These chapters in Part I we refer to simply as the contributions of “the editors.” “Contribution(s) of the editors” does not refer their individual contributions in Parts IV and V. (2) Source and locus will not be given for quotations if they are taken from the chapters that follow. Unless otherwise indicated, all quotations in this Introduction are from the named author's contribution to the present volume.

Main Themes of the Book

Alfred North Whitehead's philosophy was a protest against the compartmentalization of knowledge. A specialized subfield of philosophy focused

on Whitehead interpretation is therefore something of a paradox. Given the daunting complexity of Whitehead's writing, literal exegesis and historical scholarship aiming at an "immanent" interpretation of his thought have a continuing and obviously important role to play, but Whitehead himself would scarcely recognize such activities as his rightful legacy. A failure of Whiteheadians to be sufficiently Whiteheadian in this regard may well be the reason Whitehead's ideas have seemed at times to be threatened with extinction and mostly available in fossil form. If this is changing, it is at least partly because outsiders are storming the museum. Straightaway this has opened vast avenues of new dialogue with unsuspected partners, to which this book bears witness.

This volume brings multiple disciplinary perspectives to bear on Whitehead's psychology (which, in a way, is his whole philosophy—a metaphysics of experience) in order to analyze it in terms of relevance to contemporary consciousness studies. Accordingly, we have gathered contribution from scholars whose areas of research are diverse and often do not include Whiteheadian process philosophy as a subfield of expertise, but whose own intellectual paths have led them to recognize an important kinship with Whitehead.

The area of consciousness studies proves to be a busy intersection: a place where one can't help but meet everything from metaphysics to psychotherapy. This is not happenstance. It reflects the nature of the beast we are tracking, and we have not shied from it. This accounts for both the broad scope of the volume and the diversity of its contributions.

In important respects, this book complements David Griffin's *Unsnarling the World-Knot: Consciousness, Freedom, and the Mind-Body Problem* (1998), which grew out of an interdisciplinary conference sponsored by the Center for Process Studies at the Claremont Colleges in 1994, "Consciousness in Humans, Animals and Computers. A Scientific-Philosophical Conference." Bringing a Whiteheadian perspective to contemporary consciousness studies, Griffin effects a broad synthesis of the issues currently under debate and at the same time provides an excellent introduction to Whitehead's psychology. We reverse directions. Bringing different contemporary perspectives (including Griffin's) to bear on Whitehead's psychology, we replace synthesis with analysis and highlight the richness and polyvalence of Whitehead's ideas. Of particular concern to this volume is the role that Whitehead's process philosophy can play in providing an interpretive framework for neuropsychology, and, conversely, the role that neuropsychology can play in providing an empirical model for Whitehead's concept of process and an empirical confirmation of his theory of consciousness. According to Whitehead, consciousness is a process—a very specific kind of process that, despite its uniqueness, holds the key to understanding process as such. Consequently, a number of important findings of neuropsychology, some of them familiar, but some of them quite new and even startling, will figure decisively in these pages.

The contributions to this volume can be grouped according to a number of shared themes. Several contributors (David Griffin, Katzko, Shields, Pachalska, and MacQueen, and the editors in Part I) show that some recognizably Whiteheadian issues are at stake in the current debates about consciousness and that Whiteheadian ideas can be exploited—sometimes in ways that Whitehead could not have anticipated—to advance the debate beyond some well-known sticking points. Two of the contributors (Rosenberg and Weekes) explore the curious connection Whitehead alleges between consciousness and causation. One author proceeds by a conceptual analysis of the structure of explanatory theories, the other proceeds phenomenologically, but they both lend support to Whitehead's signature idea, refereed in chapter 4, that scientists and philosophers find consciousness very difficult to explain for the same reason that they have a problem understanding the nature of causation and the basis for induction: due to inherent constraints, the theoretical activity known as explanation tends to suppress the specifically processual aspect of becoming, which Rosenberg calls the receptive face of causation and Whitehead calls concrescence. The argument in brief: By suppressing the dynamic aspect of becoming, explanatory theories render essentially processual phenomena inscrutable. The paradigm of an "essentially processual phenomenon" would be, according to Whitehead, *experience*, of which *consciousness* is only the most sophisticated (and deceiving) sort. Others (Verley, Weekes, and the editors in chapters 3 and 4) expand systematically on Whitehead's scholarly critique of modern philosophy, which Whitehead casts almost entirely as a critique of its favorite concept, consciousness. But there are really three main themes that connect the contributions to this volume.

First, running through all the contributions to this volume is the critical insight that consciousness is *not* the *sui generis* phenomenon it is usually taken to be—in philosophical and scientific discussion as much as in the everyday understanding informed by lay sensibilities. Closely related to this principal theme is a secondary theme that connects more than half of the contributions to the volume (David Griffin, Donald Redfield Griffin, Shields, Velmans, Rosenberg, Weekes, and the editors in chapter 4). It is the question of the distribution of consciousness in the natural universe. Approaching the question from very different angles, each of the contributors just named argues that consciousness (or something much more primitive, but in the same category) is more widely distributed than customarily supposed. At the opposite extreme from the assumed exclusiveness of human consciousness is the position usually referred to as "panpsychism." In the contributions by David Griffin, Weekes, and the editors in chapter 4, Whiteheadian arguments for the universal distribution of some kind of extremely rudimentary (pre- or proto-conscious) experience are explored. In the contributions by Shields, Velmans, and Rosenberg, the possibility of such a distribution is supported with robust and original arguments that will give many readers pause.

