

Minding the Body

Robert Hanna

University of Colorado at Boulder, USA

If what[ever] consciousness [there is] spreads all over human bodies, then there won't be any temptation to use the [Cartesian] word 'ego'.

--L. Wittgenstein¹

I. Introduction

Precisely how and precisely where is human conscious experience located in the natural world? The now-familiar although still quite controversial *Extended Mind Thesis* says that there are extra-bodily vehicles of mental content. By contrast, the less familiar and very controversial *Extended Conscious Mind Thesis* says that the constitutive mechanisms of human conscious experience include both extra-neural bodily facts and also extra-bodily worldly facts. If The Extended Conscious Mind Thesis is correct, then not only is human conscious experience spread out all the way through the living human animal body, it is also spread out into the external natural and human social world itself.

In a recent paper, Andy Clark argues that

the case for [The Extended Conscious Mind Thesis] is at best unproven and that the machinery of [human] conscious experience is (probably) all in the head/CNS.²

This is a cautiously stated claim. I will re-formulate it as saying that because the arguments currently on offer for The Extended Conscious Mind Thesis fall short of decisive proof, then, all things considered, we should conclude that *the constitutive mechanisms of human consciousness are all neural*, that is, that the constitutive mechanisms of human conscious experience are all either in the brain or the central nervous system (CNS). Let us call this *The Cautious Consciousness-Is-All-Neural Thesis*.

I agree with Clark that The Extended Conscious Mind Thesis is (probably) false; but I also think that there is a sufficient reason for rejecting The Cautious Consciousness-Is-All-Neural Thesis. This sufficient reason, in turn, is the conjunction of two perhaps surprising theses about the nature and location of the constitutive mechanisms of conscious experience—(1) *The Essential Embodiment Thesis*, which says that human consciousness is *essentially embodied*, in the two-part sense that

(1a) necessarily, human consciousness is embodied,

and

(1b) necessarily, human consciousness has a full-scale neurobiological incarnation of all its states in all the conscious human animal's vital systems and vital organs—including the higher brain, brain stem, limbic system, nervous system, endocrine system, immune system, and cardiovascular system, right out to the skin, but no further than that.

—and (2) *The Deep Consciousness Thesis*, which says that

necessarily, whenever a conscious human animal is in *any* sort of mental state, then it is *also* occurrently conscious in some definite way, even if only minimally.

So according to The Deep Consciousness Thesis, occurrent human consciousness penetrates into *every* aspect of our mental lives, *including* so-called “non-conscious” or “sub-personal” information processing. The Deep Consciousness Thesis also entails a crucial distinction between

(i) rational human self-consciousness, or *reflective* human consciousness,

and

(ii) proto-rational human nonself-conscious consciousness, or *pre-reflective* human consciousness.

If The Essential Embodiment Thesis and The Deep Consciousness Thesis are both correct, then rational human self-consciousness or reflective consciousness, nonself-conscious or pre-reflective human consciousness, and the neurobiological life of the

suitably complex living human animal all mutually determine each other and jointly constitute one and the same thing—*us*, just minding our own bodies.

II. Clark, The Extended Mind, and The Not-So-Very-Extended Conscious Mind

Clark argues for The Cautious Consciousness-Is-All-Neural Thesis in four steps. **First**, he distinguishes between The Extended Mind Thesis and The Extended Conscious Mind Thesis.³ **Second**, he argues that the three basic arguments offered by contemporary radical “enactivist” defenders of The Extended Conscious Mind Thesis—

(i) the argument from “sensorimotor loops and variable neural correlates,”

(ii) the argument from “virtual representations,”

and

(iii) the argument from “dynamic entanglement”—all fall short of decisive proof for that Thesis.⁴

Third, he argues that appeals to brains-in-a-vat and other similar scenarios by orthodox “in-the-head” opponents of The Extended Conscious Mind Thesis are inadequate.⁵

Fourth and finally, he argues that there is compelling evidence for a principled distinction between the neural mechanisms of consciousness, which are constitutive, and the merely causal supports and triggers of consciousness, which includes the extra-neural living body and the external world. This part of Clark’s argument presents the bottom line, so I will quote it at some length:

[I]f indeed the physical machinery of conscious experience requires fast timescale operations and processing, and the non-neural body acts as a low pass filter preventing external (and internal, muscular) signals from directly entering into such operations and processing, then such signals are fit to play only a causal role, driving the neural systems within which the right kinds of fast binding and processing can occur. In such cases, one might have all manner of complex couplings without thereby producing an extended material basis for conscious experience. Contrast the case (discussed at length in Clark, 2007, 2008) of the possible role of gesture in the process of reasoning. There seems no reason why slow timescale gestural events should not productively interact with faster timescale neural ones so as to yield a special kind of coupled gestural-neural unfolding that is *itself* the distinctive physical engine of a certain kind of problem-solving. But within this coupled unfolding, the streaming contents of conscious experience would all depend constitutively only on the neural processing itself. The account on offer thus enables us to embrace

the kinds of claims made by Noë and others to the effect that certain experiences may only come about due to the neural systems being driven, in some distinctive way, by external signals. But it does so without being forced to the conclusion that such external sources comprise part of the most local machinery that generates the conscious experience itself. The account thus offers a principled reason for making the causal/constitutive cut, in the special case of conscious experience, in an orthodox, non-extended, kind of way.⁶

I agree completely with the first three steps of Clark's argument. Unlike Clark, however, I *also* believe that The Extended Mind Thesis is false, for reasons carefully worked out by Robert Rupert.⁷ But the falsity of The Extended Mind Thesis, of course, has no direct bearing on (what I take to be) Clark's correct claim that The Extended Mind Thesis does not entail The Extended Conscious Mind Thesis. The Extended Conscious Mind Thesis is arguably false even if The Extended Mind Thesis were true. It is the fourth and final step of Clark's argument that I want to take issue with.

More precisely, what I want to take issue with follows directly from Clark's view that

if indeed the physical machinery of conscious experience requires fast timescale operations and processing, and the non-neural body acts as a low pass filter preventing external (and internal, muscular) signals from directly entering into such operations and processing, then such signals are fit to play only a causal role, driving the neural systems within which the right kinds of fast binding and processing can occur.⁸

I grant Clark the claim that the physical machinery which requires fast timescale operations and processing does indeed constitutively enter into human conscious experience. Let us call this kind of human conscious experience, *type-1 human conscious experience*. This type of human conscious experience, in rational humans, also inherently involves the capacity for self-consciousness or reflection, so type-1 rational human conscious experience is also *human self-conscious or reflective experience*. But here's the rub: How does Clark rule out the nontrivial possibility that there is nevertheless also a *distinct kind* of human conscious experience, which I will call *type-2 human nonself-conscious or pre-reflective conscious experience*, happening in the extra-neural living

human animal body, and occurring at slower timescales, *in addition to* the type-1/self-conscious or reflective rational human experience that is happening in the brain and CNS?

More positively put now, my basic claim against Clark is that basic levels of human mental activity and representation generally assumed to belong to “the cognitive unconscious”⁹ in normal, healthy, mature rational human animals are in fact essentially nonself-consciously or pre-reflectively *conscious*. If I am correct, then human conscious experience goes all the way down to the sensorimotor ground floor of cognitive and practical agency, via the vital cord of human organismic life. Most precisely put: If The Essential Embodiment Thesis and The Deep Consciousness Thesis are both true, then Clark’s Cautious Consciousness-Is-All-Neural Thesis and also The Conscious Extended Mind Thesis are both false.

Before I get to my arguments, I want to take pre-emptive action against two possible worries about The Essential Embodiment Thesis and The Deep Consciousness Thesis.

The **first** possible worry is that because some rational human animals have “out-of-body” conscious experiences, or OBEs¹⁰—e.g., of floating high in the sky above their own bodies, looking down at them—then it follows that The Essential Embodiment Thesis and The Deep Consciousness Thesis are both false.

