

International Phenomenological Society

The Message is: There is no Medium

Consciousness Explained by Daniel C. Dennett

Review by: Daniel C. Dennett

Philosophy and Phenomenological Research, Vol. 53, No. 4 (Dec., 1993), pp. 919-931

Published by: [International Phenomenological Society](#)

Stable URL: <http://www.jstor.org/stable/2108264>

Accessed: 31/08/2014 06:51

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



International Phenomenological Society is collaborating with JSTOR to digitize, preserve and extend access to *Philosophy and Phenomenological Research*.

<http://www.jstor.org>

The Message is: There is no *Medium*

DANIEL C. DENNETT
Tufts University

Sydney Shoemaker notes that my “avoidance of the standard philosophical terminology for discussing such matters” often creates problems for me; philosophers have a hard time figuring out what I am saying and what I am denying. My refusal to play ball with my colleagues is deliberate, of course, since I view the standard philosophical terminology as worse than useless—a major obstacle to progress since it consists of so many errors trapped in the seductively lucid amber of tradition: “obvious truths” that are simply false, broken-backed distinctions, and other cognitive illusions. I want to shift the perspective of philosophy of mind, and for that task using the standard terminology would be counterproductive. Fortunately, the inevitable communication-difficulties my policy provokes are forced into the open by occasions such as this constructive confrontation, permitting me to clarify my shocking message. I am grateful to Shoemaker, and to Michael Tye, Frank Jackson and David Rosenthal, for their vigorous and sympathetic reactions to my book.

Traditional philosophy of mind, in my opinion, has a rather bad case of tunnel-vision, excusing itself from many interesting and important questions. This is excellently revealed by the puzzlement Shoemaker and Tye express about my targets. Just who am I arguing against, in my assault on the Cartesian Theater? Could it be sense-datum theorists, Shoemaker wonders: “But sense-datum theory has few defenders nowadays, and seems unlikely to be Dennett’s main target.” Tye comments “Why he should be so concerned to refute it is not immediately clear. For Dennett himself associates no specific philosophers or psychologists with Cartesian materialism, and I know of no one who endorses it.” Well, if Shoemaker and Tye know no philosophers who are guilty of the errors I am scouting, that is probably because almost no philosopher has ventured to work out a theory of consciousness in sufficient detail to confront the problems that lead more intrepid theorists, with great regularity, into the errors of Cartesian materialism. Among psychologists and neuroscientists who do aspire to model consciousness as a neural phenomenon, the temptations of “sense-datum theory” (though not, of course, under that name) are almost irresistible. When

I called Cartesian materialism the view nobody holds, I meant it ironically; whereas almost nobody would avow it, the hopes and hunches and arguments of *most* theorists—philosophers included—betray a sneaking and unrecognized commitment to the view. I have been gratified by the widespread acknowledgment of this by scientists in response to my “outing” of the theory. Antonio Damasio, one of the world’s leading theoretical neuroscientists, says, in his commentary on Dennett and Kinsbourne, “Time and the Observer: the Where and When of Consciousness in the Brain,” in *Behavioral and Brain Sciences*: “This notion, which D & K modestly refer to as a ‘prevailing view,’ is far worse than that: It informs virtually all research on mind and brain, explicitly and implicitly.”¹

But when David Rosenthal says that the Cartesian Theater model “tacitly underlies much that’s mistaken in current thinking about mind and consciousness,” I suspect he has philosophers in mind, primarily, and he’s right about them, too. The few philosophers who *have* tried to say something specific about how consciousness might fit in the brain end up endorsing Cartesian materialism implicitly (see, e.g., the commentaries by Ned Block and Robert van Gulick in the same issue of *Behavioral and Brain Sciences*), or even explicitly. Michael Lockwood, in his essay on my book,² now declares himself to be a Cartesian materialist.

Indeed, if Tye and Shoemaker want to see a card-carrying Cartesian materialist, each may look in the mirror, for their commitment to qualia or “phenomenal seemings” itself involves a covert assumption that there is a privileged medium in the brain, the medium in which these properties get instantiated—as I will now try to show briefly. When Tye says that advocates of the reality of phenomenal seemings “are not committed to holding that there is a single place in the brain in which the seemings and feelings occur” this is really just because these advocates have not developed their allegiance to this reality beyond the vaguest of handwavings. Presumably “phenomenal seemings” are seemings-to-*me* (as opposed to the non-phenomenal seemings that are just seemings-to-my-retina or seemings-to-some-unconscious-control-process), but making this distinction *requires* that one identify a privileged neural medium—the *Medium*, you might call it—transduction into which marks the onset of true phenomenal seeming.³

¹ Antonio R. Damasio, “The Selfless Consciousness,” *Behavioral and Brain Sciences*, 15, 1992, pp. 208–9.