The principal theme first: in different ways, each contribution to this volume seeks to relativize the concept of consciousness that is normally taken for granted. The reflective and attentionally focused consciousness that tends to be identified with consciousness absolutely belongs to a wide and multidimensional spectrum of conscious states—or, if someone insists on reserving the word *consciousness* for the particular apex of human experience that is reflective and attentionally focused (as, indeed, Whitehead allows), then we must say that consciousness belongs to a wide and multidimensional spectrum of experience, most of which is “unconscious” or only partially or obliquely conscious. (To know whether the word consciousness is being used in its broader or narrower sense, the reader of this volume will have to be attentive to contexts: we have not thought it appropriate to legislate uniformity in this matter on our contributors.) The point is that “paradigmatic” consciousness is only one of many kinds of consciousness/experience. Its isolation as a paradigm is the result of a variety of organic, psychological, social, and historical processes of development, refinement, and selection (some necessitated by survival and social existence, others contingent, but ossified as dogma). Its view of the world is therefore not absolute or final, but conditioned by these processes.

Weber argues that everyday consciousness of the natural attitude is not an absolute *given*, but an artifact of socialized ontogeny. It results from the constraints of utilitarian and social rationalization that operate on the individual in mostly unconscious ways. Verley, Weekes, and the editors in chapters 3 and 4 expose important ways that the preferred concept of consciousness is an artifact of biases peculiar to the modern philosophical tradition. Pachalska and MacQueen and Schweiger et al. discuss varieties of consciousness revealed by brain pathology in humans. These varieties of consciousness differ markedly from the usual paradigm of reflective and attentionally focused awareness of objects *qua* objects. The authors argue, moreover, that these varieties of consciousness are not abnormal. On the contrary, they are the normal subphases in the moment-to-moment microgenesis of consciousness, with the qualification that they are abnormally exposed because pathology has arrested the microgenetic process at a preterminal phase of realization. Velmans also stresses ways that human consciousness results from and is conditioned by processes of refinement or selection operating on a spectrum of broader and more basic kinds of consciousness/experience. He goes so far as to suggest that what we normally think of as consciousness in humans may be the sophisticated result of a highly selective release from inhibition of what is in reality a pervasive and primitive kind of awareness intrinsic to all organic matter or even to all matter, *period*. In fact, the startling generality of this conclusion is something to which Rosenberg’s analysis of causality led him for reasons wholly unrelated to Velmans’ argument: there

is, Rosenberg concludes, a primitive experiential aspect in every causal nexus, and that means in every event, *period*. Velmans wonders if the purpose of centralization in complex nervous systems isn't to prevent overload by inhibiting this primitive consciousness throughout most of the system, allowing for a selective focus on information of critical relevance. The late Donald Redfield Griffin's examination of consciousness in animals relativizes human consciousness in a more straightforward, if no less controversial way. Although he is mainly concerned with documenting ways that animals can be seen to have consciousness similar to ours, its wider distribution in the animal kingdom means that consciousness is a genus of which human consciousness is only a specific kind. Reflective human consciousness may indeed possess an epistemological privilege, but this is no longer something it can take for granted on the grounds that it defines and exhausts what consciousness is. Nor can we assume any longer that the minimum identity conditions of consciousness/experience are in any way obvious—and least of all obvious from self-conscious reflection or introspection in human beings.

The secondary theme: Since the distribution of consciousness/experience in the universe bears in an obvious way on the relative or absolute status of “paradigmatic” human consciousness, the distribution question becomes another connecting theme in this volume. The distribution of consciousness/experience cannot be divorced from the question of the minimal conditions for the existence of consciousness/experience. The more complex the conditions, the less distributed it will be. Conversely, the more distributed it is, the more the complexity of human consciousness must appear as the result of specialized constraints that exclude other, more basic kinds of consciousness/experience (or at least their foregrounded manifestation). The distribution of consciousness, alluded to by Weekes and the editors in chapter 4), is a thematic focus of contributions by David Griffin, Donald Redfield Griffin, Velmans, Shields, and Rosenberg, and it naturally leads to the hot-button issue of panpsychism, notoriously associated with Whitehead's metaphysics. This special case of the distribution question brings us to the third unifying theme of our volume.

The signature thesis of Whitehead's metaphysics is that the core of actuality is always some kind of experience. Avoiding the misnomer “panpsychism,” David Griffin has aptly dubbed this thesis *panexperientialism*.¹ It has long been common to dismiss Whitehead's panexperientialism hastily on the grounds that it is patently absurd to suggest that things like rocks and toasters have experience or that subatomic particles are conscious. But this conflates panexperientialism with panpsychism. As Griffin shows in his contribution, these objections are misdirected. First, panexperientialism distinguishes between conscious and nonconscious experience (in the same way that Velmans, for example, distinguishes between very high-grade and very

low-grade “consciousness”). While all entities that are genuine individuals (including subatomic events) are postulated to have some kind of experience, only the most complexly organized of compound individuals have conscious experience. The simplest individuals, presumably Planck-scale units of nature, have an extremely rudimentary kind of “experience” that would consist in little more than a sensitivity or responsiveness to their environment that was not 100% predictable. Second, drawing on an important clarification made by Charles Hartshorne, panexperientialism distinguishes between compound individuals, such as organisms, which are genuine integral individuals and thus have a coherently unified experience, and merely cohesive aggregates (like rocks and toasters), which have individuals as their micro-constituents, but are not themselves integral individuals and as such have *no* experience.²

It is also important to keep in mind that, according to panexperientialism, individuals *per se* are momentary events and do not endure for more than the briefest possible duration. Enduring entities, such as electrons or psyches, are made up of many such durational individuals forming a temporal series that is cumulative and characterized by overwhelming similarity between any two consecutive members. Consequently, panexperientialism does not attribute a mind or soul to anything but enduring compound individuals.³ It is only the small differential of a momentary experience that panexperientialism attributes to every individual regardless of status—compound or simple, bound within a cumulative series or not.