But, clearly, not only by hypothesis but also empirically, the evidence universally shows us that “out-of-body” conscious experiences are non-veridical illusions that occur only in *conscious rational human animals*. And in any case, the conscious experience of floating high in the sky above one’s own body, looking down at it, is not a conscious

experience of *disembodiment*. Indeed, if I am correct, then there really is *no such thing as a conscious experience of disembodiment*. On the contrary, an “out-of-body” conscious experience is simply the non-veridical illusory conscious experience of having a strange new birdlike, light airy body that floats in the air above one’s old non-birdlike, heavy earthbound body. Dreams, fantasies, or hallucinations of bodily liberation or “astral projection” are therefore phenomenologically *misdescribed* as conscious experiences of disembodiment: for even birdlike, light airy bodies are still Prisoners of Gravity.

The **second** and much more important possible worry is that The Deep Consciousness Thesis might seem to imply what Clark calls a “causal-constitutive” or “coupling-constitutive” fallacy—i.e., the fallacious inference “that just because X [causally] drives Y, X becomes partially constitutive of Y”¹¹—essentially similar to the basic error highlighted by Clark in the standard arguments for The Extended Conscious Mind Thesis, which “merely conflate causal drive and essential phenomenon-producing machinery.”¹²

But properly understood, The Essential Embodiment Thesis and The Deep Consciousness Thesis do *not* imply a “causal-constitutive” fallacy. On the contrary, the basic arguments for The Essential Embodiment Thesis and The Deep Consciousness Thesis, as we will very shortly see, have the following non-fallacious form:

“X (= the suitably neurobiologically complex living human animal body as a whole) is partially constitutive of Y (= human consciousness, whether type-1/self-conscious or reflective or type-2/nonself-conscious or pre-reflective) because

- (i) there are independent good arguments for the existence of Y, and
- (ii) there are independent good arguments for X’s essentially belonging to the causal substructure of Y.”

Here, now, are those arguments.

III. The Essential Embodiment Thesis

What is an animal? The Oxford English Dictionary tells us that the word ‘animal’ means “a living organism which feeds on organic matter, usually one with specialized sense organs and nervous system, and able to respond rapidly to stimuli.”¹³ In biology on the other hand, ‘animal’ has a more technical meaning, in that animals constitute one of the five kingdoms of living things: Monera (bacteria), Protists, Fungi, Plants, and Animals. The class of animals in this biological sense includes both vertebrates and invertebrates.

My usage of ‘animal’ in this essay, however, is a precisification of the ordinary language and scientific terms, intended to coincide with its use in cognitive ethology.¹⁴ To signal this precisification, I have coined the quasi-technical term *minded animal*. Minded animals are animals with *consciousness like ours* and *minds like ours*—or for terminological convenience from now on, “consciousness_{lo}” and “minds_{lo}.”

In a recent book, I have argued that necessarily every creature with a consciousness_{lo} is an *essentially embodied mind*.¹⁵ Essentially embodied minds are the same as minds_{lo}. An essentially embodied mind, or mind_{lo}, in turn, is an irreducible consciousness_{lo} that is also necessarily and completely neurobiologically embodied. This is to say that its irreducible consciousness_{lo} essentially *cannot* be disembodied, and that it thereby has a *full-scale neurobiological incarnation* of its conscious states in all its vital systems and vital organs—including the higher brain, brain stem, limbic system, nervous system, endocrine system, immune system, and cardiovascular system, right out to the skin. Furthermore every consciousness_{lo}, as the consciousness of an essentially embodied mind, or mind_{lo}, is fundamentally manifest as *desire-based emotion*, and, in particular, as

effective desiring—which is the kind of desiring that is also a *trying* that causes intentional action.¹⁶ So essentially embodied minds, or minds_{lo}, are always poised for trying to *do* something, and thereby always have a capacity for intentional agency.

The crucial idea of essential embodiment needs to be further elaborated. To say that every animal that has a consciousness_{lo} thereby has a full-scale neurobiological incarnation of its irreducibly conscious states in all its vital systems and vital organs—including the higher brain, brain stem, limbic system, nervous system, endocrine system, immune system, and cardiovascular system—right out to the skin, as I noted in section I, is what I call The Essential Embodiment Thesis. It is important to note that The Essential Embodiment Thesis has two logically distinct parts:

(1) the *necessary* embodiment of conscious minds_{lo} in a living organism (The Necessity Thesis),

and

(2) the *complete* neurobiological embodiment of conscious minds_{lo} in all the vital systems, vital organs, and vital processes of our living bodies (The Completeness Thesis).

The Necessity Thesis says that necessarily, conscious minds_{lo} are alive.

Negatively formulated, it says that conscious minds_{lo} cannot be dead, disembodied, or purely mechanical.

By contrast, The Completeness Thesis says that conscious minds_{lo} are fully spread out into our living bodies, necessarily including the brain, but also necessarily not *restricted* to the brain. Please note that I am *not* saying that our brains, hearts, livers, or stomachs are conscious. On the contrary, according to my view it is only *whole animals* that are conscious, not their body parts alone, and not even their brains alone. So what I am saying by asserting The Completeness Thesis is that the minded animal as a whole—

e.g., a rational human animal—is conscious *with*, or *in-and-through*, its brain, heart, liver, stomach, or whatever, right out to the skin..

One could, at least in principle, assert the Necessity Thesis and also reject the Completeness Thesis. But I want to assert both of them together. So I hold that the supposed consciousness of a causally detached brain—say, a living brain floating listlessly in a vat, as in Hilary Putnam’s famous thought-experiment¹⁷—even though it seems both conceivable and logically possible, just would not be a consciousness *like ours*. On our view, a consciousness_{lo} necessarily involves a brain that is causal-dynamically coupled with all the other vital systems, organs, and processes of our living body.

The notion of a “causal-dynamic coupling” is crucial. The Necessity Thesis and The Completeness Thesis do *not* jointly entail that consciousness_{lo} actually is or ever could be embodied in *any* causally necessary condition of our kind of consciousness, which would of course include all sorts of entities and facts outside our living bodies. That is what I call *The Embodiment Fallacy*.¹⁸ Instead The Necessity Thesis and The Completeness Thesis jointly entail that consciousness_{lo} is embodied *only* in a special kind of fully integrated dynamic system that is both causally necessary and causally sufficient for consciousness_{lo}—namely, one that has *all the same causal powers as* the vital systems, organs, and processes of our living bodies. Any such living body is the natural matrix, or natural basis, of a consciousness_{lo}.

And that point in turn raises another extremely important point that is specifically about the very idea of a “natural matrix” of a consciousness_{lo}. A natural matrix of a consciousness_{lo} is not merely a *compositional material substrate*—a mass of physical

stuff and a collection of physical parts—that necessarily accompanies and supports consciousness_{lo}. A natural matrix is instead a system of causal-dynamic relations, embedded in some or another compositional material substrate, awaiting specific activation or actualization. This means that if you significantly modify the shape of your body, or lose a limb or some other body part, *without also replacing it with an equivalent counterpart that has the same relational causal powers*, then you would also correspondingly modify or lose your mind. But the *specific bodily stuff* and the *particular body parts* are not metaphysically important. The mere matter doesn't really matter.

In *Meditations* VI, and while auspiciously wearing his Substance Dualism hat, Descartes makes a similar point:

Although the whole mind seems to be united to the whole body, I recognize that if a foot or arm or any other part of the body is cut off, nothing has been thereby taken away from the mind.¹⁹

But my *reason* for making this point is radically different from that of Descartes. In his Substance Dualist guise, Descartes holds that the mind is an absolutely homogeneous and simple unity, and thereby indivisible. *My* point, however, is about the metaphysics of living animal bodies like ours, not about the metaphysics of mental substance. Again, what I hold is that the natural matrix of consciousness_{lo} is *not* just a hunk of specific bodily stuff and *not* just a heap of particular bodily parts. Instead the natural matrix of a consciousness_{lo} is all the vital systems, organs, and processes of our living bodies, *as individuated by their relational causal powers*, that is, by what they can efficaciously do in causal community with each other and with the larger surrounding world. That these vital systems, organs, and process are in fact composed of some or another hunk of specific bodily stuff and also of some or another heap of particular bodily parts—say, specifically human body stuff and particular human body parts—is of course extremely

practically important for members of the relevant species made out of that stuff and those parts, but otherwise it is metaphysically trivial. Thus The Essential Embodiment Thesis is a thesis about the *operative neurobiological dynamics* of creatures with consciousness_{lo}, and not (except trivially) a thesis about our compositional material substrate.

