

Daniel C. Dennett

Corralling the Quicksilver

'Qualia' is an unfamiliar term for something that could not be more familiar to each of us: the *ways things seem to us*. As is so often the case with philosophical jargon, it is easier to give examples than to give a definition of the term. Look at a glass of milk at sunset; *the way it looks to you*—the particular, personal, subjective visual quality of the glass of milk is the *quale* of your visual experience at the moment. The *way the milk tastes to you then* is another, gustatory *quale*, and *how it sounds to you* as you swallow is an auditory *quale*. These various 'properties of conscious experience' are prime examples of *qualia*. Nothing, it seems, could you know more intimately than your own qualia; let the entire universe be some vast illusion, some mere figment of Descartes' evil demon, and yet what the figment is *made of* (for you) will be the *qualia* of your hallucinatory experiences. Descartes claimed to doubt everything that could be doubted, but he never doubted that his conscious experiences had qualia, the properties by which he knew or apprehended them.

The verb 'to quine' is even more esoteric. It comes from *The Philosophical Lexicon* (Dennett 1987b), a satirical dictionary of eponyms: 'quine, v. To deny resolutely the existence or importance of something real or significant'. At first blush it would be hard to imagine a more quixotic quest than trying to convince people that there are no such properties as qualia; hence the ironic title of this chapter. But I am not kidding.

My goal is subversive. I am out to overthrow an idea that, in one form or another, is 'obvious' to most people—to scientists, philosophers, lay people. My quarry is frustratingly elusive; no sooner does it retreat in the face of one argument than 'it' reappears, apparently innocent of all charges, in a new guise.

Which idea of qualia am I trying to extirpate? Everything real has properties, and since I do not deny the reality of conscious experience, I grant that conscious experience has properties. I grant moreover that each person's states of consciousness have properties in virtue of which those states have the experiential content that they do. That is to say, whenever someone experiences something as being one way rather than another, this is true in virtue of some property of something happening in them at the time, but these properties are so unlike the properties traditionally imputed to consciousness that it would be grossly misleading to call any of them the long-sought qualia. Qualia are supposed to be *special* properties, in some hard-to-define way. My claim—which can only come into focus as we proceed—is that conscious experience has no properties that are special in *any* of the ways qualia have been supposed to be special.

The standard reaction to this claim is the complacent acknowledgement that while some people may indeed have succumbed to one confusion or fanaticism or another, one's own appeal to a modest, innocent notion of properties of subjective experience is surely safe. It is just that presumption of innocence I want to overthrow. I want to shift the burden of proof, so that anyone who wants to appeal to private, subjective properties has to prove first that in so doing they are *not* making a mistake. This status of *guilty until proven innocent* is neither unprecedented nor indefensible (so long as we restrict ourselves to concepts). Today, no biologist would dream of supposing that it was quite all right to appeal to some innocent concept of *élan vital*. Of course one *could* use the term to mean something in good standing; one could use *élan vital* as one's name for DNA, for instance, but this would be foolish nomenclature, considering the deserved suspicion with which the term is nowadays burdened. I want to make it just as uncomfortable for anyone to talk of qualia—or 'raw feels' or 'phenomenal properties' or 'subjective and intrinsic properties' or 'the qualitative character' of experience—with the standard presumption that they, and everyone else, knows what on earth they are talking about.¹

What are qualia, *exactly*? This obstreperous query is dismissed by one author ('only half in jest') by invoking Louis Armstrong's legendary reply when asked what jazz was: 'If you got to ask, you ain't never gonna get to know' (Block 1978, 281). This amusing tactic perfectly illustrates the presumption that is my target. If I succeed in my task, this move, which passes muster in most circles today, will look as quaint and insupportable as a jocular appeal to the ludicrousness of a living thing—a living thing, mind you!—doubting the existence of *élan vital*.

My claim, then, is not just that the various technical or theoretical concepts of qualia are vague or equivocal, but that the source concept, the 'pre-theoretical' notion of which the former are presumed to be refinements, is so thoroughly confused that, even if we undertook to salvage some 'lowest common denominator' from the theoreticians'

proposals, any acceptable version would have to be so radically unlike the ill-formed notions that are commonly appealed to that it would be tactically obtuse—not to say Pickwickian—to cling to the term. Far better, tactically, to declare that there simply are no qualia at all.²

Rigorous arguments only work on well-defined materials, and, since my goal is to destroy our faith in the pre-theoretical or 'intuitive' concept, the right tools for my task are intuition pumps, not formal arguments. What follows is a series of fifteen intuition pumps, posed in a sequence designed to flush out—and then flush away—the offending intuitions. In the next section, I will use the first two intuition pumps to focus attention on the traditional notion. It will be the burden of the rest of the paper to convince you that these two pumps, for all their effectiveness, mislead us and should be discarded. In the following section, the next four intuition pumps create and refine a 'paradox' lurking in the tradition. This is not a formal paradox, but only a very powerful argument pitted against some almost irresistibly attractive ideas. In the next section, six more intuition pumps are arrayed in order to dissipate the attractiveness of those ideas, and the following section drives this point home by showing how hapless those ideas prove to be when confronted with some real cases of anomalous experience. This will leave something of a vacuum, and in the final section three more intuition pumps are used to introduce and motivate some suitable replacements for the banished notions.

