Sweet Dreams

Philosophical Obstacles to a Science of Consciousness

Daniel C. Dennett

Sweet Dreams

Philosophical Obstacles to a Science of Consciousness

Daniel C. Dennett

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

7 A *Fantasy Echo* Theory of Consciousness

Many years ago, a friend told me about a professor of literature who was puzzled by a final examination essay in which a student went on at some length about *fantasy echo* poetry.[1] The professor called the student in and queried him about his curiously evocative but unexplained epithet. What on earth was the student talking about and where, if one might ask, had he picked up this idea? "From your lectures, of course!" the student replied. The professor was dumfounded, but soon enough got to the bottom of the mystery: he had often referred in his lectures to late-nineteenth-century works in the *fin de siècle* style.

This gem of serendipitous misperception has been rattling around in my brain for several decades. A few months ago, it occurred to me that it really deserved a new career, and that the time was ripe for reincarnation. Eureka! For those same decades I had been yearning for a sailboat; this was the year, at last, to

1. This chapter is drawn with many deletions and revisions from a lecture I gave at the Consciousness in London Conference at The Kings College, London, April 24, 1999. That lecture included many of the passages included in earlier chapters, and these have largely been excised, leaving behind just a few expressions and arguments that may help clarify my view.

buy a boat, and its name shall be *Fantasy Echo*. What a perfect name for a 1999 dreamboat! But for various good reasons it turns out that 1999 is *not* a good year for me to buy a sailboat (chartering, once again, must suffice); it appears that my relief from slooplessness, as Quine once put it, will have to await the next millennium. What my university's fund-raisers call a *great naming opportunity* was going to slip away from me, unexploited. What a pity!

I was recounting all this to another friend recently, who startled me by pointing out that I already owned, and had been working on for years, something that could with even more justice be named *Fantasy Echo*: my theory of human consciousness. So with a little help from my friends, I am happy to unveil, at this 1999 conference on theories of consciousness, my updated and newly renamed *Fantasy Echo* theory of consciousness.

1 Fleeting Fame

This is the theory that went by the name of the Multiple Drafts Model in 1991, and has more recently been advertised by me as the "fame in the brain" (or "cerebral celebrity") model (1996b, 1998a, 2001a). The basic idea is that consciousness is more like fame than television; it is *not* a special "medium of representation" in the brain into which content-bearing events must be "transduced" in order to become conscious. It is rather a matter of content-bearing events in the brain achieving something a bit like fame in competition with other fame-seeking (or at any rate potentially fame-finding) events.

But of course consciousness couldn't be *fame*, exactly, in the brain, since to be famous is to be a shared intentional object *in*

the consciousnesses of many folk, and although the brain is usefully seen as composed of hordes of *homunculi*, imagining them to be *au courant* in just the way they would need to be to elevate some of their brethren to cerebral celebrity is going a bit too far—to say nothing of the problem that it would install a patent infinite regress in my theory of consciousness. The looming infinite regress can be stopped the way such threats are often happily stopped, not by abandoning the basic idea but by softening it. As long as your *homunculi* are more stupid and ignorant than the intelligent agent they compose, the nesting of homunculi within homunculi can be finite, bottoming out, eventually, with agents so unimpressive that they can be replaced by machines.

So consciousness is not so much *fame*, then, as *influence*—a species of relative "political" power in the opponent processes that eventuate in ongoing control of the body. In some oligarchies, perhaps, the only way to achieve political power is to be *known by the King*, dispenser of all powers and privileges. Our brains are more democratic, indeed anarchic. In the brain there is no King, no Official Viewer of the State Television Program, no Cartesian Theater, but there are still plenty of *quite* sharp differences in political power exercised by contents over time. What a theory of consciousness needs to explain is how some relatively few contents become elevated to this political power, while most others evaporate into oblivion after doing their modest deeds in the ongoing projects of the brain.

Why is *this* the task of a theory of consciousness? Because that is what conscious events *do*. They hang around, monopolizing time "in the limelight"—but we need to explain *away* this seductive metaphor, and its kin, the searchlight of attention, by explaining the *functional* powers of attention-*grabbing* without

presupposing a single attention-*giving* source. That is the point of what I call the Hard Question: And Then What Happens? Postulate activity in whatever neural structures you please as the necessary and sufficient condition for consciousness, but then take on the burden of explaining why *that* activity ensures the political power of the events it involves.