A number of contributors to this volume explore arguments, both logical and empirical, for taking panexperientialism seriously. Logical arguments of various types (metaphysical, transcendental, conceptual) are advanced by David Griffin, Shields, Rosenberg, and the editors in chapter 4. Empirical arguments must appeal to scientific evidence about the distribution of consciousness/experience. Given the roughly inverse relationship between complexity and distribution of consciousness/experience, a number of contributors take up the critical question that unavoidably arises in this context and that any serious assessment of Whitehead must address: how far downscale in complexity of organization can types of individuals be found that still appear to have some kind of experience? As Thomas Nagel has put it, “if one travels too far down the phylogenetic tree, people gradually shed their faith that there is experience there at all” (Nagel 1979, 168). David Griffin alludes in summary form to the growing range of evidence available on this important topic. Three of our contributors, Donald Redfield Griffin, Velmans, and Shields, observing a division of labor naturally suggested by their respective areas of expertise (biology/animal ethology; psychology/neuropsychology; philosophy/physics), examine empirical evidence that is in many cases startling.

Of course, as an empirical question, how far downscale in organizational complexity experience goes is something that could be answered only indirectly,

by inference from decisive clues. As Donald Redfield Griffin notes, prejudices on this question are likely to disguise themselves as disagreements about the criteria for validly inferring the existence of consciousness/experience in other life forms. Donald Redfield Griffin and Velmans both dispatch a number of specious objections to less-than-human consciousness simply by insisting that the same standards of interpretation we apply to other human beings be applied to other organisms: similar behavior and similar brain anatomy and physiology cannot be relevant in the one case and not in the other.

Regarding how far down the scale of complexity experience goes, David Griffin notes that Descartes set the cutoff point right below human beings, but that natural science has been pushing it down ever since—animal ethologists pushing it as far down the phylogenetic tree as bees, biologists as far down as single cell organisms or even as far as bacteria or DNA, and some physicists right down to the Planck-scale units of nature. In this light, panexperientialism looks like a position toward which empirical science is tending all by itself under the weight of the evidence. But it is also a position increasingly under reassessment for strictly philosophical reasons. Part of the reason for this volume is the fact that this slighted position is beginning to garner mainstream consideration. As the deadlock between dualism and materialism in consciousness studies becomes more tiresome, appreciation for the important differences between the less plausible panpsychism and the more plausible panexperientialism grows. To many, panexperientialism is looking more and more like a viable *via tertia* (or “third way”).

And we note last that an answer to an important objection to panexperientialism also emerges from this volume. The demand is rightly made: what other meaning can “experience” possibly have than the experience human users of language are readily familiar with in themselves? Consequently, if the concept of experience is attenuated and generalized so that it no longer designates what human speakers normally mean in one language or another, how can it mean anything at all? If the experience of a bacterium or an electron is totally unlike ours, what point is there in calling it experience at all? It seems that we are either saying something obviously wrong or not really saying anything at all.

A great deal of empirical research into the distribution of experience in the universe is addressed in this volume, and it suggests a much wider distribution than is traditionally conceded. But our contributors also show that within its own compass human experience is rich and multiform enough to supply the semantic Rosetta Stone needed to talk meaningfully about these nonhuman manifestations of experience. By noticing that even human experience encompasses kinds of awareness that fall far short of the lucid, objectifying consciousness of the well-socialized adult, we can free ourselves from the conceit that consciousness as such must be narrowly construed as

something uniquely human (Weber, Weekes, and the editors in Part I). If there is, even within human consciousness as we experience it now, the vestiges of qualitatively distinct kinds of consciousness corresponding to each evolutionary stratum of our brain (Pachalska and MacQueen and the work of Jason Brown they draw on), then there is no reason to think that we do not share these more basic forms of consciousness with those species that have only the more primitive brain formations. If human consciousness passed through more primitive stages in its own evolution, if it passes through cumulative phases in the recovery from unconsciousness, if it passes through distinct phases in the early motor development of the individual and its subsequent socialization, if it passes through a nested hierarchy of phases as it emerges moment by moment from the neural activity of the brain (Schweiger et al., Weber, and the literature they reference), then it makes no sense to deny that we have any criteria for generalizing an attenuated concept of consciousness beyond human experience. The important question that remains open is not whether, but how far the concept of experience can be meaningfully generalized.

In order to set the individual contributions to this book in the context of its overarching themes, the editors supply the following précis. As for the arrangement of the contributions, the editors hoped to order them in a way that would allow each to benefit the most from being read in the sequence settled on, but to some extent the order is unavoidably arbitrary, and each contribution does, in fact, stand on its own.

Précis of the Contributions

David Griffin not only provides a lucid, jargon-free overview of Whitehead's theory of consciousness, but he also manages to bring it directly into the arena of current debate. The fact that there are conceptual common denominators allowing for a meaningful, if virtual debate between Whitehead and contemporary theorists may surprise many who have, perhaps understandably, stumbled at the outset over Whitehead's dense and idiosyncratic language.⁴

Griffin enlarges on an idea very important to Whitehead: that the philosopher may not deny in theory what she presupposes in practice. Griffin notes that this fallacy involves what Apel and Habermas call a performative contradiction: asserting something that violates the conditions of possibility of making the assertion in question.⁵ Accordingly, Griffin elaborates four criteria having to do with the conditions of the possibility of theorizing. These performatively undeniable facts Griffin calls ideas of "hard-core common sense," and any adequate theory of consciousness must account for them: the idea "that conscious experience exists, that it exerts influence upon the

body, that it has a degree of self-determining freedom, and that it can act in accord with various norms.”