² Michael Lockwood, 1993. “Dennett’s Mind,” *Inquiry*, 36, pp. 59–72.

³ This medium might be anatomically spread around, but that creates further problems for the model—see Dennett and Kinsbourne, “Escape from the Cartesian Theater,” our response to the commentaries in *Behavioral and Brain Sciences*, 1992, pp. 234–47, esp. sections 1, 2, and 4.

I owe this new way of putting it to Bruce Mangan, who in his forthright attack on my book,⁴ found a formulation of Cartesian materialism of enviable clarity and vividness: consciousness, he proposed, is itself a particular physical medium in the brain. You don't have conscious visual experience, for example, as soon as light strikes your retina, for that is not yet the medium that *is* consciousness. He declines to tell us just which neural tissues compose the medium of consciousness, but at the moment when information eventually gets transduced into that medium, it enters conscious experience. In the aftermath of spectrum inversion fantasies, for instance, when the host of merely reactive dispositions are becoming re-inverted in order to normalize the merely functional or behavioral properties, it is the state of affairs in the *Medium* that determines whether or not one's qualia are still inverted. If there is no such medium, there is no fact of the matter about whether one's qualia (as distinguished to one's merely reactive dispositions) are inverted—more on this later.⁵

I turn now to respond in more detail to Tye, Jackson, Shoemaker, and Rosenthal, in that order.

I claim that verificationism makes sense precisely in the case of consciousness, citing Kant as a distinguished predecessor (p. 132n), but Tye prefers Mr. McCawber's conservatism. Tye sees, correctly, that I base my claim on the belief that "the central concept of phenomenal experience or seeming can only be understood in terms of concepts pertaining to judgment or belief." A major problem he sees with this is that it renders my "main argument" redundant. This would not worry me, since my main argument might still be the best way of drawing out the implications of my main claim for the imagination-impaired. It all hinges, as Tye notes, on whether there might be "real seeming" independent of all judging or believing. Why? Because the only way to overrule my verificationism would be to find a ground for insisting that there was some real seeming going on, about which a subject's *judgments* were inconclusive.⁶

The term "judging," which I use loosely, also has more strict senses, and Tye and Shoemaker both show what problems arise for my view if one takes me to be using the term in its strict—and perhaps standard—philosophical

⁴ "Dennett, Consciousness and the Sorrows of Functionalism," *Consciousness and Cognition*, 2, 1993, pp. 1–17. See also my reply, "Caveat Emptor," in the same issue.

⁵ When Shoemaker speaks of the duration of "the experience" (as opposed to the duration of the entire neural process from sense organ stimulation to eventual quiescence) he must be supposing that there is something like a transduction event that marks the onset of conscious experiencing as a subprocess within the larger process—and this is the essence of Cartesian materialism.

⁶ By the way, Tye slightly misdescribes his example of wheels seeming to revolve backwards. There has to be a source of phased snapshotting—film or television frames, or strobe light or, under special conditions, just the 60-cycle (or in Europe, 50-cycle) pulsing of ordinary electric light—to make rapidly turning wheels seem to rotate backwards.

sense. But Shoemaker goes on to point out, correctly, that there is a “somewhat more plausible” way of interpreting my claims that avoids these embarrassments: “let us understand these to be contentful states whose contents the subject does not necessarily endorse (as one does endorse the contents of one’s judgments).” This is what I meant, and not only do the passages Shoemaker refers to support this reading; many of my other claims demand it. As Shoemaker notes in the passage just noted, the standard concept implies the existence of a single canonical “subject” whose endorsement is essential for judgments, whereas my theory demands that there be judgment-like episodes that compete for something like eventual endorsement. So Tye is right that my view would be in deep trouble if I meant by “judgment” what philosophers standardly mean, but Shoemaker is right that there is a better interpretation of my language.