Assuming, then, that The Completeness Thesis is formulated in terms of the relational causal powers of the vital systems, organs, and processes of our living bodies, and not (except trivially) in terms of their compositional material substrate, there are at least four good reasons for defending The Essential Embodiment Thesis.

First, it seems obvious that if any of the vital systems, organs, or processes in our bodies is destroyed or permanently disabled without a functional replacement that has essentially the same relational causal powers—say, an artificial heart, a liver transplant, etc.—then our consciousness will cease to exist, precisely because the whole organism *dies*. Therefore the *existence* of consciousness_{lo} necessarily depends on its complete neurobiological embodiment.²⁰

Second, it seems equally obvious that significant changes made to the relational causal powers of any of our vital systems, organs, or processes normally produce correspondingly significant changes in the specific character of the conscious mind_{lo}. And this is as true of the *non-brain* systems as it is of the *brain* systems. A thyroid gland malfunction, hormone imbalance, adrenaline surge, or heart attack is apt to cause highly significant changes in consciousness_{lo}. Therefore the *specific character* of consciousness_{lo} also necessarily depends on its complete neurobiological embodiment.²¹

To be sure, a lobotomy or a concussive blow to the head is apt to cause *more* fundamental changes in consciousness_{lo} than a thyroid malfunction, hormone imbalance,

and so-on. And again, to be sure, the brain is centrally causally involved in every aspect of normal attentive, singly-focused, alert, self-reflective, waking consciousness_{to}. So we are not in any way denying the necessary and central causal role of the brain in the constitution of normal attentive, singly-focused, alert, self-reflective, waking consciousness_{to} and intentionality. But at the same time, I am also strongly recommending that philosophers of mind and cognitive neuroscientists should not *overemphasize* the causal role of the brain,²² to the extent that this undermines our recognition of the equally necessary role of the relational causal powers of the rest of our vital systems, organs, and processes. For example, as everyone knows, even fairly minor changes in our *digestive* processes can produce non-trivial changes in consciousness_{to}.

Think, for example, of the striking phenomenological differences between:

- (a) feeling very hungry and craving a plate of spaghetti,
- (b) feeling as if you ate just the right amount of spaghetti,

and

- (c) feeling utterly stuffed with spaghetti.

The brain obviously is centrally causally involved in these normal attentive, singly-focused, alert, waking phenomenological differences, but it seems also equally obvious that the brain does not in and of itself causally *control or determine* these differences. On the contrary, it seems obvious that the “enteric brain”—our guts—is doing much of the causally controlling and determinative work here.²³ And similar points can be made about the other non-brain vital organs, systems, and processes. Each of them can and does play a causally controlling and determining role with respect to some differences in normal attentive, singly-focused, alert, self-reflective waking consciousness_{to}, even if the brain is also centrally causally involved.

Analogously, even if every basic act of a corporation passes directly through its Chief Executive Officer, it does not follow that the CEO controls or determines the specific character of *every* such act, or even *most* of them. In fact, in a great many cases the CEO is just *the chief executive slave* of the controlling determinations of the shareholders (if it is a public company), or of the employees (if it is either an employee-owned company or unionized), or of the actual business operations of the company. So too the brain is often just *the central causal slave* of the rest of the living body.

Third, there is neurophysiological empirical evidence that supports The Essential Embodiment Thesis. For example, recent work on the neurochemistry of human emotions strongly suggests that the vital systems centrally causally involved with and embodying our basic emotions are gut-based, not brain-based.²⁴

But **fourth** and finally, probably the most compelling empirical evidence for The Essential Embodiment Thesis, precisely because it is the simplest, is the well-known fact that the “arc” of reflex action (say, someone’s pulling her hand away from something very hot) operates more quickly than the time it takes for the brain to process information sent to it via the nervous system about the body parts involved in that reflex action (say, that the subject’s hand has been seriously burned). If I am correct about The Deep Consciousness Thesis, as I will argue in section **IV**, then this is *also* a human conscious experience, although not of course a self-conscious or reflective experience—in the example of the burned hand, the subject’s hand moves *before* she self-consciously or reflectively feels the searing pain of a burn. But I do think that reflex action still has a special *phenomenology*, in the classical Nagelian senses that there is a definite something-it-is-like-to-be, and a particular point of view, for a suitably neurobiologically

complex living organism like us when, e.g., that organism is pulling her hand away from something very hot even though the self-conscious or self-reflective awareness of the searing pain of the burn has not yet emerged. Reflex action necessarily includes a proto-rational, first-order, reflexive, nonself-conscious or *pre-reflective* consciousness_{lo}, even if it does not necessarily include a rational, higher-order, and self-conscious or *self-reflective* consciousness_{lo}.

In addition to this point about pre-reflective consciousness_{lo}, a further reason to think that *reflex action* is indeed reflexively conscious, although in a pre-reflectively conscious way, is that it is possible to train oneself, through biofeedback strategies, to modulate or even suppress such reflexes.

So if this point about pre-reflective consciousness_{lo} is correct, and if we also take biofeedback data seriously, then necessarily, in cases of reflex action a pre-reflective consciousness_{lo} occurs *with* and *in-and-through* the vital systems that constitute and subserve our intentional body movements, even though by hypothesis the brain is *not* centrally causally involved in the production of these pre-reflectively conscious intentional actions. Or in other words, there is compelling empirical evidence that there is a necessary and complete neurobiological embodiment of consciousness_{lo} even when the brain is only *peripherally* causally involved. Moreover, there is a direct metaphysical pay-off from this conclusion. As W.T. Rockwell puts it:

[I]f the brain does not record certain features of a perception that the mind is nevertheless aware of, this must mean that the mind is not identical with the brain.²⁵

But this is not Dualism. For Rockwell, and also for me, the conscious mind_{lo} is not identical to the brain, and thus the conscious mind_{lo} is not reducible to the brain, not because the conscious mind_{lo} is in any way metaphysically separable either from the

brain or from the vital systems of the living body more generally, but instead just because the embodiment of consciousness_{lo} *goes much further out into the living body than the brain*, and indeed all the way out to the skin—but no further than that.

IV. The Deep Consciousness Thesis

As I noted in section I, the Deep Consciousness Thesis says that

Necessarily, whenever a conscious human animal is in *any* sort of mental state, then it is *also* occurrently conscious in some definite way, even if only minimally. So occurrent human conscious experience penetrates into *every* aspect of our mental lives, *including* so-called “non-conscious” or “sub-personal” information processing.

If The Deep Consciousness Thesis is true, then it provides the beginnings of a solution to what Ray Jackendoff aptly calls *The Mind-Mind Problem*, which is how it is ever possible for there be genuine two-way causal or semantic interaction across the theoretical and normative gap between the Conscious Mind (or “first-personal” information processing) and the Computational Mind (or so-called “non-conscious” or “sub-personal” information processing)²⁶ My solution to The Mind-Mind Problem, via with The Deep Consciousness Thesis, is that so-called non-conscious or sub-personal mental processing is still in fact *first-personal, conscious, proto-rational, and normative* mental processing even though it is *pre-reflective* and *nonsel-conscious*.²⁷

This doctrine may seem shockingly unorthodox. But properly understood, it is much less shocking than it may seem. One fundamental source of philosophical confusion in this area is that the very idea of consciousness_{lo}, or “the first-personal,” is deeply ambiguous as between

(1) *self-consciousness_{lo} or reflection,*

which is the capacity of a rational, conscious creature like us to have conscious propositional/conceptual meta-representational states, or self-describing conscious thoughts about itself, and what Evan Thompson aptly calls

(2) *sensorimotor subjectivity*,²⁸

which is the more primitive, proto-rational, and nonself-conscious or pre-reflective capacity of conscious, suitably neurobiologically complex living organisms like us to have what Nagel also aptly calls a “single point of view.”²⁹ In turn, I hold, this nonself-conscious or pre-reflective capacity of a conscious living organism like us to have a single point of view is grounded in egocentrically-centered essential embodiment, and a *primitive bodily awareness* that includes proprioception (the sense of the relative positioning of one’s own body parts and limbs, at rest or in movement), kinaesthesia (the sense of bodily movement), the senses of orientation and balance (as intrinsic aspects of proprioception or kinaesthesia), bodily pleasures and pains, tickles and itches, the feeling of pressure, the feeling of temperature, the feelings of vitality or lethargy, and so-on.