It's obvious how a statement such as “Consciousness does not exist” involves a performative contradiction,⁶ but a statement such as “Consciousness is merely an epiphenomenon” runs afoul of performative consistency, too. If the uttered statement means what the speaker intended, then it must be conceded that her mind has had an effect on her body. If the statement was affirmed because the speaker thought it was true, then she must have had the freedom to let herself be motivated by an ideal such as truth.⁷

Griffin takes a broad survey of the important players in current debates and finds that the discussion remains boxed in by the traditionally dominant paradigms of reductionist materialism and Cartesian dualism. Assessed against his four performative criteria, neither of these positions is acceptable. Materialism runs afoul of all four, and dualism runs afoul of all but the first. The contemporary debate is therefore framed by what amounts to a false dilemma. The overlooked third option or *tertium quid* would be a naturalism that was not reductionistic or, by the same token, an interactionism that was not dualistic. According to Griffin, Whitehead's position meets this requirement: “With dualists, Whitehead agrees that consciousness belongs to an entity—a mind or psyche—that is distinct from the brain, and that genuine freedom can, partly for this reason, be attributed to conscious experience. With materialists, Whitehead shares a naturalistic sensibility, thereby eschewing any even implicitly supernaturalistic solution to philosophical problems, and, partly for this reason, rejects any dualism between two kinds of actualities. Like materialists, in other words, he affirms a pluralistic monism. He thereby regards consciousness as a function of something more fundamental. And yet he, like dualists, rejects the reductionism involved in functionalism as understood by materialists.”

What makes this *tertium* possible is Whitehead's theory of experience as the core of actuality. In other words, panexperientialism is uniquely qualified to avoid the pitfalls of the materialism-dualism dichotomy. This yields an essentially transcendental argument for panexperientialism: an argument based on performative consistency as a condition of the possibility of conscious activity. Griffin alludes to two subsidiary arguments, as well. For one, he notes the trend in empirical science to cast the net of experience more and more widely. In the absence of a sufficient reason to draw a hard line at a particular point (as Velmans discusses in his contribution, as well) a *prima facie* presumption of validity should be granted to the logical extrapolation of this trend. For another, Griffin (like Velmans) notes the difficulties that arise once we draw such a hard line. It creates a discontinuity and a dualism difficult to square with the theory of evolution. Griffin even makes the case that panexperientialism alone can explain how consciousness could arise in the

course of evolution. This is no bluff since he charts the intervening phases that would lead to the gradual or staggered evolution of consciousness from the unconscious or merely incipient intentionality of experience in its most rudimentary shape. But Griffin stresses that his most important argument is the transcendental one: that *only* panexperientialism can satisfy the four performative criteria he sets out.

Michael Katzko offers us a complementary survey of current debates on consciousness, looking especially at three influential philosophers who strongly disagree with one another: David Chalmers, Daniel Dennett, and John Searle. On the surface it would seem that the positions of these three philosophers have relatively little in common, but Katzko argues that they share fundamental presuppositions. This becomes evident when we examine how each philosopher construes the problem he thinks a theory of consciousness is obligated to solve. In each case the problem defining his objective is essentially the same: the difficulty of understanding how the physical could possibly give rise to the mental. (Notice that the question is essentially about causation.) The answers they and many others in the literature give to this question are to be sure quite different. Some answer that it's *not* possible: either because there's really no such thing as the mental or because the mental isn't caused or created by the physical at all, even if it always somehow corresponds to it. Others answer that it's possible, but as yet incomprehensible, or that it's possible, but inherently incomprehensible, and so on. But all these solutions start from the same conception of the problem. Katzko sees them as so many attempts to make a virtue of necessity—having uncritically embraced false dilemmas bequeathed to them by the seventeenth century, contemporary philosophers have no choice but to countenance one side or another. As a whole, the contemporary discussion takes it for granted that what we need to do is rethink our understanding of the mental in order to render its relation to the physical unproblematic. The homogeneity and one-sidedness of the contemporary discussion becomes evident when we compare Whitehead's philosophical conviction that what we need is a new concept of the physical. Why after all should our concept of the mental do *all* the accommodating, especially when the concept of the physical to which accommodation is demanded was discredited by physics a century ago? (In this connection Shields rightly speaks in his contribution of a "cultural lag.")

Looking at the contemporary debate with the eyes of a clinician, Katzko sees a disordered discourse, hamstrung by arbitrary and unacknowledged limitations. Using Whitehead's framework of concepts to make this diagnosis, he shows how the operative concepts of the current debate (mind, the physical, intentionality, qualia) illustrate many of the fallacies described by Whitehead (misplaced concreteness, simple location, vacuous actuality).

The common denominator of these fallacies is the methodological mistake of commencing investigation with abstractions to which one subsequently attempts to reattach what was left out by appealing to other abstractions. The alternative is to begin inquiry with the complete context of the concrete experience in which theoretical investigation operates (including such things as what the investigation presupposes “in practice”), seeking from the outset a generalization that is inclusive rather than exclusive. Whitehead agrees with thinkers such as Bergson, James, and Bradley that prereflective experience is characterized by an unbroken wholeness to which reflection must always do justice. Whenever analytic abstractions prescind⁸ from this wholeness and treat the world as a set of typological isolates that can be recombined to “explain” concrete phenomena—which the reader will recognize as the resolute-compositive method that inaugurates modern thought—it will subsequently be difficult, not to say impossible, to understand how things are nevertheless interconnected in nontrivial ways. Katzko shows how well the now popular concept of intentionality illustrates this problem. It reflects an attempt to reattach the contextuality and relatedness that was left out of “mind” when it was conceived as a kind of substance to begin with. And whenever it is supposed, for example, that the “content” of experience does not entail the existence of the “external” world, the mind (or consciousness) is being treated, at least implicitly, as an autonomous entity—that is, as a substance.