The Special Properties of Qualia

Intuition pump 1: watching you eat cauliflower. I see you tucking eagerly into a helping of steaming cauliflower, the merest whiff of which makes me feel faintly nauseated, and I find myself wondering how you could possibly relish *that taste*, and then it occurs to me that, to you, cauliflower probably tastes (must taste?) different. A plausible hypothesis, it seems, especially since I know that the very same food often tastes different to me at different times. For instance, my first sip of breakfast orange juice tastes much sweeter than my second sip if I interpose a bit of pancakes and maple syrup, but after a swallow or two of coffee, the orange juice goes back to tasting (roughly? exactly?) the way it did with the first sip. Surely we want to say (or think about) such things, and surely we are not wildly wrong when we do, so . . . surely it is quite OK to talk of *the way the juice tastes to Dennett at time t*, and ask whether it is just the same as or different from *the way the juice tastes to Dennett at time t'* or *the way the juice tastes to Jones at time t*.

This 'conclusion' seems innocent, but right here we have already made the big mistake. The final step presumes that we can isolate the qualia from everything else that is going on—at least in principle or for the sake of argument. What counts as *the way the juice tastes to x* can be distinguished, one supposes, from what is a mere accompaniment,

contributory cause, or byproduct of this 'central' way. One dimly imagines taking such cases and stripping them down gradually to the essentials, leaving their common residuum, the way things look, sound, feel, taste, smell to various individuals at various times, independently of how they are subsequently disposed to behave or believe. The mistake is not in supposing that we can in practice ever or always perform this act of purification with certainty, but the more fundamental mistake of supposing that there is such a residual property to take seriously, however uncertain our actual attempts at isolation of instances might be.

The examples that seduce us are abundant in every modality. I cannot imagine, will never know, could never know, it seems, how Bach sounded to Glenn Gould. (I can barely recover in my memory the way Bach sounded to me when I was a child.) And I cannot know, it seems, what it is like to be a bat (Nagel 1974), or whether you see what I see, colourwise, when we look up at a clear 'blue' sky. The homely cases convince us of the reality of these special properties—those subjective tastes, looks, aromas, sounds—that we then apparently isolate for definition by this philosophical distillation.

The specialness of these properties is hard to pin down, but can be seen at work in *intuition pump 2: the wine-tasting machine*. Could Gallo Brothers replace their human wine-tasters with a machine? A computer-based 'expert system' for quality control and classification is probably within the bounds of existing technology. We now know enough about the relevant chemistry to make the transducers that would replace taste buds and olfactory organs (delicate colour vision would perhaps be more problematic), and we can imagine using the output of such transducers as the raw material—the 'sense data' in effect—for elaborate evaluations, descriptions, classifications. Pour the sample in the funnel and, in a few minutes or hours, the system would type out a chemical assay, along with commentary: "a flamboyant and velvety Pinot, though lacking in stamina"—or words to such effect. Such a machine might well perform better than human wine-tasters on all reasonable tests of accuracy and consistency the wine-makers could devise,³ but *surely* no matter how 'sensitive' and 'discriminating' such a system becomes, it will never have, and enjoy, what *we* do when we taste a wine: the qualia of conscious experience. Whatever informational, dispositional, functional properties its internal states have, none of them will be special in the way qualia are. If you share that intuition, you believe that there are qualia in the sense I am targeting for demolition.

What is special about qualia? Traditional analyses suggest some fascinating second-order properties of these properties. First, since one *cannot say* to another, no matter how eloquent one is and no matter how co-operative and imaginative one's audience is, exactly what way one is currently seeing, tasting, smelling, and so forth, qualia are *ineffable*—in fact the paradigm cases of ineffable items. According to tradition, at least part of the reason why qualia are ineffable is that they are

intrinsic properties—which seems to imply *inter alia* that they are somehow atomic and unanalysable. Since they are ‘simple’ or ‘homogeneous’ there is nothing to get hold of when trying to describe such a property to one unacquainted with the particular instance in question.

Moreover, verbal comparisons are not the only cross-checks ruled out. Any objective, physiological, or ‘merely behavioral’ test—such as those passed by the imaginary wine-tasting system—would of necessity miss the target (one can plausibly argue), so all interpersonal comparisons of these ways of appearing are (apparently) systematically impossible. In other words, qualia are essentially *private* properties. And, finally, since they *are* properties of *my experiences* (they are not chopped liver, and they are not properties of, say, my cerebral blood flow—or haven’t you been paying attention?), qualia are essentially directly accessible to the consciousness of their experiencer (whatever that means), or qualia are properties of one’s experience with which one is intimately or directly acquainted (whatever that means), or ‘immediate phenomenological qualities’ (Block 1978) (whatever that means). They are, after all, the very properties the appreciation of which permits us to identify our conscious states. So, to summarize the tradition, qualia are supposed to be properties of a subject’s mental states that are

- (1) ineffable
- (2) intrinsic
- (3) private
- (4) directly or immediately apprehensible in consciousness.

Thus are qualia introduced onto the philosophical stage. They have seemed to be very significant properties to some theorists because they have seemed to provide an insurmountable and unavoidable stumbling block to functionalism or, more broadly, to materialism or, more broadly still, to any purely ‘third-person’ objective viewpoint or approach to the world (Nagel 1986). Theorists of the contrary persuasion have patiently and ingeniously knocked down all the arguments, and said most of the right things, but they have made a tactical error, I am claiming, of saying in one way or another: “We theorists can handle *those qualia* you talk about just fine; we will show that you are just slightly in error about the nature of qualia.” What they ought to have said is: “What qualia?”