I am pleased to see that Tye grants that if my arguments against “phenomenal realism” were successful, I really would have explained consciousness. So it all does come down to whether my “concoction” can take the place of “real phenomenal consciousness.” And he notes that *if* you believe in the latter, there is a huge (but not necessarily unbridgeable) task remaining to be done. I take that to be in itself a ground for suspicion—nothing more—of the concept of phenomenal consciousness. Add to that suspicion the positive reasons I have given for discarding the concept of phenomenal consciousness, and my “concoction” begins to look attractive, I think, especially if it can offer the beginnings of positive accounts (as I claim it does) of some undeniably ubiquitous features of phenomenology—the awfulness of pain, the facts of color preference, etc. Faced with a choice between brute inexplicability on the one hand, and a concoction that begins to explain what needs explaining, I think the choice is obvious.

I entirely agree with **Jackson’s** discussion of the ambiguities hidden in the standard questions about the reality of this and that, ambiguities which excuse—indeed prohibit—straightforward answers. And I accept his breakdown into the A, B, and C questions, and his conclusion that it is the C question, the “truth-maker” question, that is compulsory for materialists, and hence for me.

The truth-maker question for materialists is to specify what it is about a person’s material nature (widely understood so as not to exclude, for instance, material environment) which necessitates their psychological nature. (p. 901)

But when he goes on to proclaim a sharp distinction between “behaviorism, functionalism, eliminativism, and instrumentalism as competing answers,” he sees differences that don’t hold up, in my opinion. My answer, in any event, mixes elements from all of these, and denies that there is any good reason to cleave to a less eclectic answer. Matching the brevity of Jackson’s

thumbnail account of the differentia is my thumbnail account of how to merge all four:

Some traditionally well-regarded mental states should be *eliminated*; in other words, only a reformed folk psychology stands in need of materialistic reduction. Now we must deal with the leftovers: what makes it true that people have the *real* mental states is facts about their *behavioral* dispositions and capacities, but these facts can be perspicuously and efficiently expressed only from the intentional stance, an *instrument* of prediction (and explanation). As for *functionalism*, in its defensible version, it is not really an alternative to behaviorism, but simply a reflection of the tight constraint behavioral capacities (as described from the intentional stance) place on internal states. So let me confirm Jackson's surmise that I am his behaviorist; I unhesitatingly endorse the claim that "necessarily, if two organisms are behaviorally exactly alike, they are psychologically exactly alike."

Jackson thinks this view is a non-starter because of the problems he thinks are raised by Ned Block's imagined gigantic look-up table, which, he notes, I do not address in my book. I do address it elsewhere, however, in "Can Machines Think?"⁷ I offer what I consider to be conclusive grounds for dismissing Block's putative counterexample. Briefly, the sane theorist is entitled to lean a little on some obvious if arguably non-necessary facts. Block claims that a huge look-up table could always "in principle" provide the innards governing any behavioral regularities whatever, and intuition proclaims that we would not consider anything controlled by such a mere look-up table to have psychological states. (If I discovered that you were in fact controlled by such a giant look-up table, I would conclude that you were not a person at all, but an elaborate phony.) But as Alan Turing recognized when he proposed his notoriously behavioristic imitation game, the Turing Test, this "in principle" possibility is not really a possibility at all. A look-up table larger than the visible universe, accessed at speeds trillions of times in excess of the speed of light, is not a serious possibility, and nothing less than that would suffice. What Turing realized is that for *real time responsiveness* in an unrestricted Turing Test, there is only one *seriously* conceivable architecture: one that *creates* its responses locally, on the fly, by processes that systematically uncover the meaning of the inputs, given its previous history, etc., etc. Any such "creative" architecture, I submit, would meet our intuitive demands for intelligence—for "having a psychology"—since it would have plenty of transitional states that accomplished the necessary analysis and construction. So Turing's simple, "external" behavioral test actually puts exactly the right sort of constraint on the internal states:

⁷ "Can Machines Think," in M. Shafto, ed., *How We Know*, New York: Harper and Row, 1985, pp. 121–45. See also "Fast Thinking," chapter 9 in *The Intentional Stance*, Cambridge, Massachusetts: MIT Press, 1987.

not too parochial or “chauvinistic” (to use Block’s term), and not too lax—since Block is simply wrong about the possibility of the giant look-up table. (If “possibilities in principle” count here, they count everywhere, in which case an analogue of Block’s counterexample also goes to establish that no chemist has succeeded in establishing that there is H₂O in the Pacific Ocean.) Now Turing did not bother spelling this all out; it takes someone like Block to dream up the loophole that forces this retrospectively obvious point into explicitness.