The attractiveness of the idea of a special medium of consciousness is not simply a persistent hallucination. It is not *entirely* forlorn, as we can see by pursuing the analogy with fame a bit further. Fame—in the world, not in the brain—is not what it used to be. The advent of new media of communication has in fact radically changed the nature of fame, and of political power, in our social world, and something interestingly analogous may have happened in the brain. That, in any case, is my speculative proposal. As I have argued over and over again, *being in consciousness* is not like *being on television*; one can be on television and be seen by millions of viewers, and still not be famous, because one's television debut does not have the proper *sequelae*. Similarly, there is no special area in the brain where representation is, by itself, sufficient for consciousness. It is always the *sequelae* that make the difference. (And Then What Happens?)

My inspiration for the fame-in-the-brain analogy was, of course, Andy Warhol:

> In the future, everybody will be famous for fifteen minutes.

What Warhol 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 it more carefully than usual, we see that something has been stretched 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, 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 citizen be viewed for fifteen minutes or less 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 *for five seconds* (not merely attended-to-by-millions of eyes but famous)? There are in fact hundreds if not thousands of people who every day pass through the state of being viewed, for a few seconds, by millions of people. Consider the evening news, presenting a story about the approval of a new drug. An utterly anonymous doctor is seen (by millions) plunging a hypodermic into the arm of an utterly anonymous patient—that's being on television, but it isn't fame!

Several philosophers have risen to the bait of my rhetorical question and offered counterexamples to my implied claim about the duration of fame. Here is how somebody could be famous for fifteen seconds: he goes on international TV, introduces himself as the person who is about to destroy our planet and thereupon does so. Oh, they got me! But notice that this example actually works in my favor. It draws attention to the importance of the normal *sequelae*: the only way to be famous for less than a longish time is to destroy the whole world in which your fame would otherwise reverberate. And if anybody

wanted to cavil about whether that was *really* fame, we could note how the question could be resolved in an extension of the thought experiment. Suppose our antihero presses the button and darn, no nuclear explosion! *And then what happens?* The world survives, and in it we either observe the normal *sequelae* of fame or we don't. In the latter case, we would conclude, retrospectively, that our candidate's bid for fame had simply failed, in spite of his widely broadcast image. (Maybe nobody was watching, or paying attention.) The important point of the analogy is that consciousness, like fame, is a *functionalistic* phenomenon: handsome is as handsome *does*.

The importance of such echoes, of *reverberation*, of *return-trips*, of *reminding*, of *recollectability*, is often noted by writers on consciousness. Here is Richard Powers, for instance:

> To remember a feeling without being able to bring it back. This seemed to me as close to a functional definition of higher-order consciousness as I would be able to give her. (1995, p. 228)

But this is "higher-order" consciousness, isn't it? What about "lower-order" consciousness? Might the echo-capacity be wholly absent therefrom? The idea that we can identify a variety of consciousness that is logically independent of the echo-making power has many expressions in the recent literature. It is even tempting to suppose that this lower or simpler variety of consciousness is somehow a normal precondition for echo-making. It is, perhaps, the very feature that echoes when there *are* echoes. A particularly popular version is Ned Block's proposed distinction between *phenomenal consciousness* and *access consciousness*. Fame in the brain provides, perhaps, a useful way of thinking about the "political" access that some contents may have to the reins of power in the ongoing struggle to control

the body, but it has nothing to say about the brute, lower-order, *what-it-is-like-ness* of phenomenal consciousness.

What is it like *to whom*? As I have often said, in criticism of Block's attempted distinction, once you shear off all implications about "access" from phenomenal consciousness, you are left with something apparently indistinguishable from phenomenal *un*consciousness. Consider an example. As a left-handed person, I can wonder whether I am a left-hemisphere-dominant speaker or a right-hemisphere-dominant speaker or something mixed, and the only way I can learn the truth is by submitting myself to objective, "third-person" testing. I don't "have access to" this intimate fact about how my own mind does its work. It escapes all my attempts at introspective detection, and might, for all I know, shunt back and forth every few seconds without my being any the wiser (see chapter 4). This is just one of many—indeed countless—"intrinsic" properties that the events occurring in my brain have that, by being entirely inaccessible to me, are paradigms of unconscious properties. The challenge facing those who want to claim that some among these "intrinsic" properties are the properties of phenomenal *consciousness* is to show what makes them different (without making any appeal to "access" or echo-making power).