There is good reason to think that it is primitive bodily awareness that provides the foundation for our sense of self, and that the capacity to have a single point of view begins with proto-rational sensorimotor subjectivity rather than rational self-conscious or self-reflective understanding. As Mark Johnson has pointed out, one central feature of our bodily life is our experience of containment and boundedness.³⁰ From the beginning of our lives, we move in and out of a variety of bounded spaces, including cribs, rooms, and vehicles. We manipulate objects and place them in containers, and also experience our own bodies as containers into which we put food. We also experience physical containment in our environment insofar as things envelop us and we experience ourselves as differentiated and separated from that which lies beyond us. Our essential embodiment

gives a very definite character to our perceptual experience and establishes a center and a periphery, so that one's world radiates out from one's body. As one interacts with the world from this perceptual and experiential egocentric standpoint, one develops an implicit awareness of the distinction between self and other.³¹ And as Maxine Sheets-Johnstone has noted, the first realizations we have of ourselves as infants are realizations about bodies—the body that we are and the bodies we are not.³²

Drawing upon J.J. Gibson's notion of the "ecological self," other theorists have maintained that through somatic proprioception, the perceiver gains not only roughgrained but also finegrained or even hyper-finegrained information about her position, movement, limb disposition, and other bodily properties. The pragmatic and egocentric spatial framework involved in sensorimotor subjectivity does not correspond precisely to objective measurements, but instead is a body-centered and perspectival spatial framework that involves experiential reference to one's body. Such ecological self-awareness "manifests itself as an integrated or global sense of where [she is] spatially in relation to the immediate environment" and what she is capable of doing.³³ Proprioceptive awareness keeps track of the body as the center and focal point of body-relative egocentric space, which is what makes orientation and action possible.³⁴ Objects are both perceptually situated, as well positioned as potentially manipulable, by virtue of their positioning and orientation in relation to our bodies. In this way, primitive bodily awareness shapes our earliest understanding of the world in perception and action, as well as our sense of self, so that sensorimotor subjectivity ultimately paves the way for more sophisticated forms of cognition, including self-consciousness_{lo} and reflection.

The crucial point here is that self-consciousness_{lo} or reflective consciousness requires sensorimotor subjectivity and nonself-conscious or pre-reflective consciousness_{lo}, but sensorimotor subjectivity and nonself-conscious or pre-reflective consciousness_{lo} does *not* require self-consciousness_{lo} or reflection. At least some non-human animals—e.g., Nagel’s bat—and all normal human infants have sensorimotor-subjective or pre-reflectively conscious states that are *not* also self-conscious or reflective. And even among normal adult human animals, sensorimotor subjectivity or pre-reflective consciousness_{lo} is frequently present even though self-consciousness_{lo} or reflection is not.

In the early stages of emotional experience, before one has had the opportunity to consider and reflect on what is happening, one nonetheless is occurrently conscious. Suppose that you wake up to the sound of glass shattering in your living room, and because your attention is directed elsewhere and you are focused on trying to figure out what caused the sound, you may very well not self-consciously or self-reflectively experience fear. While Jesse Prinz³⁵ describes this as a case of *unconscious* emotion, I believe that to be afraid, one must already be approaching the world from a single point of view that intrinsically involves subjective experience and which is grounded in one’s egocentrically-centered and spatially oriented essential embodiment. The sound’s presentation and affective influence involve pre-reflective bodily awareness and have a phenomenal character, so that even when one’s attention is directed elsewhere, one still feels the fear. It makes little sense to say that one’s fear is unconscious and unfelt, for there is *something it is like* to live through that scary moment. To see that one does consciously hear the sound and feel afraid before noticing it or turning one’s attention to

it, consider the fact that the content of the experience is access conscious insofar as one is poised to make use of one's pre-reflective emotional experience in guiding one's actions. On our view, sensorimotor subjectivity penetrates into every aspect of our mental lives, fully including the so-called "non-conscious" or "sub-personal" information processing that goes on in emotional experience. Although this information processing is indeed nonself-conscious or pre-reflective, it still involves subjective experience that is fully first-personal, conscious, and centered around a particular point of view.

Again, and despite the fact that I am a rational, self-conscious or reflectively conscious animal, when I am skillfully driving my car and drinking hot coffee without spilling it, but also thinking intensely about philosophy, the conscious acts or states that skillfully control my driving and my coffee-drinking are sensorimotor-subjective or pre-reflectively conscious but not in any way self-conscious or reflective. If they were, then I would most probably spill the hot coffee all over myself, and drive off the road into the ditch.

Since, presumably, *everyone* would agree that normal human infants and at least some non-human animals are conscious animals but not also self-conscious or self-reflective animals, and also that it is possible for rational, self-conscious, self-reflective animals like us skillfully to drive a car and at the same time drink hot coffee consciously and pre-reflectively but not self-consciously or reflectively, then at least implicitly everyone *already* concedes a distinction between sensorimotor subjectivity and meta-representational, self-conscious or reflectively conscious subjectivity.

And this is just one example of an action that we carry out "automatically" and "without even thinking about it." Other such perceptual and motor skills include our

ability to walk, navigate our way through doorways, reach and grasp for objects, hold our bodies upright in a sitting position, and even type on our computers. While I cook dinner, I am focused on what I am doing as I measure out ingredients, stir the sauce, cut vegetables, and get jars from the pantry. As I am working, I consciously feel my body parts and limbs in their relative positions. I feel my weight shift from one leg to the other as I stand in front of the stove, I feel my wrist move as I stir the sauce, and I feel my fingers move as I cut the vegetables. But I do not pay much attention to how my wrist or my hand feels unless they become awkward or uncomfortable. Skills automatized through experience and habitual practices, though they may seem to involve non-conscious or sub-personal mental processing, are still in fact occurrently conscious.

Another example of the dissociation between sensorimotor subjectivity and meta-representational, self-conscious or self-reflective subjectivity is the so-called “cocktail-party effect,” i.e., cases in which one screens out the sounds of conversations other than one’s own. If one’s name is mentioned, however, one’s attention immediately shift to that conversation, which indicates that one must already have had an auditory consciousness of what was being said. Pre-reflective consciousness often does not involve attention or focusing, and instead is peripheral. Although one experiences other people’s conversations implicitly and without self-reflective awareness of having heard them, one nonetheless lives through the state of being affected by their speech and is poised to make use of this auditory information. This strongly suggests that even in cases where what one hears from other people’s conversations is not accessed explicitly, it nonetheless is consciously experienced and accessible.

Hence it is not so very shocking after all for me to hold that all mental states, even “tacit” computational information processing states, are also occurrently conscious. All I am saying is that even “tacit” computational information processing involves *sensorimotor subjectivity*, and nonself-conscious or pre-reflective consciousness_{lo}, but not meta-representational, self-conscious or reflective subjectivity. Moreover, as I understand these phenomena, sensorimotor subjectivity is the unmediated and *essentially non-conceptual*³⁶ mode of bodily consciousness that forms the necessary basis for all other conscious states.