My challenge strikes some theorists as outrageous or misguided because they think they have a much blander and hence less vulnerable notion of qualia to begin with. They think I am setting up and knocking down a straw man, and ask, in effect: “Who said qualia are ineffable, intrinsic, private, directly apprehensible ways things seem to one?” Since my suggested fourfold essence of qualia may strike many readers as tendentious, it may be instructive to consider, briefly, an apparently milder alternative: qualia are simply ‘the qualitative or phenomenal

features of sense experience[s], in virtue of having which they resemble and differ from each other, qualitatively, in the ways they do' (Shoemaker 1982, 367). Surely I do not mean to deny *those* features.

I reply: it all depends on what 'qualitative or phenomenal' comes to. Shoemaker contrasts *qualitative* similarity and difference with 'intentional' similarity and difference—similarity and difference of the properties an experience represents or is of'. That is clear enough, but what then of 'phenomenal'? Among the non-intentional (and hence qualitative?) properties of my visual states are their physiological properties. Might these very properties be the qualia Shoemaker speaks of? It is supposed to be obvious, I take it, that these sorts of features are ruled out, because they are not 'accessible to introspection' (S. Shoemaker, personal communication). These are features of my visual *state*, perhaps, but not of my visual *experience*. They are not *phenomenal* properties.

But then another non-intentional similarity some of my visual states share is that they tend to make me think about going to bed. I think this feature of them *is* accessible to introspection—on any ordinary, pre-theoretical construal. Is that a phenomenal property or not? The term 'phenomenal' means nothing obvious and untendentious to me, and looks suspiciously like a gesture in the direction leading back to ineffable, private, directly apprehensible ways things seem to one.⁴

I suspect, in fact, that many are unwilling to take my radical challenge seriously, largely because they want so much for qualia to be acknowledged. Qualia seem to many people to be the last ditch defence of the inwardness and elusiveness of our minds, a bulwark against creeping mechanism. They are sure there must be *some* sound path from the homely cases to the redoubtable category of the philosophers, since otherwise their last bastion of specialness will be stormed by science.

This special status for these presumed properties has a long and eminent tradition. I believe it was Einstein who once advised us that science could not give us the *taste* of the soup. Could such a wise man have been wrong? Yes, if he is taken to have been trying to remind us of the qualia that hide forever from objective science in the subjective inner sancta of our minds. There are no such things. Another wise man said so—Wittgenstein (1958, especially pp. 91–100). Actually, what he said was:

The thing in the box has no place in the language-game at all; not even as a *something*; for the box might even be empty.—No, one can 'divide through' by the thing in the box; it cancels out, whatever it is (p. 100);

and then he went on to hedge his bets by saying "It is not a *something*, but not a *nothing* either! The conclusion was only that a nothing would serve just as well as a something about which nothing could be said" (p. 102). Both Einstein's and Wittgenstein's remarks are endlessly amenable to exegesis, but, rather than undertaking to referee this War of the

Titans, I choose to take what may well be a more radical stand than Wittgenstein's.⁵ Qualia are not even 'something about which nothing can be said'; 'qualia' is a philosophers' term which fosters nothing but confusion,⁶ and refers in the end to no properties or features at all.

The Traditional Paradox Regained

Qualia have not always been in good odour among philosophers. Although many have thought, along with Descartes and Locke, that it made sense to talk about private, ineffable properties of minds, others have argued that this is strictly nonsense—however naturally it trips off the tongue. It is worth recalling how qualia were presumably rehabilitated as properties to be taken seriously in the wake of Wittgensteinian and verificationist attacks on them as pseudo-hypotheses. The original version of *intuition pump 3: the inverted spectrum* (Locke 1959) is a speculation about two people: how do I know that you and I see the same subjective colour when we look at something? Since we both learned colour words by being shown public coloured objects, our verbal behaviour will match *even if we experience entirely different subjective colours*. The intuition that this hypothesis is systematically unconfirmable (and undisconfirmable, of course) has always been quite robust, but some people have always been tempted to think technology could (in principle) bridge the gap.

Suppose, in *intuition pump 4: the Brainstorm machine*, there were some neuroscientific apparatus that fits on your head and feeds your visual experience into my brain (as in the movie, *Brainstorm*, which is not to be confused with the book, *Brainstorms*). With eyes closed I accurately report everything you are looking at, except that I marvel at how the sky is yellow, the grass red, and so forth. Would this not confirm, empirically, that our qualia were different? But suppose the technician then pulls the plug on the connecting cable, inverts it 180 degrees, and reinserts it in the socket. Now I report the sky is blue, the grass green, and so forth. Which is the 'right' orientation of the plug? Designing and building such a device would require that its 'fidelity' be tuned or calibrated by the normalization of the two subjects' reports—so we would be right back at our evidential starting point. The moral of this intuition pump is that no intersubjective comparison of qualia is possible, even with perfect technology.

So matters stood until someone dreamt up the presumably improved version of the thought experiment: the *intrapersonal inverted spectrum*. The idea seems to have occurred to several people independently (Gert 1965; Putnam 1965; Taylor 1966; Shoemaker 1969, 1975; Lycan 1973). Probably Block and Fodor (1972) have it in mind when they say "It seems to us that the standard verificationist counterarguments against the view that the 'inverted spectrum' hypothesis is conceptually incoherent are not persuasive" (p. 172). In this version, *intuition pump 5: the*

neurosurgical prank, the experiences to be compared are all in one mind. You wake up one morning to find that the grass has turned red, the sky yellow, and so forth. No one else notices any colour anomalies in the world, so the problem must be in you. You are entitled, it seems, to conclude that you have undergone visual colour qualia inversion (and we later discover, if you like, just how the evil neurophysiologists tampered with your neurons to accomplish this).