It is the echo-making power, after all, that we invariably appeal to when we try to motivate the claims we make about the consciousness not just of others, but of ourselves. Proust famously elevated the really quite delicate aroma of *madeleines*, almond cookies, for its power to provoke in him vivid memories and emotions from his childhood. The inviting aroma of classroom library paste (safe, edible!) has a similar effect on me. Contrast it with the aroma of, say, the Formica desktop at which

I sat in second grade. But, you protest, it doesn't have an aroma! Well, it does, but not an evocative one for me, not one whose coming and going I can even detect under normal circumstances. It is an aroma that is, at best, *subliminal*—beneath the threshold of my consciousness. What if, nevertheless, it could be shown that the presence or absence of that desktop in my olfactory environment had a subtle biasing effect on my performance on some cognitive task—it might, for instance, bias me in favor of thinking first of the most classroom-relevant meanings of ambiguous words. If so, we'd have a quandary: was this, like "blindsight in normals," a case of *un*conscious echo, or a proof that the aroma of the Formica was indeed part of the background (the Background, to some) of my boyhood consciousness? Either way, it is the presence of an echo, however faint, that provides whatever motivation the latter view has. The believer in phenomenal consciousness stripped even of this echo-making power has a tough sell: the coming and going of the aroma is a change in phenomenal consciousness in spite of the subject's total obliviousness—lack of access—to it.

I said earlier that the idea of consciousness depending on a special medium of representation in the brain is not entirely forlorn, and with these clarifications of our intuitions behind us, I am ready to tackle that issue. Television and fame are two entirely different sorts of things—one's a medium of representation and one isn't—but the sorts of fame made possible by television are interestingly different from earlier sorts of fame, as we have recently been told 'til we're sick of it. Consider the phenomena of Princess Diana, O. J. Simpson, and Monica Lewinsky. In each case a recursive positive feedback became established, dwarfing the initial triggering event and forcing the world to wallow in wallowing in wallowing in reactions to

reactions to reactions to the coverage in the media of the coverage in the media of the coverage in the media, and so forth. Did similar fame-phenomena occur in the preelectronic age? The importance of publicity had been appreciated for millennia—secret coronations, for instance, have always been shunned, for the obvious reasons. There have long been sites of recursive reaction, such as the page of letters to the editor in the *Times* (of London, and to a lesser extent, the *New York Times*). But these were still relatively slow, "narrow band" (as we say nowadays) channels of communication, and they reached a small but influential segment of the populace. In the preelectronic age, were there people who were famous for being famous? It is the capacity for the combined modern media to capture *anything* and turn it into a ubiquitously "accessible" or "influential" topic through sheer echoic amplification that strikes some observers as a novel (and perhaps alarming) social phenomenon, and I want to suggest that a similar family of innovations in the brain may lie behind the explosive growth in *reflective power* that I take to be the hallmark of consciousness.

2 Instant Replay

At this point in earlier discussions of this topic, the loyal opposition notes that I am impressed—perhaps overimpressed—with the power of *self*-consciousness, or *reflective* or *introspective* consciousness, at the expense of just plain animal *sentience* or, echoing Block again, *phenomenal* consciousness, but when I talk of reflective power here, I am *not* talking about the highly intellectual (and arguably language-dependent) capacity for—shall we say—*musing* about our *musings*. I'm talking about the capacity of a dog, for instance, to be reminded of its owner or its

tormentor by an aroma that provokes an echo that provokes a reidentification. But if *that* is all I'm talking about, then the objection still stands: my notorious claim that human consciousness is largely a culturally borne "meme machine" is refuted by the example of the dog!

Not so fast. It would be refuted—or at least somewhat displaced—by the dog if we could be sure that the reminding aroma really does operate by triggering in the dog the sort of echoic, Proustian events that we report to each other. But there may well be simpler hypotheses that explain the dog's delighted (or hostile) arousal when the aroma hits its nostrils. What else might be going on when a dog "recognizes" somebody by aroma? Does—can—the dog *recollect* the earlier encounter? Are dogs capable of *episodic memory*, or is there just summoned up in the dog a "visceral" echo, of either joy or fear? Minimal recognition of this sort need not involve recollection in our own case, so it need not involve recollection in the case of other species. It need not bring in its ensemble the Proustian trappings and surroundings of the earlier encounter that normally—but not always—decorate our own episodes of episodic memory.