Hidden behind this abrupt response to Block is another, which is in some regards more useful, even though it begins by conceding to him the “possibility in principle” of the giant look-up table. Let us suppose that Descartes’ evil demon, with his infinite powers, hand-coded this look-up table—nothing less could accomplish such a super-astronomical task. Now pause to consider how this table is—must be—organized. There is going to be plenty of structure in this vast storehouse of raw data—discernible in the *behavioral* regularities it guarantees. In short, there is, in this vast memory, a *virtual machine* of tremendous complexity. (My computer has only two *real* disk drives, which contain real disks that really spin, and that are divided into sectors that can be written on and read from by electromagnetic processes. But it also has a virtual disk drive which is simply a portion of RAM—it behaves exactly as if it were a disk drive. It “passes the Turing Test for Disk Drives” when subjected to interactions by all the software that I can run on it.) All the internal structure and process that Block finds intuitively to be required for genuine psychology is present in virtual form in the giant look-up table. Is the structure “real”? Yes, in one very important sense: it has to have been designed—barring more miracles. The “r and d” process that the demon used to configure all the data in the look-up table is at least as elaborate as the “r and d” process that configures the actual processes and structures in an actual brain. What I as a card-carrying behaviorist (of this special sort) insist upon is that the difference between virtual and “real” machines is important as biology, important as materials science, but not important as psychology.⁸ For example, if it turned out that left-handers, or women, used virtual brain machines for some cognitive tasks that others used real brain machines for, this would be a fascinating discovery, but it would certainly not impugn the mentality of left-handers or women, even with regard to the competence in question.

So there is no real conflict between my endorsement of behaviorism and my endorsement as well of what we may call virtual machine functionalism—they come to the same thing. Jackson’s brief objection to the latter is that “relatively simple machines” could encode multiple drafts and instantiate Joycean machines, and manifestly they would not be conscious, so I owe

⁸ For more on this theme, see my discussion of Mangan’s views in “Caveat Emptor,” op. cit.

an account of the kind of complexity that necessitates consciousness. I agree, but I thought I gave that account: the sort of complexity that matters is the sort that could sustain a highly contentful and predictive intentional stance characterization. No simple machine can do that—for it amounts to passing the Turing Test. This is particularly apparent, I would have thought, in my discussion of zimpos (pp. 309–14). Among zombies, only a zimbo could conceivably pass the Turing Test; a zimbo has a highly complex profile as an indefinitely higher-order intentional system. But, say I, that is *sufficient* for consciousness (pp. 405–6). Jackson is absolutely right in claiming that I, as a materialist, must answer the truth-maker question. He has not seen that in the course of the book, I have done so, quite forthrightly—if not quite in the format philosophers usually provide their answers.

Shoemaker once remarked to me in a letter responding to “Quining Qualia” that I seemed to have been placed on this earth to raise his blood pressure. I fear my reply to his current essay will confirm that suspicion. Yes, Sydney, I am saying you are the victim of an illusion, and while your thought experiments vividly illustrate your state of mind, they do not show it to be undeluded. Please *don't* take my words with a pinch of salt!

Ned Block figures once again as a ghost at the banquet—his “inverted Earth” thought experiment is supposed to secure a conclusion that is also secured by some of Shoemaker’s own arguments. I like Shoemaker’s version better, since, as he notes, it is shorter, which will hasten its demise. “All inverted spectrum scenarios do,” Shoemaker notes, “is to highlight and dramatize a point that should be obvious without them—that there is an aspect of our experiences... that is not captured by behavioral and functional characterizations.” Too true: that is all such scenarios do—they dramatize a presumption without providing any non-question-begging support for it. If you already believe in the “obvious” point, then you are bound to go along with the gag, but I flatly deny the presumption. Let us look, then, at Shoemaker’s rendition of the case of partial inversion to see where I can drive in the wedge. Post-surgically our subject reports that the shades of color “within a very small range” have “switched places.” As Shoemaker notes, this would be a change in the structure of her quality space, and would be behaviorally detectable in lots of ways. (She would see some vivid color boundaries where she used to see indistinguishably smooth fades or blends, for instance.)