Here it seems at first—and indeed for quite a while—that qualia are acceptable properties after all, because propositions about them can be justifiably asserted, empirically verified, and even explained. After all, in the imagined case, we can tell a tale in which we confirm a detailed neurophysiological account of the precise etiology of the dramatic change you undergo. It is tempting to suppose, then, that neurophysiological evidence, incorporated into a robust and ramifying theory, would have all the resolving power we could ever need for determining whether or not someone's qualia have actually shifted.

But this is a mistake. It will take some patient exploration to reveal the mistake in depth, but the conclusion can be reached—if not secured—quickly with the help of *intuition pump 6: alternative neurosurgery*. There are (at least) two different ways the evil neurosurgeon might create the inversion effect described in *intuition pump 5*:

1. Invert one of the 'early' qualia-producing channels, e.g. in the optic nerve, so that all relevant neural events 'downstream' are the 'opposite' of their original and normal values. *Ex hypothesi* this inverts your qualia.
2. Leave all those early pathways intact and simply invert certain memory-access links—whatever it is that accomplishes your tacit (and even unconscious) comparison of today's hues, with those of yore. *Ex hypothesi* this does *not* invert your qualia at all, but just your memory-anchored dispositions to react to them.

On waking up and finding your visual world highly anomalous, you should exclaim 'Egad! *Something* has happened! Either my qualia have been inverted or my memory-linked qualia reactions have been inverted. I wonder which!'

The intrapersonal, inverted spectrum thought experiment was widely supposed to be an improvement, since it moved the needed comparison into one subject's head. But now we can see that this is an illusion, since the link to earlier experiences, the link via memory, is analogous to the imaginary cable that might link two subjects in the original version.

This point is routinely—one might say traditionally—missed by the constructors of 'intrasubjective, inverted spectrum' thought experiments, who suppose that the subject's *noticing the difference*—surely a vivid experience is discovery by the subject—would have to be an instance of (directly? incorrigibly?) recognizing the difference *as a shift*

in qualia. But as my example shows, we could achieve the same startling effect in a subject without tampering with his presumed qualia at all. Since *ex hypothesi* the two different surgical invasions can produce exactly the same introspective effects, while only one operation inverts the qualia, nothing in the subject's experience can favour one of the hypotheses over the other. So unless he seeks outside help, the state of his own qualia must be as unknowable to him as the state of anyone else's qualia: hardly the privileged access or immediate acquaintance or direct apprehension the friends of qualia had supposed 'phenomenal features' to enjoy!

The outcome of this series of thought experiments is an intensification of the 'verificationist' argument against qualia. If there are qualia, they are even less accessible to our ken than we had thought. Not only are the classical intersubjective comparisons impossible (as the *Brainstorm* machine shows), but we cannot tell in our own cases whether our qualia have been inverted—at least not by introspection. It is surely tempting at this point—especially to non-philosophers—to decide that this paradoxical result must be an artefact of some philosophical misanalysis or other, the sort of thing that might well happen if you took a perfectly good pre-theoretical notion—our everyday notion of qualia—and illicitly stretched it beyond the breaking point. The philosophers have made a mess; let them clean it up; meanwhile we others can get back to work, relying as always on our sober and unmetaphysical acquaintance with qualia.

Overcoming this ubiquitous temptation is the task of the next section, which will seek to establish the unsalvageable incoherence of the hunches that lead to the paradox by looking more closely at their sources and their motivation.

Making Mistakes about Qualia

The idea that people might be mistaken about their own qualia is at the heart of the ongoing confusion and must be explored in more detail, and with somewhat more realistic examples if we are to see the delicate role it plays.

Intuition pump 7: Chase and Sanborn. Once upon a time there were two coffee-tasters, Mr Chase and Mr Sanborn, who worked for Maxwell House.⁷ Along with half a dozen other coffee-tasters, their job was to ensure that the taste of Maxwell House coffee stayed constant, year after year. One day, about six years after Chase had come to work for Maxwell House, he confessed to Sanborn:

I hate to admit it, but I'm not enjoying this work anymore. When I came to Maxwell House six years ago, I thought Maxwell House coffee was the best-tasting coffee in the world. I was proud to have a share in the responsibility for preserving that flavour over the years. And we've done our job well; the coffee tastes just the same today as it

tasted when I arrived. But, you know, I no longer like it! My tastes have changed. I've become a more sophisticated coffee drinker. I no longer like *that taste* at all.

Sanborn greeted this revelation with considerable interest. 'It's funny you should mention it,' he replied, 'for something rather similar has happened to me.' He went on:

When I arrived here, shortly before you did, I, like you, thought Maxwell House coffee was tops in flavour. And now I, like you, really don't care for the coffee we're making. But *my* tastes haven't changed; my . . . *tasters* have changed. That is, I think something has gone wrong with my taste buds or some other part of my taste-analyzing perceptual machinery. Maxwell House coffee doesn't taste to me the way it used to taste; if only it did, I'd still love it, for I still think *that taste* is the best taste in coffee. Now I'm not saying we haven't done our job well. You other tasters all agree that the taste is the same, and I must admit that on a day-to-day basis I can detect no change either. So it must be my problem alone. I guess I'm no longer cut out for this work.

Chase and Sanborn are alike in one way at least: they both used to like Maxwell House coffee, and now neither likes it. But they claim to be different in another way. Maxwell House tastes to Chase just the way it always did, but not so for Sanborn. But can we take their protestations at face value? Must we? Might one or both of them simply be wrong? Might their predicaments be importantly the same and their apparent disagreement more a difference in manner of expression than in experiential or psychological state? Since both of them make claims that depend on the reliability of their memories, is there any way to check on this reliability?