These added details are not just decorations, of course. We human beings rely on them to confirm to ourselves that we are indeed remembering, and not just imagining or guessing. Did I ever meet C. I. Lewis? Yes, once. He was a very old man, and I was a freshman at Wesleyan, in 1959–1960, and he came to give a lecture or two there. I didn't know anything about him at the time, but my philosophy professor had encouraged me to attend, just to see a great man. It was in the Honors College, I recall, and he sat down to read his paper (and I was sitting on the north side of the room facing him, as I recall)—but I don't recall what his paper was about at all. I was more impressed by

the respect he was shown by all in attendance than by anything he actually said. The next year, I read two of his books in Roderick Firth's epistemology class at Harvard, which cemented the recollections I'm now reviewing, and of course completely contaminated any memories I might have otherwise thought I had about what he'd said the year before. Now perhaps dogs have similar reflective episodes in their inner lives; if they do, then surely they are just as conscious as we are, in every sense. But I hypothesize—this is the empirical going-out-on-a-limb part of my view—that they do not. Events in their echo chambers damp down to nothing after a couple of reverberations, I suspect. Why? Because they do not need such an echo chamber for anything, and it is not a by-product of anything else they need, and it would be expensive. For nonhuman animals, I suspect, efficiency and timeliness are the desiderata that dictate short, swift, ballistic trajectories of contents. As the business consultants say, the goal is: *Up your throughput!*

But we human beings got sidetracked. We developed a habit of "replaying events in our minds" over and over, and this habit, initially "wasteful" of time and energy, is very likely the source of one of our greatest talents: episodic memory and "one-shot learning" that is not restricted to special cases. (The Garcia Effect is one such special case; rats made nauseous while eating a novel-smelling food have a remarkable Proust effect indeed: they develop an instant distaste for anything with that smell.)

Scientists who use animals in experiments know that in order to teach a new habit, a new discrimination, to an animal, they will typically have to repeat a training or conditioning episode, sometimes three or four times, sometimes hundreds or even thousands of times, before the animal reliably extracts the desired content. There is "one-shot" learning of particularly

galvanizing lessons, but can the learner later recall the episode or just the lesson? Might it be that our familiar human talent for reliving pastel versions of our earlier experiences is in large part a learned trick? The events we can readily recall from our lives are actually a rather limited subset of what happens during our waking lives. (Or can you dredge up what you were thinking about while you brushed your teeth last Wednesday?) Episodes in real life happen just once, without (external) repetition, but perhaps our habit of immediately reviewing or rehearsing whatever grabs our attention strongly is a sort of inadvertent self-conditioning that drives these events into the imaginary "storehouse of episodic memory" (it is certainly not an organ or subsystem of the brain). The hypothesis is that until you've acquired the habit of such "instant replay," permitting the choice bits of daily life to reverberate for a while in the brain, you won't have any episodic memory. This could account for "infantile amnesia," of course, and a further, independent hypothesis is that it is a humans-only phenomenon, an artifact of habits of self-stimulation that other species can't acquire in the normal course of things.

Episodic memory is not for free. One idea is that it is the very echoic power that makes episodic memory possible. Animals remember thanks to multiple repetitions of stimuli in the world. We remember, it seems, one-shot, but really, it isn't just one-shot. What we remember is stuff that has been played and replayed and replayed obsessively in our brains. (Note that a feature, not a bug, in this account is that although some repetition is indeed all too familiar to us as *conscious* repetition, the repetition that *elevates* a content to the clout of conscious recallability is largely not conscious. Indeed, there is no need for a sharp dividing line between conscious and unconscious

repetitions. No bright line need distinguish true fame from mere behind-the-scenes influence.) It is the echo that creates the capacity for long-term episodic memory. We are used to using these trappings as confirmations of our own convictions that we *are* recollecting. Did you ever meet Carnap? Yes, I reply. "It was at UCLA, in 1965 or '66, I would guess. It was in the corridor outside the philosophy department, and as best I recall, Alfred Tarski and Richard Montague were talking with Carnap. I asked somebody who the people with Tarski were, and when they told me, I just couldn't resist going up and barging in and just shaking their hands."

This instant-reply habit itself has its amusing analogue in the world of electronic media. Before the existence of *videotape*, being on television was not a particularly echoic phenomenon. The programs were broadcast "live" and once they were over, they were over—echoing for awhile in the memories and discussions of the audience, but quick to damp out and slide into oblivion. Newsreels at the cinema were different. Newspapers were different. They preserved for review the events of the day. Until memory was added, radio and television were *not* the sort of media that could provide a suggestive hint about the structure—and media—of consciousness, since their contents were utterly evanescent, no better, really, than the flitting images on the blank wall of the *camera obscura*—except in the memories of those who witnessed them.