He adds parenthetically that this could not be the result of memory tampering, but this is simply not so. Shoemaker is taken in by the ordinary understanding of memory (a passive storehouse that materials enter *after* they have passed through the pre-processing stage and entered consciousness). We should remind ourselves that any transient informational effect in the course of perceptual processing is from one perspective a memory-effect. As Shoemaker himself puts it, “something akin to memory” is involved in ev-

ery case of motion *perception*, for instance. Even such familiar phenomena as complementary-color after-images are thus in *one* sense the result of memory-tampering. Now it is no escape from this to insist that even if these “early” effects are *called* memory, there must be some later effect that is not “merely” a memory effect (but rather a real qualia inversion), for this manifestly begs the question!

But set that difficulty aside. When Shoemaker goes on to imagine our subject making only a “semantic” accommodation to her newly inverted circumstances, he must be distinguishing such mere semantic accommodations from the host of other dispositional accommodations that might or might not be made. After the crimson-chartreuse switch, and the subject’s merely semantic accommodation to it, does chartreuse ineluctably bring to her mind fair Harvard or not? Suppose not; then Shoemaker is wrong when he claims that after a total series of such inversions “we can suppose that together with the semantic accommodation to that change it would restore the person’s ‘reactive dispositions’ to what they were initially.” When objectively crimson things are put before her, she unhesitatingly calls them crimson, but few or none of her old Harvard-related associations are stimulated. (Some associations may be stimulated *via* the newly made bridge to the *word* “crimson,” but this bridge may well carry only a fraction of the earlier traffic.)

So we must try to suppose all such association-inversions do go along with the surgery. But this ruins the story. For instance, now the subject will not be able to confirm (to herself—let alone to the rest of us!), as we go along on the series of shifts, that, e.g., the earlier qualia-shifts haven’t spontaneously *reverted*. What? Can’t she “just *tell*,” when the latest bout of surgery obliterates a color-boundary problem she’d been noting, which of the two “possible” inversions occurred? No—unless you beg the question. *Ex hypothesi*, nothing will seem to be “strangely” colored post-surgically, even though color contrasts that existed before the surgery will have disappeared. Alternatively, Shoemaker might propose that there would be an initial sense of strangeness or novelty for some of the post-surgical colors, but the “strangeness wears off” (as it must, to restore all reactive dispositions), but then what grounds could he give to support the obligatory claim that the process of “strangeness wearing off” is not in fact the process of “qualia re-inversion”? This was the point I was making (too swiftly, apparently) in my discussion of the imaginary case of the fatal shade of blue that reminded the subject of a car crash (pp. 395ff).

Shoemaker notes in his footnote 6 that there are empirical objections to spectrum inversion that I do not raise. Actually, in my own way I do allude to them in simplified form—in my discussion of color preferences and, again, in the blue car example. Shoemaker proposes to “finesse such objections” by some science fiction, noting that “surely no defender of physical-

ism or functionalism will want to rest his case on the claim that we are not like the envisaged creatures.” Guess again. What *anchors* our naive sense that there are such properties as qualia are the multiple, asymmetrical, interdependent set of reactive dispositions by which we acquaint ourselves with the sensible world. Our sense that the color red has, as it were, an identity, a “personality” all its own is *due to* the host of *different* associations that go with each color. Shoemaker’s envisaged creatures, lacking all such reactive landmarks in their dispositional make-up, would not think they had qualia at all—what it was like to have one sort of experience would not differ at all from what it was like to have a different one! So I claim; if Shoemaker thinks otherwise, he needs to find a supporting argument.⁹

Supposing that there is something in addition to these complex families of reactive dispositions is falling for an illusion, plain and simple—an illusion of subtraction, you might call it. Shoemaker seems to go along with the natural but treacherous assumption that reactive dispositions must involve the person reacting *to* a quale, presented somehow to the reactor, and causing, by its presentation, the reaction. (The *given* is then *taken*.) For instance, here’s how pain works: the pain-networks produce (somewhere central?) the awfulness quale, which is then the very property to which “one” reacts with abhorrence. My view is that this confuses cause and effect; it is the reactions that *compose* the “introspectable property” and it is *through reacting* that one “identifies” or “recognizes” the property. Consider the fascinating case of infant “cuteness.” When you look at a little baby, is there a cuteness *quale* in your experience? It turns out (most probably) that our appreciation of the cuteness of infants has a deep biological base. Some animals have offspring that need significant amounts of parental care after birth; during this vulnerable period, they all have foreshortened “baby-faces”—short snouts, big eyes, etc. Animals that eschew parental care are born with “adult” head shapes. Current theory maintains that the baby-face features evolved to *trigger hard-wired nurturing routines*—in adult dinosaurs (those that cared for their young) as well as in adult birds, horses and us. You couldn’t get much more “reactive behavioral-dispositional” than that. Now do you suppose that *in addition* to the propensity to trigger such parental tenderness-behavior there is a cuteness-quale produced, to which you are reacting when you