If we were sufficiently careful about the distinction between essentially non-conceptual and sensorimotor subjectivity or pre-reflective (i.e., type-2) consciousness on the one hand, and the conceptually-driven, self-conscious or self-reflective (type-1) consciousness on the other, then by means of The Deep Consciousness Thesis it is arguable that even the deeply puzzling and much-discussed phenomenon of *blindsight*³⁷—in which some brain-damaged subjects who introspectively report an inability to see are also able to point with some accuracy to objects in the self-professedly blind parts of their visual fields³⁸—can be plausibly explained in a way that avoids several of the more philosophically questionable implications of the standard explanations. For according to The Deep Consciousness Thesis, it can be held that not only the roughgrained sensorimotor ability manifest in actual blindsight, but also the finegrained or hyper-finegrained—respectively, in the thought-experimental cases of what Ned Block calls “superblindsight” and “superduperblindsight”³⁹—sensorimotor connection between what blindsighters perceive in space, and their ability to point to it, discriminate it, or track it, is guided and mediated by the information carried by or

contained in essentially non-conceptual and sensorimotor-subjective or pre-reflectively (i.e., type-2) conscious vision, even though they lack self-conscious or reflective (i.e., type-1) conscious vision for that cognitive and practical task.⁴⁰

Otherwise put, according to the Deep Consciousness Thesis-based proposal, in blindsight the frontline information-processing mechanisms of the eyes and related areas of the wider brain-body system (whose neural operations are, perhaps, localized in the parietal lobe) are relevantly and relatively undamaged and continue to transmit sensorimotor-subjective or pre-reflectively conscious visual information, even though the corresponding downstream mechanisms for processing self-conscious or reflective visual information (whose neural operations are, perhaps, localized in the temporal lobe) have broken down. Blindsighters, after all, *have their eyes open and are working under well-lit conditions*. Blindsighters would then be best and most coherently characterized as “sighted” in *one* sense of conscious vision, but also “blind” in *another* sense of conscious vision, instead of being paradoxically characterized as being both “blind” and “sighted” in the *same* sense of conscious vision.

If The Deep Consciousness-based proposal is correct, then blindsighters subjectively experience conceptually-driven, self-conscious or self-reflective blindness via the more sophisticated “what”-sensitive downstream processing mechanisms of the brain-body system, but also subjectively experience essentially non-conceptual, sensorimotor-subjective, pre-reflectively conscious sight via the simpler “where”-sensitive processing mechanisms of the eyes and related parts of the brain-body system. Blindsighted subjective experience, presumably, has its own unique sort of phenomenal character and thus its own unique “something it is like to be *for* the organism.” Otherwise

put, presumably, blindsight is a determinate kind of conscious perceptual experience, just as ordinary seeing is a determinate kind of conscious perceptual experience. The blindsighted person obviously is not unconscious, and therefore (it seems to me) obviously is consciously feeling and doing something in a determinately specific way when she “blindsees” an object.

Furthermore, the notion of a divided consciousness is already theoretically familiar from well-known experiments involving divided attention tasks and the dissociated cognitive abilities of neo-commissurotomy patients (i.e., “split brain” patients whose *corpus callosum*, the primary neural connection between the two brain hemispheres, has been surgically severed), and functionally similar agnosias. Most importantly for the present purposes, there are the well-known Milner and Goodale data in favor of the hypothesis that there are two relatively distinct visual pathways of information processing, the *ventral stream* and the *dorsal stream*. The ventral stream is localized in the temporal lobe and supports so-called “conscious”—or I would want to say, in correction of that crucially ambiguous term, *conceptually-driven and self-consciously or reflectively (i.e., type-1) conscious*—visual perception. And the dorsal stream is localized in the parietal lobe and supports so-called “non-conscious”—or as I would want to say, in correction of that crucial misnomer, *essentially non-conceptual and nonself-consciously or pre-reflectively (i.e., type-2) conscious*—visual perception.⁴¹ So what this proposal entails is that in blindsight the ventral stream, as the support for *one* kind of conscious vision, is significantly compromised, while the dorsal stream, as the support for the *other* kind of conscious vision, remains uncompromised.

This way of thinking about blindsight, in turn, would neatly avoid the further and even deeper paradox that in blindsight a brute, non-conscious, non-unified, purposeless mental processing somehow exerts roughgrained, finegrained, or hyper-finegrained control over our essentially embodied cognitive and practical intentional agency. If this were true, then blindsighters would be nothing but natural automata with respect to their blindseeing activities. But it seems more than just implausible to hold that blindsighted people are nothing but naturally mechanized puppets or robots in the blind areas of their self-conscious or self-reflective visual fields, but *otherwise* really intentional agents. On the contrary, it seems abundantly clear to me, that blindsighted people are rational human *real intentional agents* who are genuinely visually conscious in those areas, and also genuinely choose and act under the relevant experimental conditions, such that they are ultimate sources of their own intentional body movements, which are thereby “up to them,” and such that they are also causally and morally *responsible* for their movements. After all, the scientists in blindsight experiments are certainly not *overwhelming manipulators* like, e.g., the evil cognitive scientist in *The Manchurian Candidate*.⁴² So our intuitive, smooth attribution of *responsibility for their intentional body movements* to the blindsighted subjects is good prima facie evidence *against* their being nothing but naturally mechanized puppets or robots in the blind areas of their conceptually-driven and self-consciously or reflectively (i.e., type-1) conscious visual fields, hence also good prima facie evidence *against* their being non-conscious cognizers in that domain, and also good prima facie evidence *in favor of* their being essentially non-conceptual and nonself-consciously or pre-reflectively (i.e., type-2) conscious cognizers in that very domain.

It is true that both blindsighted conscious experience and also blindsighted choosing and doing occur in a way that is in some determinate respects sharply and intrinsically phenomenologically, semantically, and biologically/neurobiologically different from the visual consciousness and intentional visual activity of normal conceptually-driven and self-consciously or self-reflectively (i.e., type-1) consciously sighted people. So blindsighters have an essentially non-conceptual and sensorimotor-subjective or pre-reflective visual (i.e., type-2) consciousness that is just like those of ordinary self-consciously or self-reflectively sighted people, but at the same time those blindsighters simply differ determinately, specifically, and sharply from ordinary sighted people at the cognitive and practical level that is inherently guided and mediated by *conceptual* content. At the same time, however, no one doubts that, other things being equal, blindsighters are operating normally as rational human animals and real intentional agents during the course of the blindsight experiments. So, in effect, all that The Deep Consciousness Thesis is saying is that blindsighters are rational human animals and intentional agents *all the way down*. Their higher-level conceptually-driven and self-consciously or self-reflectively (i.e., type-1) conscious cognitive activity and their lower-level blindsighted essentially non-conceptual and nonself-consciously or pre-reflectively (i.e., type-2) conscious cognitive activity are *not* two essentially separate processes—one causally closed inherently ghostly and immaterial process (pure epiphenomenal rationality), and another causally closed inherently mechanical and material process (pure mechanical animality), as the classical Cartesian Dualist Interactionist conception of the mind implies. Rational human animals are rational *and* humanly animal *and* type-2

conscious all the way through and all the way down, *including* the essentially non-conceptual blindseeing activities of blindsighters.

This explanation of blindsight, correspondingly, suggests a new way of explaining the equally puzzling phenomenon of “filling-in.”⁴³ Filling-in is the fact that our visual field presents itself as rich and continuous *even though* we have blind spots on our retinas. Various possible solutions to the puzzle have been offered. My solution is that filling-in is, in effect or even essentially, the *cognitive converse* of blindsight. In blindsight, the cognitive subject has sensorimotor-subjective or pre-reflectively conscious vision without self-conscious or self-reflective vision—that is, she has sensorimotor-subjective or pre-reflectively conscious vision via the simpler processing mechanisms of the eyes, together with self-conscious or self-reflective blindness via the more sophisticated processing mechanisms of the downstream brain-body system. Conversely, in filling-in, cognitive subjects have self-conscious or self-reflective vision *without* sensorimotor-subjective or pre-reflectively conscious vision, that is, they have self-conscious or self-reflective vision via the more sophisticated processing mechanisms of the downstream brain-body system, together with sensorimotor-subjective or pre-reflectively conscious *blindness* via the simpler processing mechanisms of the eyes. If this explanation is correct, then it will smoothly bind together blindsight and filling-in within a single theoretical framework.