Let me sum up. I have ventured (1) the empirical hypothesis that our capacity to relive or rekindle contentful events is the most important feature of consciousness—indeed, as close to a defining feature of consciousness as we will ever find; and (2) the empirical hypothesis that this echoic capacity is due in large part to habits of self-stimulation that we pick up from human

culture, that the Joycean machine in our brains is a virtual machine made of memes. These are independent claims. If the meme-hypothesis were roundly defeated by the discovery—the confirmation—of just such echoic systems at play in the brains of nonhuman animals, I would then agree, for that very reason, that the species having those echo-chambers were conscious in just about the way we are—because that's what I say consciousness is. The price I'd pay for that verdict is the defeat of my bold claim about software and *virtual* machines, but I'd still be getting a bargain, since the other side would be relying on the fame theory of consciousness *as a theory of consciousness* in order to establish the relevance—to riddles about consciousness—of their discoveries.

8 Consciousness: How Much Is That in *Real* Money?

Consciousness often seems to be utterly mysterious. I suspect that the principle cause of this bafflement is a sort of accounting error that is engendered by a familiar series of challenges and responses. A simplified version of one such path to mysteryland runs as follows:

Phil: What is consciousness?

Sy: Well, some things—such as stones and can-openers—are utterly lacking in any *point of view*, any *subjectivity* at all, while other things—such as you and me—do have points of view: private, perspectival, interior ways of being apprised of some limited aspects of the wider world and our bodies' relations to it. We lead our lives, suffering and enjoying, deciding and choosing our actions, guided by this "first-person" access that we have. To be conscious is to be an agent with a point of view.

Phil: But surely there is more to it than that! A cherry tree has limited access to the ambient temperature at its surface, and can be (mis-)guided into blooming inopportunely by unseasonable warm weather; a robot with video camera "eyes" and microphone "ears" may discriminate and respond aptly to hundreds of different aspects of its wider world; my own immune system

can sense, discriminate, and respond appropriately (for the most part) to millions of different eventualities. Each of these is an agent (of sorts) with a point of view (of sorts) but none of them is conscious.

Sy: Yes, indeed; there is more. We conscious beings have capabilities these simpler agents lack. We don't just notice things and respond to them; we *notice* that we notice things. More exactly, among the many discriminative states that our bodies may enter (including the states of our immune systems, our autonomic nervous systems, our digestive systems, and so forth), a subset of them can be discriminated in turn by higher-order discriminations which then become sources of guidance for higher-level control activities. In us, this recursive capacity for self-monitoring exhibits no clear limits—beyond those of available time and energy. If somebody throws a brick at you, you see it coming and duck. But you also discriminate the fact that you *visually* discriminated the projectile, and can then discriminate the further fact that you can tell visual from tactile discriminations (usually), and then go on to reflect on the fact that you are also able to recall recent sensory discriminations in some detail, and that there is a difference between experiencing something and recalling the experience of something, and between thinking about the difference between recollection and experience and thinking about the difference between seeing and hearing, and so forth, 'til bedtime.

Phil: But surely there is more to it than that! Although existing robots may have quite paltry provisions for such recursive self-monitoring, I can readily imagine this particular capacity being added to some robot of the future. However deftly it exhibited its capacity to generate and react appropriately to

"reflective" analyses of its underlying discriminative states, it wouldn't be conscious—not the way we are.

Sy: Are you sure you can imagine this?

Phil: Oh yes, absolutely sure. There would be, perhaps, some sort of *executive* point of view definable by analysis of the power such a robot would have to control itself based on these reactive capacities, but this robotic subjectivity would be a pale shadow of ours. When it uttered "it seems to me . . . ," its utterances wouldn't really mean anything—or at least, they wouldn't mean what I mean when I tell you what it's like to be me, how things seem to me.

Sy: I don't know how you can be so confident of that, but in any case, you're right that there is more to consciousness than that. Our discriminative states are not just discriminable; they have the power to provoke preferences in us. Given choices between them, we are not indifferent, but these preferences are themselves subtle, variable, and highly dependent on other conditions. There is a time for chocolate and a time for cheese, a time for blue and a time for yellow. In short (and oversimplifying hugely), many if not all of our discriminative states have what might be called a dimension of affective valence. We care which states we are in, and this caring is reflected in our dispositions to change state.