⁹ Shoemaker also claims to have shown elsewhere that the similarity relations between properties such as colors “are parasitic on the similarity relations between features of perceptual states of those observers in a way that assigns to those features the role of being qualia.” In the papers in question, the view he actually argues against, the “intentionalist view,” maintains simply that experiences have no introspectable properties other than intentional ones (“Qualities and Qualia: What’s in the Mind?” *Philosophy and Phenomenological Research*, 50, 1990, pp. 109–31. See p. 114). But I agree that in addition to the pure intentional properties (which Shoemaker and I both claim can be handled by functionalism), there are others—the dispositional properties!

find yourself cooing and cuddling? As Ivan Fox has said, a quale thrown into that gap falls right through.

First let me confirm that **Rosenthal's** summary description of the Multiple Drafts Model is an excellent translation of my view into terms that many philosophers will find more familiar. His own view, as he notes, is very close kin to mine—the influences on each other run back and forth through our discussions over the years—and we understand each other's positions better than most other philosophers do. So why do I find it so hard to say why I still prefer my version to his? My discussion (pp. 314–20) of what I find unacceptable in his view is, I think, both the most difficult and least successful part of my book. Perhaps it will help if first I tally up what I think is importantly right in his view.

Like me, he sees the need for a theory that *builds* human consciousness out of functional (unconscious) properties of the information-flow in an animal (or robot) control system. And the property that matters most, he notes, is the one which permits us to distinguish reporting from (merely) expressing, for example: the property created by a general and iterable capacity for self-monitoring. This is nicely captured “from a distance” one might say, in the folk-psychological terms of thoughts about thoughts—HOTs in his terminology.

The problems arise, I think, when one tries to take this everyday characterization too seriously as a working model of actual inner processes, for implicit in folk psychology is the idea of a single, unified subject, whose “endorsement” of these HOTs (recalling Tye and Shoemaker's discussions of judgments) would require a more centralized organization of the brain than the facts permit. The heart of Rosenthal's model has a spurious specificity:

A mental state can have many mental effects without becoming conscious, but not if it causes a HOT. Having a thought about something is one way of being conscious of it. So if one comes noninferentially to have a thought that one is in a particular mental state, that state becomes conscious.

This noninferential process is what we might call a sub-perceptual process, presumably, and such processes occur ubiquitously in the brain—but intuitively only the grandest or most important of these should count as consciousness-ensuring. HOTs may themselves be unconscious (Rosenthal's master stroke), but they have to be, somehow, more globally influential than the merely local-effect, run-of-the-mill self-monitoring outcomes that crowd the brain. How could the distinction be drawn? I think Rosenthal, in spite of his own good warnings about ancient errors, is still somewhat in the thrall of the idea of consciousness as a property that has sharp boundaries in the brain—and in the world. Provoked by his way of putting things, I have

come up with yet another analogy to help make my case that the tradition has been *radically* wrong about what kind of a phenomenon consciousness is.

Andy Warhol anticipated a future in which each person would be famous for fifteen minutes. What he nicely captured in this remark was a *reductio ad absurdum* of a certain (imaginary) concept of fame. Would that be *fame*? Has Warhol described a logically possible world? If we pause to think about his example more carefully than usual, we see that what makes the remark funny is that he has stretched something beyond the breaking point. It is true, no doubt, that thanks to the mass media, fame can be conferred on an anonymous citizen almost instantaneously (Rodney King comes to mind), and thanks to the fickleness of public attention, fame can evaporate almost as fast. But Warhol's rhetorical exaggeration of this fact carries us into the absurdity of Wonderland. We have yet to see an instance of someone being famous for just fifteen minutes, and in fact we never will. Let some one citizen be viewed for fifteen minutes by hundreds of millions of people, and then—unlike Rodney King—be utterly forgotten. To call that fame would be to misuse the term (ah yes, an “ordinary language” move, and a good one, if used with discretion).