Whitehead does not deny the great practical and technological triumphs of the resolute-compositive method, but he thinks it contributes little to philosophical understanding. It is not possible to explain concreteness—what Aristotle called *tode ti*—as a collocation of abstractions. The task of philosophy is therefore not to explain the concrete by means of the abstract, but to explain how abstractions arise from the analytic partitioning of concrete experience. Katzko shows how the partitioning of experience preferred by the current debate, lacking a self-conscious methodological grounding in the holism of prereflective experience, is often an arbitrary throw-back to platitudes of the seventeenth century.

The late Donald Redfield Griffin’s contribution goes a long way toward assuaging the uneasiness noted by philosopher Thomas Nagel that “if one travels too far down the phylogenetic tree, people gradually shed their faith that there is experience there at all” (Nagel 1979, 168). Griffin does not make the specifically Whiteheadian distinction between conscious and unconscious experience, although he does not rule it out either.⁹ Defining consciousness as “subjectively experiencing feelings or thoughts,” he shows that accumulating evidence strongly suggests that many species of animals have consciousness. His discussion touches on apes, parrots, dolphins, and bees.

Griffin suggests that evidence of animal consciousness is routinely ignored because of an overriding philosophical prejudice “that there is no

conceivable way in which valid, objective evidence about conscious experiences of other species can ever be obtained.” With regard to species very unlike our own, the prejudice takes the stronger form that conscious thinking simply could not be possible. The result is a double standard, where something naturally accepted as evidence of conscious thinking in the case of humans is dismissed in the case of animals. Griffin reviews evidence from neuropsychology, fieldwork in animal ethology, and experimental work where animals have been trained to communicate.

The evidence from neuropsychology is arresting. Human consciousness does not appear to be associated with any neural structure or function unique to the human brain, but rather with widely distributed, but coordinated activity engaging large areas of the brain. There is no obvious reason why activity of this sort must be limited to brains of the highest complexity. In any case, the close similarity between animal and human nervous systems “means that there is no inherent reason why animal brains cannot produce conscious experiences.”

But there is also positive neuroscientific evidence of animal consciousness, at least in monkeys. Just as the phenomenon of blindsight in humans has been used to clarify the important difference between registering and responding to information from the environment, on the one hand, and being conscious of it, on the other, so, too, the evidence of blindsight in monkeys warrants a similar interpretation. Furthermore, the discovery of “mirror neurons” in monkeys suggests that they sometimes entertain possibilities, thinking about behaviors they could or would like to perform. If this interpretation is correct, it has important consequences. It is very hard (maybe impossible) to understand counterfactual ideation as information processing or as stimulus-response conditioning, and Griffin draws very near to Whitehead’s technical understanding of consciousness when he suggests that mirror neurons, in providing evidence that monkeys sometimes think about what is *not* the case, but possible, provide evidence of consciousness. There is, moreover, direct evidence that monkeys are conscious when attentionally focused on what *is* the case. It is commonly claimed that monkeys, while they may “know” many facts that are important in their lives, do not *know* that they know them. Reflexivity is thus taken to be a necessary condition of consciousness that monkeys supposedly lack. However, an ingeniously designed experiment demonstrates that monkeys are able to know whether they have remembered a particular piece of useful information and to optimize their strategies for getting food in light of this higher-order knowledge.

It cannot be stressed enough how close Griffin’s discussion of consciousness in monkeys comes to Whitehead’s very abstract analysis of consciousness. For Whitehead, experience is conscious in one of two basic cases: (1) when we “feel” the absence of a difference between a thing and

the description it satisfies or (2) when we “feel” the difference between a thing and a description it doesn’t satisfy. To take an arbitrary example (the old cat and the mat), the verbal transcription in the former case would be *the cat is on the mat*; in the latter case, *the cat is not on the mat*. But *the cat is on the mat* is actually abbreviated. Something logically irrelevant, but psychologically crucial has been left out. The verbal analogue of the *consciousness* that the cat is on the mat would have to reflect the actual state of affairs as the absence of a (potential) difference between the thing indicated and its description: the cat is *not* not on the mat. In other words, ideation cannot be conscious unless it involves a counterfactual element. The counterfactual element is denied, but the unrealized possibility of its truth is what makes consciousness of facts possible. To be conscious of a fact is to “experience” that the possibility of its falsehood exists, but is not realized (note that what is experienced is therefore a proposition). Such an experience is possible, according to Whitehead, because there is, in addition to the physical element that supplies the basis of experience, a purely mental element that supplies the necessary modal and logical functions.¹⁰

In the case of Griffin’s monkeys, they apparently understand that their recollection *could be wrong* (possibility), but are confident in specific cases that it is *not wrong* (unrealized possibility), and in other cases they understand that they no longer recall or that the recollection is no longer reliable. In the former case, where error is the unrealized possibility, they are conscious of knowing. In the latter case, where knowledge is the unrealized possibility, they are conscious of not knowing.¹¹