My reason for introducing two characters in the example is not to set up an interpersonal comparison between how the coffee tastes to Chase and how it tastes to Sanborn, but just to exhibit, side by side, two poles between which cases of intrapersonal experiential shift can wander. Such cases of intrapersonal experiential shift, and the possibility of adaptation to them, or interference with memory in them, have often been discussed in the literature on qualia, but without sufficient attention to the details, in my opinion. Let us look at Chase first. If we fall in for the nonce with the received manner of speaking, it appears at first that there are the following possibilities:

- (a) Chase's coffee-taste qualia have stayed constant, while his reactive attitudes to those qualia, devolving on his canons of aesthetic judgment, etc., have shifted—which is what he seems, in his informal, casual way, to be asserting.
- (b) Chase is simply wrong about the constancy of his qualia; they have shifted gradually and imperceptibly over the years, while his standards of taste have not budged—in spite of his delusions about having become more sophisticated. He is in the state Sanborn claims to be in, but just lacks Sanborn's self-knowledge.

- (c) Chase is in some predicament intermediate between (a) and (b); his qualia have shifted some *and* his standards of judgment have also slipped.

Sanborn's case seems amenable to three counterpart versions:

- (a) Sanborn is right; his qualia have shifted, due to some sort of derangement in his perceptual machinery, but his standards have indeed remained constant.
- (b) Sanborn's standards have shifted unbeknownst to him. He is thus misremembering his past experiences, in what we might call a nostalgia effect. Think of the familiar experience of returning to some object from your childhood (a classroom desk, a tree-house) and finding it much smaller than you remember it to have been. Presumably as you grew larger your internal standard for what was large grew with you somehow, but your memories (which are stored as fractions or multiples of that standard) did not compensate, and hence, when you consult your memory, it returns a distorted judgment. Sanborn's nostalgia-tinged memory of good old Maxwell House is similarly distorted. (There are obviously many different ways this impressionistic sketch of a memory mechanism could be implemented, and there is considerable experimental work in cognitive psychology that suggests how different hypotheses about such mechanisms could be tested.)
- (c) As before, Sanborn's state is some combination of (a) and (b).

I think that everyone writing about qualia today would agree that there are all these possibilities for Chase and Sanborn. I know of no one these days who is tempted to defend the high line on infallibility or incorrigibility that would declare that alternative (a) is—and must be—the truth in each case, since people just cannot be wrong about such private, subjective matters.⁸

Since quandaries are about to arise, however, it might be wise to review in outline why the attractiveness of the infallibilist position is only superficial, so it will not recover its erstwhile allure when the going gets tough. First, in the wake of Wittgenstein (1958) and Malcolm (1956, 1959), we have seen that one way to buy such infallibility is to acquiesce in the complete evaporation of content (Dennett 1976). "Imagine someone saying: 'But I know how tall I am!' and laying his hand on top of his head to prove it" (Wittgenstein 1958, 96). By diminishing one's claim until there is nothing left to be right or wrong about, one can achieve a certain empty invincibility, but that will not do in this case. One of the things we want Chase to be right about (if he is right) is that he is not in Sanborn's predicament, so if the claim is to be reviewed as infallible it can hardly be because it declines to assert anything.

There is a strong temptation, I have found, to respond to my claims in this chapter more or less as follows: "But after all is said and done,

there is still something I know in a special way: I know *how it is with me right now.*" But if absolutely nothing follows from this presumed knowledge—nothing, for instance, that would shed any light on the different psychological claims that might be true of Chase or Sanborn—what is the point of asserting that one has it? Perhaps people just want to reaffirm their sense of proprietorship over their own conscious states.

The infallibilist line on qualia treats them as properties of one's experience one cannot in principle misdiscover, and this is a mysterious doctrine (at least as mysterious as papal infallibility) unless we shift the emphasis a little and treat qualia as *logical constructs* out of subjects' qualia judgements: a subject's experience has the quale *F* if and only if the subject judges his experience to have quale *F*. We can then treat such judgments as constitutive acts, in effect, bringing the quale into existence by the same sort of license as novelists have to determine the hair colour of their characters by fiat. We do not ask how Dostoevski knows that Raskolnikov's hair is light brown.

There is a limited use for such interpretations of subjects' protocols, I have argued (Dennett 1978a, 1979, especially pp. 109–110, 1982), but they will not help the defenders of qualia here. Logical constructs out of judgments must be viewed as akin to theorists' fictions, and the friends of qualia want the existence of a particular quale in any particular case to be an empirical fact in good standing, not a theorist's useful interpretive fiction, else it will not loom as a challenge to functionalism or materialism or third-person objective science.

It seems easy enough, then, to dream up empirical tests that would tend to confirm Chase and Sanborn's different tales, but if passing such tests could support their authority (that is to say, their reliability), failing the tests would have to undermine it. The price you pay for the possibility of empirically confirming your assertions is the outside chance of being discredited. The friends of qualia are prepared, today, to pay that price, but perhaps only because they have not reckoned how the bargain they have struck will subvert the concept they want to defend.