If that is not obvious, then let me raise the ante: could a person be famous—not merely attended-to-by-millions-of-eyes, but famous—for five seconds? Every day there are in fact hundreds if not thousands of people who pass through the state of being viewed, for a few seconds, by millions of people. Consider the evening news on television, presenting a story about the approval of a new drug. Accompanying Dan Rather's voice-over, an utterly anonymous nurse is seen (by millions) plunging a hypodermic into the arm of an utterly anonymous patient. Now that's fame—right? Of course not. Being seen on television and being famous are different sorts of phenomena; the former has technologically sharp edges that the latter entirely lacks.

What I am arguing, in my attack on the Cartesian theater, is that being an item in consciousness is *not at all* like being on television; it is, rather, a species of mental fame. Almost literally. Consciousness is cerebral celebrity—nothing more and nothing less. Those contents are conscious that persevere, that monopolize resources long enough to achieve certain typical and “symptomatic” effects—on memory, on the control of behavior and so forth. Not every content can be famous, for in such competitions there must be more losers than winners. And unlike the world of sports, winning is everything. There are no higher honors to be bestowed on the winners, no Hall of Fame to be inducted into. In just the way that a Hall of Fame is a redundant formality (if you are already famous, election is superfluous, and if you are not, election probably won't make the difference), there is no induction or transduction into consciousness beyond the influence already secured by winning the competition and thereby planting lots of hooks into the ongoing operations of the brain.

Instantaneous fame is a disguised contradiction in terms, and it follows from my proposed conception of what consciousness is that an instantaneous flicker of consciousness is also an incoherent notion. Those philosophers who see me as underestimating the power of future research in neuroscience when I claim that no further discoveries from that quarter could establish that there was indeed a heretofore undreamt-of variety of evanescent—but genuine—consciousness might ask themselves if I similarly undervalue the research potential of sociology when I proclaim that it is inconceivable that sociologists could discover that Andy Warhol’s prediction had come true. This could only make sense, I submit, to someone who is still covertly attached to the idea that consciousness (or fame) is the sort of semi-mysterious property that might be discovered to be present by the tell-tale ticking of the phenomometer or famometer (patent pending!).¹⁰

Consider a question nicely parallel to the Orwell-Stalin stumpers I declare to be vacuous: did Jack Ruby become famous before Lee Harvey Oswald died? Hmm. Well, hundreds of millions witnessed him shooting Oswald on “live” television, and certainly he *subsequently* became famous, but did his fame begin at the instant his gun-toting arm hove into view, while Oswald was still alive and breathing? I for one do not think there is any fact we don’t already know that could conceivably shed light on this question of event ordering.

Cerebral celebrity is tantamount to being the object of a HOT, but is a more promising model for future development. Gone is the specific class of internal episodes, the thoughts about thoughts, identifiable in principle by their vehicular shape and their ingredients (like boxcars containing boxcars). In their place is something that would do the same work—structures that

¹⁰ Owen Flanagan has made the claim in detail, in *Consciousness Reconsidered*, Cambridge, Massachusetts: MIT Press, 1992, esp. pp. 14–15, pp. 82–83. He supposes, for instance, that the 40-hertz phase-locked oscillation championed by Singer, von der Malsburg, and others, and suggested by Crick and Koch as a key mechanism in consciousness, might serve as the missing ingredient that could trump my short-sighted verificationism. “Never say never” he advises me. Well, now that neuroscience’s love affair with 40-hertz is beginning to cool off (because of problems that were always inherent in the idea—people were asking too much of it), let’s look more closely at the prospects. If Crick and Koch are taken to be proposing a mechanism for *securing cerebral celebrity*—the underlying mechanism by which some contents win and others lose in competition—then they are not offering a rival view, but merely specifying details I left blank in my sketch. Ignore the problems and suppose they are right (I would love to have an account of the detailed mechanisms, after all). Notice that it is logically impossible for a 40-hertz oscillation mechanism to resolve temporal onset questions below the two-pulse minimum of 25msec, and any plausible competition-for-entrainment model would surely require considerably more time—moving us inexorably into a window of indeterminacy of the size I postulated: several hundred milliseconds. Alternatively, if Flanagan supposes that Crick and Koch are claiming that 40-hertz entrainment *causes a subsequent state-change* (which could be an instantaneous transduction into some new medium, for instance), then he is simply insisting on the very concept of consciousness I am challenging—that consciousness requires entrance into a “charmed circle.”

more or less guarantee the mutual accessibility of various contents to each other and to the mechanisms that can use them. That was to have been the *function* of thoughts about thoughts, after all, but this way we eliminate the middleman who has to *have* those HOTs and know what to do about them.