Griffin makes his strongest case on the basis of animal communication. He asks only that we accept as evidence of consciousness in animals what we take as evidence of consciousness in humans. Even conceding a single standard, however, many will deny that animal communication has the requisite parity: it is not symbolic, lacks semantic content, and lacks displacement. (“Displacement” means “convey[ing] information about something displaced in space or time from the situation where the communication takes place.”) Let the reader note that we are here talking about the “decisive clues” alluded to earlier that would warrant an inference from patent behavior (in this case: communication) to the existence of consciousness, which is necessarily something latent. The decisive features communication must have to warrant such an inference appear to be (1) a symbolic character, (2) semantic content, and (3) displacement. For it is precisely the prominence of these three features in human communication that compels us to view it as expressing subjective experiences, and these three features, so it is alleged, are conspicuously absent from animal communication. But, in fact, these features *are* attested. The alarm calls of vervet monkeys convey specific semantic information about the types of predators, not just emotional arousal, and the famous “waggle

dance” of honeybees displays all three critical features. The waggle dance is an elaborate symbolic code, specific enough to convey precise semantic content that is displaced, flexible enough to serve multiple purposes (e.g., finding the best nectar or finding the best location for a new hive), and it involves an extensive exchange of information among dancers leading through reciprocal adjustments to a final group decision.

Griffin concedes that most examples of animal communication, such as the alarm calls of the vervet monkeys, are indeed examples of direct reactions to the current situation in which the animal finds itself, whereas humans “often think and communicate about past occurrences or what may happen in the future.” Consciousness is strongly associated with displacement for the same reason it is strongly associated with counterfactual ideation. It is hard to see how sensitivity and responsiveness to what is *not* present could be a matter of unconscious information processing or stimulus-response conditioning. Finding displacement in the communication of social insects therefore poses significant challenges to conventional assumptions about consciousness.

George W. Shields makes skillful use of the methods and resources of Analytic philosophy to argue for panexperientialism, thereby disarming some of its most self-confident critics—Anglo-American philosophers who think panexperientialism violates basic sureties of logically rigorous and scientifically informed analysis. Sophistication in formal logic and command of hard science often make Analytic criticism formidable. Shields meets this criticism on its own terms, presenting formally rigorous arguments and hard empirical evidence in favor of panexperientialism.

The first part of his paper focuses on logical and philosophical arguments for panexperientialism. Shields examines what panexperientialism means and proposes the following as minimal criteria: that every genuine individual has a physical presence in space-time and is related internally to its environment.¹² Whitehead’s analytic unit of experience—the prehension—is therefore an internal relation. Following Hartshorne, Shields argues that internal relations translate logically into strict implications. From this analysis he infers what the denial of panexperientialism amounts to: an ontology of exclusively external relations and a logic of entirely open possibilities. There would be no restrictions on the conjunction or separation of individuals in this world (this is what is meant by “open possibilities” in this context). The only necessity would be the completely symmetrical logical necessities of identity and noncontradiction.

Shields is happy to continue the strain or argument begun by David Griffin when he analyzed the idea of performative consistency in terms of “hard-core common sense.” Shields calls “assumptions which we presume in our practice universally or nearly universally” “deep protocols of common sense” and under this rubric extends Griffin’s list to include four more items:

“that (1) our experience as temporally conceptualized into ‘past,’ ‘present,’ and ‘future’ is coherent; (2) the act of remembering is in principle not the same as the act of imagining; (3) causal influence is objectively real; and (4) a ‘skeptical solipsism of the present’ is false.” Shields presents what he calls a “reduction to pragmatic absurdity” by showing that the external relations ontology violates these four common sense commitments. Since the relation between any two events is wholly external and contingent it becomes impossible to understand how the present could have anything to do with the past. A present event, such as the act of remembering the past, would necessarily be independent of anything that actually happened in the past, undercutting the concepts of “past,” “memory,” and “causal influence” at one stroke and sealing present consciousness hermetically in the present moment. Panexperientialism, by accepting internal relations, avoids these difficulties. Shields notes that panexperientialism also avoids the cardinal problems typically afflicting materialism and dualism, namely, the emergence of qualia out of matter and the possibility of the mind acting causally on matter.

Shields considers and rebuts in some detail six objections raised against panexperientialism: (1) that it implies that things like rocks have thoughts (only it doesn’t), (2) that it implies that the behavior of the elementary constituents of matter would not be predictable through their physical properties alone (only they aren’t), (3) that elementary particles are completely identical whatever their past histories and thus could not have any interior states (only they aren’t and so could), (4) that attributing any kind of feeling, however qualified or attenuated, to micro-constituents of matter violates the linguistic protocols for meaningful use of terms such as feeling (but not just the concept of feeling is generalized—the criteria for its attribution are generalized as well, yielding predicates such as “openness to the environment” or “internal relatedness,” which are still “psychological” predicates without being strained usage), (5) that any adequate physicalism must be tantamount to epiphenomenalism (but epiphenomenalism cuts *against* physical science because it is anti-evolutionary, implying that “animals and humans evolved with persistent natural selection of entirely superfluous mental entities”); and (6) that the existence of unproblematic forms of emergence, such a liquidity from molecules, shows that proto-experiential “elements” are not needed to explain the emergence of the experience we are familiar with (but experience is unlike liquidity in the relevant respect because the latter is a kind of emergence that can be understood and predicted from its antecedent elements, while experience notoriously cannot).