Consider how we could shed light on the question of where the truth lies in the particular cases of Chase and Sanborn, even if we might not be able to settle the matter definitively. It is obvious that there might be telling objective support for one extreme version or another of their stories. Thus if Chase is unable to re-identify coffees, teas, and wines in blind tastings in which only minutes intervene between first and second sips, his claim to *know* that Maxwell House tastes just the same to him now as it did six years ago will be seriously undercut. Alternatively, if he does excellently in blind tastings, and exhibits considerable knowledge about the canons of coffee style (if such there be), his claim to have become a more sophisticated taster will be supported. Exploitation of the standard principles of inductive testing—basically Mill's method of differences—can go a long way toward indicating what sort

of change has occurred in Chase or Sanborn—a change near the brute perceptual processing end of the spectrum or a change near the ultimate reactive judgment end of the spectrum. And as Shoemaker (1982) and others have noted, physiological measures, suitably interpreted in some larger theoretical framework, could also weight the scales in favour of one extreme or the other. For instance, the well-studied phenomenon of induced illusory boundaries (see figure 17.1) has often been claimed to be a particularly ‘cognitive’ illusion, dependent on ‘top-down’ processes and hence, presumably, near the reactive judgment end of the spectrum, but recent experimental work (Von der Heydt et al. 1984) has revealed that ‘edge detector’ neurons *relatively* low in the visual pathways—in area 18 of the visual cortex—are as responsive to illusory edges as to real light-dark boundaries on the retina, suggesting (but not quite proving, since these might somehow still be ‘descending effects’) that illusory contours are not imposed from on high, but generated quite early in visual processing. One can imagine discovering a similarly ‘early’ anomaly in the pathways leading from taste buds to judgment in Sanborn, for instance, tending to confirm his claim that he has suffered some change in his basic perceptual—as opposed to judgmental—machinery.

But let us not overestimate the resolving power of such empirical testing. The space in each case between the two poles represented by possibility (a) and possibility (b) would be occupied by phenomena that were the product, somehow, of two factors in varying proportion: roughly, dispositions to generate or produce qualia and dispositions to react to the qualia once they are produced. (That is how our intuitive picture of qualia would envisage it.) Qualia are supposed to affect our action or behaviour only via the intermediary of our judgments about




Figure 17.1 Induced illusory contours

them, so any behavioural test, such as a discrimination or memory test, since it takes acts based on judgments as its primary data, can give us direct evidence only about the *resultant* of our two factors. In extreme cases we can have indirect evidence to suggest that one factor has varied a great deal, the other factor hardly at all, and we can test the hypothesis further by checking the relative sensitivity of the subject to variations in the conditions that presumably alter the two component factors. But such indirect testing cannot be expected to resolve the issue when the effects are relatively small—when, for instance, our rival hypotheses are Chase's preferred hypothesis (a) and the minor variant to the effect that his qualia have shifted *a little* and his standards *less than he thinks*. This will be true even when we include in our data any unintended or unconscious behavioural effects, for their import will be ambiguous. (Would a longer response latency in Chase today be indicative of a process of "attempted qualia renormalization" or "extended aesthetic evaluation?")

The limited evidential power of neurophysiology comes out particularly clearly if we imagine a case of adaptation. Suppose, in *intuition pump 8: the gradual post-operative recovery*, that we have somehow "surgically inverted" Chase's taste bud connections in the standard imaginary way: post-operatively, sugar tastes salty, salt tastes sour, etc. But suppose further—and this is as realistic a supposition as its denial—that Chase has subsequently compensated—as revealed by his behaviour. He now *says* that the sugary substance we place on his tongue is sweet, and no longer favours gravy on his ice-cream. Let us suppose the compensation is so thorough that on all behavioural and verbal tests his performance is indistinguishable from that of normal subjects—and from his own pre-surgical performance.

If all the internal compensatory adjustment has been accomplished early in the process—intuitively, pre-qualia—then his qualia today are restored to just as they were (relative to external sources of stimulation) before the surgery. If on the other hand some or all of the internal compensatory adjustment is post-qualia, then his qualia have not been renormalized *even if he thinks they have*. But the physiological facts will not in themselves shed any light on where in the stream of physiological process twixt tasting and telling to draw the line at which the putative qualia appear as properties of that phase of the process. The qualia are the 'immediate or phenomenal' properties, of course, but this description will not serve to locate the right phase in the physiological stream, for, echoing intuition pump 6, there will always be at least two possible ways of interpreting the neurophysiological theory, however it comes out. Suppose our physiological theory tells us (in as much detail as you like) that the compensatory effect in him has been achieved by an *adjustment in the memory-accessing process* that is required for our victim to compare today's hues to those of yore. There are *still* two stories that might be told:

- (I) Chase's current qualia are still abnormal, but thanks to the revision in his memory-accessing process, he has in effect adjusted his memories of how things used to taste, so he no longer notices any anomaly.
- (II) The memory-comparison step occurs just prior to the qualia phase in taste perception; thanks to the revision, it now *yields* the same old qualia for the same stimulation.

In (I) the qualia contribute to the input; in effect, to the memory comparator. In (II) they are part of the output of the memory comparator. These seem to be two substantially different hypotheses, but the physiological evidence, no matter how well developed, will not tell us on which side of memory to put the qualia. Chase's introspective evidence will not settle the issue between (I) and (II) either, since *ex hypothesi* those stories are not reliably distinguishable by him. Remember that it was to confirm or disconfirm Chase's opinion that we turned to the neurophysiological evidence in the first place. We can hardly use his opinion in the end to settle the matter between our rival neurophysiological theories. Chase may think that he thinks his experiences are the same as before *because* they really are (and he remembers accurately how it used to be), but he must admit that he has no introspective resources for distinguishing that possibility from alternative (I), on which he thinks things are as they used to be *because* his memory of how they used to be has been distorted by his new compensatory habits.