Another body of evidence that supports The Deep Consciousness Thesis is the research on implicit memory among amnesiacs and Alzheimer’s patients, as well as normal subjects. Unlike explicit memory, which requires self-reflective recollection of previous experience, implicit memory is characterized by a lack of reflective self-

awareness in the act of recollection. Subjects demonstrate implicit memory when their previous experiences facilitate performance on a task even though they do not (or are unable to) intentionally recollect those experiences. This has led some theorists to characterize implicit memory as non-conscious or sub-personal. However, on my view, memory of recent experiences cannot be expressed in the complete absence of conscious awareness, though of course it can be expressed in the absence of self-conscious, self-reflective, higher-order recollection. While explicit memory involves the ability *self-consciously* or *reflectively* to recall or recognize past experiences or recently processed information, implicit memories cannot be “looked up” at will or intentionally remembered to be used for action and practical reasoning. Instead, implicit memory involves the activation of sensory and motor capacities needed to perform specific tasks. Some theorists have argued that infants and very young children are capable of implicit memory alone, and studies have shown that implicit memory may be largely unimpaired even as a person’s powers of explicit memory decline with age or deteriorate more radically as a result of amnesia or Alzheimer’s disease.

One mode of implicit memory is skill learning, which consists of learned, seemingly automatic capacities or skills. In *The Influence of Habit on the Faculty of Thinking* (1929) Maine de Biran described how, after sufficient repetition, a habit can be executed seemingly automatically, without awareness of the act of itself or of the previous experiences that allowed one to learn the habit. He had in mind two sorts of habitual expressions: repeated movements (executed via mechanical memory), and feelings (sensitive memory). These repeated movements and feelings are eventually executed with such automaticity that “we no longer see the voluntary action which directs

them and we are absolutely unaware of the source that they have.”⁴⁴ I believe that this is because the frontline information-processing systems that take in stimuli or perceptual input from the external environment are associated or linked with the relevant motor skills in the individual’s sensorimotor-subjective experience, even though the corresponding downstream mechanism for processing self-conscious or reflective awareness of this linkage is not involved or has broken down. This is not self-conscious or reflectively conscious memory that one could report, but instead sensorimotor-subjective, pre-reflective memory. In view of the fact what the phenomenologists like to call our “lived bodies”⁴⁵ frequently implicitly remember what our higher-order thinking processes do not, it is not surprising that so many colleges and universities now explicitly emphasize the classical Dewey-inspired thesis of “learning by doing” on the new-and-improved basis of contemporary cognitive psychology and Philosophy of Mind.

Evidence of perceptual-motor skill learning in amnesiac patients supports these claims. For example, as early as 1845, British physician Robert Dunn observed how, during an amnesiac state, one of his patients learned how to make dresses. Even though she had no explicit memory of having made any dresses and could not recollect what she had done from day to day, each day she would set to work utilizing the skills she had learned. This strongly suggests that she implicitly remembered how to make dresses, even though she was unable to recollect having done so the day before in any sort of explicit or self-reflective way.⁴⁶ Another example, this one from the 1960s, is the well-known case of the amnesiac named H.M., a bilaterally lesioned frontal patient who could acquire motor skills such as mirror tracking, even though he did not explicitly remember and was unable to recollect or report on having previously performed such tasks. Still,

his performance improved at a normal rate, which indicates that he had an implicit memory of these past experiences.⁴⁷

More recently, a study by Gillsky, Schacter, and Tulving showed that an amnesiac patient could learn how to program a microcomputer despite the patient's persistent failure to remember explicitly having ever worked on a microcomputer.⁴⁸ As in the previously described cases, here memory seems to be stored and expressed through sensorimotor systems. His implicit memory of a past pattern of agency, including knowledge of how to perform this skill remained part of his sensorimotor-subjective awareness even though he could not self-consciously or self-reflectively recollect having learned it. However, as I have argued, this does *not* itself entail that implicit memory is "non-conscious" or "sub-personal." According to my view, there is good reason to understand the tacit computational informational processing involved in executing a previously learned skill as pre-reflectively conscious memory, and to suppose that when he encountered the programming task again, after already having learned how to perform it, he approached it with a background bodily orientation that made the task at hand feel *familiar*. This amnesia patient is not unconscious, and therefore is performing the task at hand in a determinately specific way. Moreover, as he programs the computer, he consciously experiences making use of his past learning, so that his present subjective experience involves a definite *aspect of familiarity* and thus is very different from what it might have been if he had never encountered this task before. This in turn strongly indicates that amnesiac patients who exhibit implicit memory are indeed genuinely consciously remembering something, even if they cannot self-consciously or self-reflectively recollect and make reports about their past experiences.

Another mode of implicit memory is repetition priming, which occurs when past experience with a stimulus facilitates a subject's ability to process that stimuli on subsequent occasions. In one type of implicit memory test, subjects are presented with a set of words and then later asked to complete partial words with whatever comes to mind. Studies have shown that subjects tend to complete words to match words to which they were previously exposed. The performance of Alzheimer's patients on such word-identification priming tests consistently remains intact. What is particularly interesting in this connection is that on category-exemplar priming tests, which rely much more heavily on conceptual processing and ask subjects to name items belonging to a particular category, performance *was* impaired. This indicates that Alzheimer's disease affects some forms of repetition priming more than others, and that there is a "functional dissociation between conceptual versus perceptual processes in priming."⁴⁹ Only to the extent that repetition priming draws on perceptual, bodily processes does it tend to be preserved in Alzheimer's patients. This further supports my characterization of implicit memory as non-self-consciously or pre-reflectively conscious memory.

Implicit memory also seems to play a role in affective and social phenomena. One example of what Maine de Biran termed "sensitive memory" is the famous case of an amnesiac woman who refused to shake hands with her physician after he pricked her with a pin, even though she did not explicitly remember that he had done so. It seems that the woman had formed an association between the physician's presence and her own bodily discomfort, though this association was not formed through reflection or conceptualization. Instead, she implicitly remembered the bodily sensations aroused by the physician and this affected her future behavior when she was again presented with

that individual. More generally, it seems clear that contact with a stimulus one has perceived previously triggers the same bodily changes and emotional feelings it has in the past, even if one does not explicitly recollect one's past experience.

In Charlie Kaufman and Michel Gondry's 2004 film, *Eternal Sunshine of the Spotless Mind*, the main characters are strangely drawn to each other and seem to have a special connection, despite their very different personalities and the fact that they appear not to have met before. As the film unfolds, the audience learns that Clementine and Joel have undergone a procedure to erase their memories of each other when their relationship turns sour. Although they have no clear or explicit memory of having known each other, let alone having had a relationship and electing to have their memories dry-cleaned, ultimately they are drawn together despite their forgotten history. Do Joel and Clementine remember each other? Here the answer seems obviously to be both "yes" and "no," and so the distinction between implicit and explicit memory, when combined with The Deep Consciousness Thesis, appears to be an intuitively natural way of conceptually explicating this fictional scenario. Although Joel and Clementine are amnesiacs with respect to their relationship and have no explicit memories of their time together, their experience of meeting on the beach at the beginning of the film clearly involves an aspect of familiarity. Perhaps this includes the familiar sound of her voice, or the familiar way he smells, or the rhythm of their gait as they stroll down the beach. But it seems clear that in some sense Joel and Clementine do remember each other, and this intuitively suggests that we do not always remember things in an abstract, intellectual, or detached manner. In many and perhaps even all cases, memory is intimately bound up with bodily and perceptual processes, and thereby includes what we know essentially "by heart."

Indeed, several studies appear to show that stimuli that are not explicitly perceived or remembered nonetheless influence subsequent affective responses. In one study, subjects were exposed to geometric shapes at durations too brief to permit “conscious” (i.e., self-conscious or reflectively conscious) perception. While subjects showed no explicit memory of the shapes on subsequent tests, they did demonstrate implicit memory “by showing a reliable preference for the previously exposed shapes on a test” which asked them which of the shapes they liked better.⁵⁰ Similarly, subjects who had been given subliminal exposures to hostile words later rated a target person more negatively than those who had not been exposed to these words. Here background bodily orientation, which has been impacted by the exposure to hostile words, affects a subject’s subsequent affect and judgments, even though she does not self-reflectively remember having seen the hostile words. Likewise, patients with facial recognition deficits (prosopagnosia) exhibit stronger galvanic skin responses to familiar faces, despite the fact the fact that they do not explicitly recognize any faces as familiar.