The second part of Shields’ paper looks at arresting empirical data in support of panexperientialism. In pushing down the lower threshold of “conscious” experience, Donald Griffin got us to social insects. Velmans, arguing that even single cell organisms might have some kind of phenomenal

awareness, got us to the very bottom of the phylogenetic tree. Shields, drawing on startling but well-confirmed empirical findings, provides the final turn of the screw that anchors experience in the Planck-scale units of nature. He cites first the phenomenon of neuroplasticity. Although it had long been dogma that the brain is hardwired once and for all in early childhood, recent research has documented ongoing alterations in the adult brain, including the growth of new neurons, as a result of sensory and cognitive input. Important to Shields' argument is the remarkable discovery that the brain's plasticity is susceptible to clinical manipulation. Attentional therapies such exercises as "observing" undesirable thoughts and emotions in an impartial manner and then refocusing attention repeatedly on alternative thoughts, have not only been shown to work, but PET scans have now revealed alterations of the brain's neural system corresponding to the behavioral changes. Shields argues that this is just a special case of the weird but documented quantum phenomenon known as the "Quantum Zeno Effect," where observation increases the probability that a given quantum state will *not* change: "the more frequently and rapidly you observe a physical system in a certain selected way, the more you 'lock in' a certain physical state of the system." The power of attentional therapies to decrease the probability of unwanted thoughts and emotions would thus result from "locking in" the alternatives by repeated ideational exercises. Shields' appeal to quantum mechanics is not entirely speculative since, as he points out, the release of neurotransmitters is regulated by processes so microscopic that quantum mechanical principles do indeed apply. The provocative conclusion is that attention to one's own thinking, like the observation of experimental setups, has the power to alter the probability that one rather than another superimposed wave function will be actualized. This kind of "top-down" causal influence is precisely what Whitehead's panexperientialism is designed to explain. Shields doubts that classical materialism can make any sense of these phenomena at all.¹³

Max Velmans also takes up the critical question of how far downscale in complexity types of individuals can be found that still have some kind of experience. Velmans approaches the question in the context of a larger question about the evolution of consciousness. Which sorts of entities are thought to have consciousness determines to a large extent when consciousness must have evolved and what biological refinements to the evolving organism are specifically responsible for it. Velmans notes that theories about the distribution of consciousness range from the ultraconservative (only humans have it) to the extravagantly libertarian (everything has it—panpsychism). While ultraconservative theories traditionally drew their support from theology, more contemporary versions "are based on the supposition that higher mental processes of the kinds unique to humans are necessary for consciousness of any kind." Velmans is skeptical not only of these ultraconservative views, but

ultimately of any degree of conservatism on this point. He examines a variety of rationales for the claim that consciousness depends on brain complexity or higher order cognitive processes and argues that “[s]uch views confuse the necessary conditions for the *existence* of consciousness with the added conditions required to support its many *forms*.”

While Velmans ultimately endorses the “extravagantly libertarian” view, it must be noted that his vocabulary does not strictly conform to the usage of Whitehead or contemporary Whiteheadians. Although he does distinguish between conscious and unconscious information processing (and elsewhere [2000] between conscious and unconscious *mind*), he does not make the terminological distinction between conscious and nonconscious experience or between panpsychism and panexperientialism. Nevertheless, he makes nearly equivalent distinctions by stressing the widely differing degrees of complexity manifested by the forms that consciousness takes: from mere feeling to conceptually articulated consciousness of self and world. At the lowest extreme, consciousness in Velmans’ acceptation is tantamount to Whitehead’s nonconscious experience. Accordingly, the form of panpsychism he advocates is very close to panexperientialism. This becomes evident when we revisit his distinction between conscious and unconscious information processing in light of his final reflections on the nature of focal-attentive consciousness. As we discuss below, Velmans marshals an arresting reason why unconscious information processing may simply be information processing in which a diffuse, primitive consciousness has been suppressed. So in the end, feeling may well be a naturally occurring feature of all biological processes, which regulate themselves through information extracted from their internal and external environments.

Velmans advances two principle arguments against conservative distribution theories. First, following a tradition that includes Thomas Huxley and Charles Sherrington, he points out that conservative distribution implies a discontinuity theory of evolution. At some point consciousness must “appear [. . .] (out of nothing) through some random mutation in complex life forms that happen[s] to confer a reproductive advantage.” Typically, it is thought that consciousness is linked to the evolution of the neocortex. However, there is nothing unique to cortical cells that might be responsible for consciousness. Indeed, as Sherrington observes, cells in the frog embryo destined to be brain can often be replaced with others, such as skin cells from the back, and still develop into brain. This leaves us with the assumption that consciousness must have something to do with neural *organization*. But the strong evidence for the gradual evolution of the human brain makes it unlikely that consciousness sprang fully formed at any point in the brain’s slow accretion of structural/functional complexity.

Velmans’ second argument against conservative distribution theories is, by contrast, so untraditional that it upends conventional objections to

panexperientialism. Velmans brings neurophysiology and phenomenology together in a startling way. He reminds us that a great number of the synapses in the brain must be inhibitory. Otherwise the nervous system would be in a constant state of universal excitement after the first signal. At the same time we know that consciousness would be impossible if the vast amount of simultaneous information streaming in to the mind/brain were not limited and filtered down to something that it could manageably attend to.¹⁴ From this Velmans infers the possibility that consciousness may be a naturally occurring feature of all neural representations. However, the more complex the nervous system, the more necessary it would be to *inhibit* consciousness of all but the most important information to prevent overload and confusion. In this case, rather than adding something to unconscious representations to make them conscious, attention¹⁵ would correspond to a highly selective release of consciousness from inhibition.