Faced with their subject's systematic neutrality, the physiologists may have their own reasons for preferring (I) to (II), or vice versa, for they may have *appropriated* the term 'qualia' to their own theoretical ends, to denote some family of detectable properties that strike them as playing an important role in their neurophysiological theory of perceptual recognition and memory. Chase or Sanborn might complain—in the company of more than a few philosophical spokesmen—that these properties the neurophysiologists choose to call 'qualia' are not the qualia they are speaking of. The scientists' retort is: "If we cannot distinguish (I) from (II), we certainly cannot support either of your claims. If you want our support, you must relinquish your concept of qualia."

What is striking about this is not just that the empirical methods would fall short of distinguishing what seem to be such different claims about qualia, but that they would fall short *in spite of being better evidence than the subject's own introspective convictions*. For the subject's own judgments, like the behaviours or actions that express them, are the resultant of our two postulated factors, and cannot discern the component proportions any better than external behavioural tests can. Indeed, a subject's 'introspective' convictions will generally be *worse* evidence than what outside observers can gather. For if our subject is—as most are—a 'naive subject', unacquainted with statistical data about his own case

or similar cases, his immediate, frank judgments are, evidentially, like any naive observer's perceptual judgments about factors in the outside world. Chase's intuitive judgments about his qualia constancy are no better off, epistemically, than his intuitive judgments about, say, lighting intensity constancy or room temperature constancy—or his own body temperature constancy. Moving to a condition inside his body does not change the intimacy of the epistemic relation in any special way. Is Chase running a fever or just feeling feverish? Unless he has taken steps to calibrate and cross-check his own performance, his opinion that his fever-perception apparatus is undisturbed is no better than a hunch. Similarly, Chase may have a strongly held opinion about the degree to which his taste-perceiving apparatus has maintained its integrity, and the degree to which his judgment has evolved through sophistication, but, pending the results of the sort of laborious third-person testing just imagined, he would be a fool to claim to know—especially to know directly or immediately—that his was a pure case (a), closer to (a) than to (b), or a case near (b).

Chase is on quite firm ground, epistemically, when he reports that *the relation* between his coffee-sipping activity and his judging activity has changed. Recall that this is the factor that Chase and Sanborn have in common: they used to like Maxwell House; now they do not. But unless he carries out on himself the sorts of tests others might carry out on him, his convictions about what has stayed constant (or nearly so) and what has shifted *must be sheer guessing*.

But then qualia—supposing for the time being that we know what we are talking about—must lose one of their 'essential' second-order properties: far from being directly or immediately apprehensible properties of our experience, they are properties whose changes or constancies are either entirely beyond our ken, or inferrable (at best) from 'third-person' examinations of our behavioural and physiological reaction patterns (if Chase and Sanborn acquiesce in the neurophysiologists' sense of the term). On this view, Chase and Sanborn should be viewed not as introspectors capable of a privileged view of these properties, but as autopsychologists, theorists whose convictions about the properties of their own nervous systems are based not only on their 'immediate' or current experiential convictions, but also on their appreciation of the import of events they remember from the recent past.

There are, as we shall see, good reasons for neurophysiologists and other 'objective, third-person' theorists to single out such a class of properties to study. But they are not qualia, for the simple reason that one's epistemic relation to them is *exactly* the same as one's epistemic relation to such external, but readily—if fallibly—detectable, properties as room temperature or weight. The idea that one should consult an outside expert, and perform elaborate behavioural tests on oneself to confirm what qualia one had, surely takes us too far away from our

original idea of qualia as properties with which we have a particularly intimate acquaintance.

So perhaps we have taken a wrong turning. The doctrine that led to this embarrassing result was the doctrine that sharply distinguished qualia from their (normal) effects on reactions. Consider Chase again. He claims that coffee tastes "just the same" as it always did, but he admits—nay insists—that his reaction to "that taste" is not what it used to be. That is, he pretends to be able to divorce his apprehension (or recollection) of the quale—the taste, in ordinary parlance—from his different reactions to the taste. But this apprehension or recollection is itself a reaction to the presumed quale, so some sleight of hand is being perpetrated—innocently no doubt—by Chase. So suppose instead that Chase had insisted that precisely *because* his reaction was now different, the taste had changed for him. (When he told his wife his original tale, she said "Don't be silly! Once you add the dislike you change the experience"—and the more he thought about it, the more he decided she was right.)

Intuition pump 9: the experienced beer drinker. It is familiarly said that beer, for example, is an acquired taste; one gradually trains oneself—or just comes—to enjoy that flavour. What flavour? The flavour of the first sip? No one could like *that* flavour, an experienced beer drinker might retort:

Beer tastes different to the experienced beer drinker. If beer went on tasting to me the way the first sip tasted, I would never have gone on drinking beer! Or to put the same point the other way around, if my first sip of beer had tasted to me the way my most recent sip just tasted, I would never have had to acquire the taste in the first place! I would have loved the first sip as much as the one I just enjoyed.

If we let this speech pass, we must admit that beer is *not* an acquired taste. No one comes to enjoy *the way the first sip tasted*. Instead, prolonged beer drinking leads people to experience a taste they enjoy, but precisely their enjoying the taste guarantees that it is not the taste they first experience.⁹

But this conclusion, if it is accepted, wreaks havoc of a different sort with the traditional philosophical view of qualia. For if it is admitted that one's attitudes towards, or reactions to, experiences are in any way and in any degree constitutive of their experiential qualities, so that a change in reactivity *amounts to* or *guarantees* a change in the property, then those properties, those 'qualitative or phenomenal features', cease to be 'intrinsic' properties and in fact become paradigmatically extrinsic, relational properties.