In an attempt to understand the distinction between explicit and implicit memory, some theorists have claimed that while explicit memory relies heavily on conceptually-driven processes, implicit memory draws mostly from data-driven processes (ones initiated or guided by the information or data presented in test materials.)⁵¹ What is crucial to add to this account, of course, is that implicit memory is activated by motor activity or affective experience. It draws heavily from bodily-driven processes, and this is why it shows up repetition priming effects and in amnesiacs ability to remember *how* to perform various tasks. As in the case of blindsight, subjects who exhibit implicit memory demonstrate through their emotional response or performance of a task that they possess

a certain kind of knowledge, and yet are not reflectively aware that they possess this knowledge, nor can they gain explicit access to it.

But, again, why do I not characterize implicit memory as “non-conscious” or “sub-personal,” as so many other theorists have done? To repeat, this is because I conceive of implicit memory as a sensorimotor-subjective mode of memory that is grounded in primitive bodily awareness and pre-reflective consciousness_{io}. For one thing, it seems clear that implicit memory guides subjects’ intentional actions. As in the case of blindsight, it is highly implausible to hold that subjects relying on implicit memory are mere mechanized puppets or robots with respect to these memories or movements. In addition, implicit memory has its own unique sort of phenomenal character, so that there is *something it is like* for the amnesiac to remember how to program a microcomputer; *something it is like* to remember having being pricked with a pin by the physician; and *something it is like* to see a familiar face or familiar shape even though one cannot explicitly, self-consciously, or self-reflectively recollect having seen it before. In other words, in all such cases there is an aspect of familiarity. Matthew Ratcliffe describes “affective familiarity or the feeling of familiarity” as a “structural constituent of intentional states, which plays a distinctive role in shaping world experience.”⁵² He characterizes it as a background bodily orientation rather than an internal report on bodily states. Feelings of bodily states contribute to the way the world appears and how we interact with that world, so that our bodily dynamics and comportment are bound up with our understanding of and engagement with our surroundings. In cases of implicit memory, certain situations are experienced as calling for a particular sort of response, so that one’s performance of a task involves a feeling of familiarity. The lived body takes

over and is freely in control: our proto-rationally intelligent and pre-reflectively consciously lived bodies know what to do even though our rationally self-conscious or reflective intellect does not.

With this background bodily orientation in place, we remember much of our past experiences, but not necessarily in the somewhat detached or distanced representational manner involved in explicit memory. Nonetheless, implicit memory is a distinctive kind of conscious cognitive experience having a certain specific phenomenal character, just as explicit memory is a determinate kind of conscious cognitive experience having a certain specific phenomenal character. Introspectively, we *can* tell the difference, even if this is not infallible and we are occasionally mistaken. But the salient difference in the veridical cases is this: While only explicit memory involves higher-order consciousness or self-consciousness, reflection, categorization, and conceptualization, both implicit and explicit memory alike rely heavily on nonself-conscious perceptual and sensorimotor capacities, and on essentially non-conceptual content. Thus there is good reason to think that pre-reflective consciousness_{to} is crucial for memory, and that the amnesiac is consciously feeling and doing something in a distinctively specific way when she implicitly remembers something.

V. Conclusion

Clark has simply assumed, without argument, that all human conscious experience is type-1/self-conscious or reflective conscious experience, and has not in any way ruled out the nontrivial possibility that there are could also be type-2/nonself-conscious or pre-reflective conscious experience occurring in the extra-neural living human body, constitutively involving of course the brain and central nervous system, but

also constitutively extending beyond them into all the other vital systems. The arguments I developed in sections **III and IV**, if sound, collectively suffice to show that there is *both* type-1/self-conscious or reflective conscious experience occurring in the brain and central nervous system at fast timescales *and also* type-2/nonself-conscious or pre-reflective conscious experience occurring at slower timescales throughout the rest of living human body, including of course all the vital systems that biologically and dynamically enclose and enfold the brain and CNS. The clear difference in timescales is *not* indicative of a “constitutive/causal cut,” but instead *is* clearly indicative of an essential difference in *types of human consciousness*. And that vindicates The Essential Embodiment Thesis and the The Deep Consciousness Thesis alike.

I conclude, then, that Clark’s argument for The Cautious Consciousness-Is-All-Neural Thesis fails, that The Cautious Consciousness-Is-All-Neural Thesis is in fact false, that The Extended Conscious Mind Thesis is also false, and that The Essential Embodiment Thesis and The Deep Consciousness Thesis are both true. If so, then the “joy” of human conscious experience, whether type-1/self-conscious or reflective conscious experience, or type-2/nonself-conscious or pre-reflective conscious experience, occurs everywhere in our living bodies, constitutively including of course the brain and the central nervous system, *and also* constitutively including all the other vital systems of the living body, right out to the skin, *but no further out than that*.

One last, and deeper, thing. Wittgenstein points out in the epigraph for this essay that the philosophical move of spreading consciousness all over the human body can radically transform our fundamental conception of ourselves. We can then realize that *neither* are we ontologically inflated Cartesian mental substances or “thinking things”

(a.k.a. Cartesian Dualism), *nor* are we ontologically reduced brains-plus-central-nervous-systems (a.k.a. “Cartesian Materialism”⁵³), because there simply are *no such noumenal ego-things*, and therefore to think so is either just a Cartesian Dualist myth or its contrary flipside, a Cartesian Materialist myth. Nor, however, are we *extended* conscious minds. Quite literally, The Extended Conscious Mind Thesis *goes too far* in its otherwise laudable attempt to avoid Cartesian Dualism and Cartesian Materialism alike.

Contrary to Cartesian Dualism, Cartesian Materialism, and The Extended Conscious Mind Thesis, then, we are nothing more and nothing less than *rational human minded animals*, living, moving, and having our being all the way out to the limits of our “human, all too human” skin, *but no further out than that*. We cognize ourselves, the world, and one another. And we intentionally act in the world, both in relation to ourselves, and also in relation to one another. But *only* our own living bodies are consciously *minded* by us. As Wittgenstein’s deep remark intimates, the hardest thing of all for the Philosophy of Mind is to see ourselves as we really are.

NOTES

- ¹ L. Wittgenstein, "Notes for Lectures on Private Experience and Sense Data," in J. Klagge and A. Nordmann (eds.), *Philosophical Occasions* (Indianapolis, IN: Hackett, 1993), pp. 200-288, at p. 225.
- ² A. Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," *Mind* 118 (2009): 963-993, at p. 987.
- ³ Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," pp. 964-969.
- ⁴ Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," pp. 969-980.
- ⁵ Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," pp. 980-983.
- ⁶ Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," pp. 986-987, underlining added.
- ⁷ See R. Rupert, "Challenges to the Hypothesis of Extended Cognition," *Journal of Philosophy* 101 (2004): 389-428; and R. Rupert, *Cognitive Systems and the Extended Mind* (Oxford: Oxford Univ. Press, 2009).
- ⁸ Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," p. 986.
- ⁹ See, e.g., J. Kihlstrom, "The Cognitive Unconscious," *Science* 237 (1987): 1445-1452.
- ¹⁰ See, e.g., the article at URL = <http://en.wikipedia.org/wiki/Out-of-body_experience>.
- ¹¹ Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," p. 982.
- ¹² Clark, "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head," pp. 982-983.
- ¹³ J. Hawkins and R. Allen (eds.), *Oxford Encyclopedic English Dictionary* (Oxford: Clarendon/Oxford Univ. Press, 1991), p. 52.
- ¹⁴ See C. Allen and M. Bekoff, *Species of Mind* (Cambridge: MIT Press, 1997); D.R. Griffin, *Animal Minds* (Chicago: Univ. of Chicago Press, 2001); D.R. Griffin, *Animal Thinking* (Cambridge: Harvard Univ. Press, 1984); and D.R. Griffin, *The Question of Animal Awareness* (New York: Rockefeller Univ. Press, 1976).
- ¹⁵ See R. Hanna and M. Maiese, *Embodied Minds in Action* (Oxford: Oxford Univ. Press, 2009), esp. chs. 1-2.
- ¹⁶ See Hanna and Maiese, *Embodied Minds in Action*, esp. chs. 3-5.
- ¹⁷ See H. Putnam, *Reason, Truth, and History* (Cambridge: Cambridge Univ. Press, 1981), ch. 1.
- ¹⁸ See Hanna and Maiese, *Embodied Minds in Action*, section 8.1.