Phil: But surely there is more to it than that! When I contemplate the luscious warmth of the sunlight falling on that old brick wall, it's not just that I prefer looking at the bricks to looking down at the dirty sidewalk beneath them. I can readily imagine outfitting our imaginary robot with built-in preferences for every possible sequence of its internal states, but it would still not have anything like my conscious *appreciation* of the visual poetry of those craggy, rosy bricks.

Sy: Yes, I grant it; there is more. For one thing, you have metapreferences; perhaps you wish you could stop those sexual associations from interfering with your more exalted appreciation of the warmth of that sunlight on the bricks, but at the same time (roughly) you are delighted by the persistence of those saucy intruders, distracting as they are, but . . . what was it you were trying to think about? Your stream of consciousness is replete with an apparently unending supply of associations. As each fleeting occupant of the position of greatest influence gives way to its successors, any attempt to halt this helter-skelter parade and monitor the details of the associations only generates a further flood of evanescent states, and so on. Coalitions of themes and projects may succeed in dominating "attention" for some useful and highly productive period of time, fending off would-be digressions for quite a while, and creating the sense of an abiding self or ego taking charge of the whole operation. And so on.

Phil: But surely there is more to it than that! And now I begin to see what is missing from your deliberately evasive list of additions. All these dispositions and metadispositions to enter into states and metastates and metametastates of reflection about reflection could be engineered (I dimly imagine) into some robot. The trajectory of its internal state-switching could, I suppose, look strikingly similar to the "first-person" account I might give of my own stream of consciousness, but those states of the robot would have no actual *feel*, no *phenomenal* properties at all! You're still leaving out what the philosophers call qualia.

Sy: Actually, I'm still leaving out *lots* of properties. I've hardly begun acknowledging all the oversimplifications of my story so

far, but now you seem to want to preempt any further additions from me by insisting that there are properties of consciousness that are altogether different from the properties I've described so far. I thought I *was* adding "phenomenal" properties in response to your challenge, but now you tell me I haven't even begun. Before I can tell if I'm leaving these properties out, I have to know what they are. Can you give me a clear example of a phenomenal property? For instance, if I used to like a particular shade of yellow, but thanks to some traumatic experience (I got struck by a car of that color, let's suppose), that shade of yellow now makes me very uneasy (whether or not it reminds me explicitly of the accident), would this suffice to change the *phenomenal* properties of my experience of that shade of yellow?

Phil: Not necessarily. The *dispositional* property of making you uneasy is not itself a phenomenal property. Phenomenal properties are, by definition, not dispositional but rather intrinsic and accessible only from the first-person point of view . . .

Thus we arrive in mysteryland. If you *define* qualia as *intrinsic properties* of experiences considered in isolation from all their causes and effects, logically independent of all dispositional properties, then they are logically guaranteed to elude all broad functional analysis—but it's an empty victory, since there is no reason to believe such properties exist. To see this, compare the *qualia* of experience to the *value* of money. Some naive Americans can't get it out of their heads that dollars, unlike francs and marks and yen, have *intrinsic value* ("How much is that in *real* money?"). They are quite content to "reduce" the value of other currencies in dispositional terms to their exchange rate with dollars (or goods and services), but they have a hunch that dollars are different. Every dollar, they declare, has something

logically independent of its functionalistic exchange powers, which we might call its *vim*. So defined, the *vim* of each dollar is guaranteed to elude the theories of economists forever, but we have no reason to believe in it—aside from the heartfelt hunches of those naive Americans, which can be explained without being honored.

Some participants in the consciousness debates simply demand, flat out, that their intuitions about phenomenal properties are a nonnegotiable starting point for any science of consciousness. Such a conviction must be considered an interesting symptom, deserving a diagnosis, a datum that any science of consciousness must account for, in the same spirit that economists and psychologists might set out to explain why it is that so many people succumb to the potent illusion that money has intrinsic value.

> There are many properties of conscious states that can and should be subjected to further scientific investigation right now, and once we get accounts of them in place, we may well find that they satisfy us as an explanation of what consciousness is. After all, this is what has happened in the case of the erstwhile mystery of what *life* is. Vitalism—the insistence that there is some big, mysterious extra ingredient in all living things—turns out to have been not a deep insight but a failure of imagination. Inspired by that happy success story, we can proceed with our scientific exploration of consciousness. If the day arrives when all these acknowledged debts are paid and we plainly see that something big is missing (it should stick out like a sore thumb at some point, if it is really important) those with the unshakable hunch will get to say they told us so. In the meantime, they can worry about how to fend off the diagnosis that they, like the vitalists before them, have been misled by an illusion.

References