The implications of this argument are dramatic. As Velmans notes, cognitive psychology has demonstrated that most human information processing takes place unconsciously. Naturally, this leads cognitive psychology to seek the specific conditions that distinguish conscious from unconscious processing. It asks, in other words: why does consciousness emerge at some particular threshold in the mind/brain's cognitive functioning? This question is perhaps no less vexed than the question of the threshold at which consciousness emerges in the course of evolution. Both questions presuppose a discontinuity: on the one hand, a discontinuity in the evolution of consciousness (a diachronic discontinuity), on the other, a discontinuity in the distribution of consciousness (a synchronic discontinuity). The synchronic discontinuity takes two forms. There is the discontinuity between organisms that do and those that do not have consciousness, and, within the nervous system of organisms that do have consciousness, there is the discontinuity between conscious and unconscious processes. These two kinds of synchronic discontinuity are closely related. Without having to make any particular assumptions about the relation between organic processes *available* to consciousness and organic processes that *result* in consciousness, we can nonetheless say it is only because consciousness does not extend (in either sense) to the vast majority of organic processes in our own brains and bodies that we resist the idea that very similar processes in other life forms might be conscious (in either sense).

But if what human beings normally experience as consciousness is only the selective release of an aboriginally pervasive consciousness from its systemic inhibition in complex nervous systems, then unconscious information processing would not be different in kind from—that is, discontinuous with—conscious processing. Unless specifically inhibited, a kind of rudimentary “consciousness” would attach to all organic information processing. Neither

at some point in evolutionary history nor at some level of neural activity would consciousness come into being *de novo*.

Of course this rudimentary consciousness would not involve attention or any kind of reflective, objectifying, or thematically motivated awareness. It would be a diffuse, nonconceptual, nonobjectifying, and nonreflective feeling of qualia—just the sort of experience Whitehead calls “unconscious” and generalizes to all events in nature. On this telling, discontinuity—the abrupt emergence of consciousness from something unconscious—would be an illusion arising from attentional consciousness, the necessary flip side of which is the suppression or exclusion of diffuse consciousness. The highly restricted access of our own consciousness to what it’s like to be a living organism then leaves us in a poor position to appreciate what we have in common with less complex forms of life.

We see how this argument leads to a new opening in the vexed problem of the evolutionary value of consciousness. It would not be consciousness, but complexity of information processing that confers a reproductive advantage on certain organisms in certain environments. Such complexity would go hand in hand with increasing complexity of the attendant consciousness. But increasing complexity in the nervous system would actually become counter-productive for the organism unless the attendant consciousness of all but a narrow selection of the increasingly diffuse aggregate experience embraced in this complexity was suppressed, thus yielding our familiar attentionally focused consciousness as well as the illusion that it is the addition of something altogether new on top of an otherwise unconscious cognitive processing.

For Velmans, an evolutionary account of human consciousness is therefore possible, but only by upending the way the question is usually posed. It would not be consciousness so much as its suppression that under certain circumstances confers a selective advantage. Attentional consciousness evolves when diffuse consciousness begins to pose a selective *disadvantage*. We should note in passing how close this comes to Whitehead’s thesis that complexity of experience, which he calls width and depth through harmony and intensity, requires a great deal of the data actually given to an entity to enter into its experience only negatively, that is, as something excluded, suppressed, diminished, or transmuted.

This unanticipated turn of the argument stands Thomas Nagel’s reservations on their head. What needs to be explained is not how there possibly could be simple phenomenal consciousness at the low end of the phylogenetic tree, but why it is lacking throughout most of the nervous system of organisms at the high end. “[P]henomenal consciousness (of any kind) might only require representation. If so, even simple invertebrates might have some rudimentary awareness, in so far as they are able to represent and, indeed, respond to certain features of the world.” Empirically, it is as yet

impossible to rule out even more remote seeming possibilities: "If the ability to represent and respond to the world, or the ability to modify behavior consequent on interactions with the world are the criteria for consciousness then it may be that consciousness extends not just to simple invertebrates (such as *Planaria*) but also to unicellular organisms, fungi and plants." The upshot of the continuity theory of the evolution of consciousness is essentially panexperientialist: "In the cosmic explosion that gave birth to the universe, consciousness co-emerged with matter and co-evolves with it. [...] On this view, evolution accounts for the different *forms* that consciousness takes. But, consciousness, in some primal form, did not emerge at any particular stage of evolution. Rather, it was there from the beginning. Its emergence, with the birth of the universe, is neither more nor less mysterious than the emergence of matter, energy, space and time."

Gregg Rosenberg also focuses on the case for panexperientialism, but he takes a very different approach from our other contributors. In a precisely executed analysis of the semantic structure of explanatory theories and of causality in particular, Rosenberg surprises us with an altogether original argument for panexperientialism. Rosenberg observes that Whitehead's panexperientialism is "a reaction to the void created by his rejection of Vacuous Actuality," which he explains with admirable lucidity: "A Vacuous Actuality would be a fundamental reality that is purely structural and quantifiable, with no intrinsic nature of its own that escapes the formal description of a pattern. The rejection of Vacuous Actuality amounts to the assertion that the entities of fundamental physics, for instance, are more than mere dynamic quantities, mere information structures in the vacuum. It is the rejection of the now popular information-theoretic 'It from bit' view for understanding the essential nature of the physical world."

Rosenberg notes that the rejection of Vacuous Actuality and the endorsement of panexperientialism by process philosophy look on the surface "like positions of insight, or even faith, not sufficiently motivated by argumentation." The purpose of his paper "is to put more argumentation in place to support the rejection of Vacuous Actuality and the panexperientialist reaction to that rejection."

Rosenberg begins with a logical analysis of different kinds of relations. He is especially interested in the kind of relationships that define the explanatory structure of scientific theories. He illustrates how they are typically conceptual relations in which the *relata* mutually presuppose one another. For example, in economics, goods and services are things that consumers and producers barter. But consumers and producers are, in turn, simply people occupying distinct positions in the system of bartering goods and services. In biology, a heritable characteristic (gene) is one that parents pass from their generation to the next, but a parent is an organism that passes