¹⁹ R. Descartes, *Meditations on First Philosophy*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch, in R. Descartes, *The Philosophical Writings of Descartes*, 3 vols., (Cambridge: Cambridge Univ. Press, 1985), vol. II, pp. 3-62, at p. 59, AT: 86.

²⁰ The relevant set of neurobiological properties alone is not a *sufficient* condition of the existence of a consciousness_{io}, however. Instead, the existence of a consciousness_{io} is *jointly hylomorphically constituted* by relevant mental and neurobiological properties. See Hanna and Maiese, *Embodied Minds in Action*, section 8.1.

²¹ Just as in the case of the existence of consciousness_{io}, so too the relevant set of neurobiological properties alone is not a sufficient condition of the *specific character* of a consciousness_{io}. Both the existence and specific character of a consciousness like ours are jointly hylomorphically constituted by relevant mental and neurobiological properties. Again, see Hanna and Maiese, *Embodied Minds in Action*, section 8.1.

²² For example, the 29 January 2007 issue of *Time* magazine was entirely devoted to the topic, *The Brain: A User's Guide*, and included supportive articles by or interviews with many leading contemporary philosophers of mind and cognitive neuroscientists.

²³ See, e.g., M. Gershon, *The Second Brain* (New York: HarperCollins, 1998).

²⁴ See, e.g., A. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain* (New York: Avon Books, 1994); A. Damasio, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness* (San Diego, CA: Harcourt, 1999); A. Damasio, *Looking for Spinoza: Joy, Sorrow, and the Feeling Brain* (San Diego, CA: Harcourt, 2003); C. Pert, *Molecules of Emotion* (New York: Scribner, 1997); and J. Prinz, *Gut Reactions: A Perceptual Theory of Emotion* (New York: Oxford, 2004).

²⁵ T. Rockwell, *Neither Brain Nor Ghost: A Nondualist Alternative to the Mind-Brain Identity Theory* (Cambridge: MIT Press, 2005), p. 47, italics in the original.

²⁶ See, e.g., R. Jackendoff, *Consciousness and the Computational Mind* (Cambridge: MIT Press, 1987).

²⁷ See J. Bermúdez, "Nonconceptual Content: From Perceptual Experience to Subpersonal Computational States," in Y. Gunther (ed.), *Essays on Nonconceptual Content* (Cambridge: MIT Press, 2003), pp. 184-216. Bermúdez holds that subpersonal states have non-conceptual content, but would not agree that they are also conscious.

²⁸ See E. Thompson, "Sensorimotor Subjectivity and the Enactive Approach to Experience," *Phenomenology and the Cognitive Sciences* 4 (2005): 407-427.

²⁹ T. Nagel, "What is like to be a bat?," in *Mortal Questions* (Cambridge: Cambridge Univ. Press, 1979), pp. 165-180, at 166-167.

³⁰ M. Johnson, *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason* (Chicago: University Of Chicago Press, 1990), p. 21.

³¹ Johnson, *The Body in the Mind*, p. 124.

³² M. Sheets-Johnstone, "Phenomenology and Agency: Methodological and Theoretical Issues in Strawson's 'The Self,'" *Journal of Consciousness Studies* 6 (1999): pp. 48-69, at p. 63.

³³ S. Gallagher and A. Marcel, "The Self in Contextualized Action," *Journal of Consciousness Studies* 6 (1999): pp. 4-30, at p. 21.

³⁴ J. Bermúdez, *The Paradox of Self-Consciousness*, (Cambridge, MA: MIT Press, 2000), p. 145.

³⁵ See J. Prinz, “Are Emotions Feelings?,” *Journal of Consciousness Studies* 12 (2005): 9-25.

³⁶ See, e.g., e.g., R. Hanna, “Kant and Nonconceptual Content,” *European Journal of Philosophy* 13 (2005): 247-290; R. Hanna, “Kantian Non-Conceptualism,” *Philosophical Studies* 137 (2008): 41-64; R. Hanna and M. Chadha, “Non-Conceptualism and the Problem of Perceptual Self-Knowledge,” *European Journal of Philosophy* 17 (2010); and R. Hanna, “Beyond the Myth of the Myth: A Kantian Theory of Non-Conceptual Content,” *International Journal of Philosophical Studies*, forthcoming.

³⁷ See, e.g., L. Weiskrantz, *Blindsight* (Oxford: Clarendon/Oxford Univ. Press, 1986).

³⁸ Strikingly there is also a “blindimagination” analogue of blindsight, in which a brain-damaged subject reports a complete loss of conscious mental imagery, yet is able to score at least as well or better on “mental manipulation” tests (i.e., involving the mental comparison of two imagined figures, etc.) than subjects whose capacity for conscious imagination is normal. One other unusual feature of the blindimagination data is that although normal subjects take longer to compare imagined figures in direct proportion to the degree of difference between the angle of perspective on the two objects (the bigger the difference, the longer it takes the subjects to recognize the objects as the same or different), the blindimager always takes the same amount of time to produce the same answer. See URL = <http://discovermagazine.com/2010/mar/23-the-brain-look-deep-into-mind.s-eye/>. Many thanks to Devon Belcher for bringing these studies to my attention.

³⁹ See N. Block, “Concepts of Consciousness,” in Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*, pp. 206-218, at 211.

⁴⁰ Analogously, in the “blindimagination” case (see note 38 above), I would want to say that the subject still possesses conscious mental imagery and is manipulating it nonself-consciously or pre-reflectively, but has lost the ability to be self-consciously or reflectively aware of it. The fact that normal subjects take longer to compare imagined figures in direct proportion to the degree of difference between the angle of perspective on the two objects, while the blindimager always takes the same amount of time to produce the same answer, could then be explained by saying that the extra cognitive layer of self-conscious or reflective reporting on mental imagery actually slows subjects down and increases their response times.

⁴¹ D. Milner and M. Goodale, *The Visual Brain in Action* (Oxford: Oxford Univ. Press, 1995).

⁴² (Directed by J. Frankenheimer, 1962).

⁴³ See, e.g., L. Pessoa, E. Thompson, and A. Noë, “Finding out about filling in: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception,” *Behavioral and Brain Sciences* 21 (1998): 723-748.

⁴⁴ M. de Biran, *The Influence of Habit on the Faculty of Thinking*, (Baltimore: Williams and Wilkins, 1929), p. 73.

⁴⁵ See, e.g., E. Husserl, *Crisis of European Sciences and Transcendental Phenomenology*, trans. D. Carr (Evanston: Northwestern Univ. Press, 1970), §62, p. 217:

Everyone experiences the embodiment of souls in original fashion only in his own case. What properly and essentially makes up the character of a living body I experience only in my own living body, namely, in my constant and immediate holding-sway [over my surroundings] through this physical body alone.... Obviously it is only in this way that I have perceptions and, beyond this, other experiences of objects in the world.... It is only *my* being-as-ego, as holding-sway, that I actually experience as itself, in its own essence; and each person experiences only his own.

⁴⁶ D. Schacter, "Implicit Memory: History and Current Status," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13:3 (1987), pp. 501-518, at 503.

⁴⁷ Other studies have shown that amnesiacs exhibit near-normal learning of skills such as reading of mirror-inverted script and puzzle solving. See Schacter, "Implicit Memory: History and Current Status," p. 509.

⁴⁸ E.L. Gilsky, D.L. Schacter, and E. Tulving, "Computer Learning by Memory-Impaired Patients: Acquisition and Retention of Complex Knowledge," *Neuropsychologica*, 24 (1986): 313-328.

⁴⁹ D.A. Fleischman, et al., "Implicit Memory and Alzheimer's Disease Neuropathology," *Brain* 128 (2005): 2006-2015, at p. 2011.

⁵⁰ Schacter, "Implicit Memory: History and Current Status," p. 506.

⁵¹ Schacter, "Implicit Memory: History and Current Status," p. 511.

⁵² M. Ratcliffe, "Interpreting Delusions," *Phenomenology and the Cognitive Sciences*, 3 (2004): 25-48, at p. 26.

⁵³ See Rockwell, *Neither Brain Nor Ghost: A Nondualist Alternative to the Mind-Brain Identity Theory*, esp. ch. 1.