Properties that 'seem intrinsic' at first often turn out on more careful analysis to be relational. Bennett (1965) is the author of *intuition pump 10: the world-wide eugenics experiment*. He draws our attention to phenolthio-urea, a substance which tastes very bitter to three-fourths of humanity, and as tasteless as water to the rest. Is it bitter? Since the

reactivity to phenol-thio-urea is genetically transmitted, we could make it paradigmatically bitter by performing a large-scale breeding experiment: prevent the people to whom it is tasteless from breeding, and in a few generations phenol would be as bitter as anything to be found in the world. But we could also (in principle) perform the contrary feat of mass 'eugenics' and thereby make phenol paradigmatically tasteless—as tasteless as water—without ever touching phenol. Clearly, public bitterness or tastelessness is not an intrinsic property of phenol-thio-urea but a relational property, since the property is changed by a change in the reference class of normal detectors.

The public versions of perceptual 'qualia' all *seem* intrinsic, in spite of their relationality. They are not alone. Think of the 'felt value' of a dollar (or whatever your native currency is). "How much is that in *real* money?" the American tourist is reputed to have asked, hoping to translate a foreign price onto the scale of 'intrinsic value' he keeps in his head. As Elster (1985) claims, "there is a tendency to overlook the implicitly relational character of certain monadic predicates." Walzer (1985) points out that "a ten-dollar bill might seem to have a life of its own as a thing of value, but, as Elster suggests, its value implicitly depends on 'other people who are prepared to accept money as payment for goods'." But even as one concedes this, there is still a tendency to reserve something subjective, felt value, as an 'intrinsic' property of that ten-dollar bill. But as we now see, such intrinsic properties cannot be properties to which a subject's access is in any way privileged.

Which way should Chase go? Should he take his wife's advice and declare that since he cannot stand the coffee anymore, it no longer tastes the same to him (it used to taste good and now it tastes bad)? Or should he say that really, in a certain sense, it does taste the way it always did, or at least it sort of does—when you subtract the fact that it tastes so bad now, of course?

We have now reached the heart of my case. The fact is that we have to ask Chase which way he wants to go, and there really are two drastically different alternatives available to him *if we force the issue*. Which way would *you* go? Which concept of qualia did you "always have in the back of your mind," guiding your imagination as you thought about theories? If you acknowledge that the answer is not obvious, and especially if you complain that this forced choice drives apart two aspects that you had supposed united in your pre-theoretic concept, you support my contention that there is no secure foundation in ordinary 'folk psychology' for a concept of qualia. We *normally* think in a confused and potentially incoherent way when we think about the ways things seem to us.

When Chase thinks of 'that taste' he thinks equivocally or vaguely. He harkens back in memory to earlier experiences, but need not try—or be able—to settle whether he is including any or all of his reactions or excluding them from what he intends by 'that taste'. His state then

and his state now are different—that he can avow with confidence—but he has no ‘immediate’ resources for making a finer distinction, nor any need to do so.¹⁰

This suggests that qualia are no more essential to the professional vocabulary of the phenomenologist (or professional coffee-taster) than to the vocabulary of the physiologist (Dennett 1978b). To see this, consider again the example of my dislike of cauliflower. Imagine now, in *intuition pump 11: the cauliflower cure*, that someone offers me a pill to cure my loathing for cauliflower. He promises that after I swallow this pill cauliflower will taste exactly the same to me as it always has, but I will like that taste. “Hang on,” I might reply, “I think you may have just contradicted yourself.” But in any event I take the pill and it works. I become an instant cauliflower-appreciator, but if I am asked which of the two possible effects (Chase-type or Sanborn-type) the pill has had on me, I will be puzzled, and will find nothing in *my experience* to shed light on the question. Of course I recognize that the taste is (sort of) the same—the pill has not made the cauliflower taste like chocolate cake, after all—but at the same time my experience is so different now that I resist saying that cauliflower tastes the way it used to taste. There is in any event no reason to be cowed into supposing that my cauliflower experiences have some intrinsic properties behind, or in addition to, their various dispositional, reaction-provoking properties.

“But in principle there has to be a right answer to the question of how it is, intrinsically, with you now, even if you are unable to say with any confidence” Why? Would one say the same about all other properties of experience? Consider *intuition pump 12: visual field inversion created by wearing inverting spectacles*, a phenomenon which has been empirically studied for years. (G. M. Stratton published the pioneering work in 1896, and J. J. Gibson and I. Kohler were among the principal investigators; for an introductory account, see Gregory 1977.) After wearing inverting spectacles for several days subjects make an astonishingly successful adaptation. Suppose we pressed on them this question: “Does your adaptation consist in your re-inverting your visual field or in your turning the rest of your mind upside-down in a host of compensations?” If they demur, may we insist that there has to be a right answer, even if they cannot say with any confidence which it is? Such an insistence would lead directly to a new version of the old inverted spectrum thought experiment: “How do I know whether some people see things upside-down (but are perfectly used to it), while others see things right-side-up?”

Only a very naive view of visual perception could sustain the idea that one’s visual field has a property of right-side-upness or upside-downness *independent of one’s dispositions to react to it*—‘intrinsic right-side-upness’ we could call it (see my discussion of the properties of the ‘images’ processed by the robot, SHAKEY, in Dennett 1982). So not all properties of conscious experience invite or require treatment as ‘intrinsic

sic' properties. Is there something distinguishing about a certain subclass of properties (the 'qualitative or phenomenal' subclass, presumably) that forces us to treat them—unlike subjective right-side-upness—as intrinsic properties? If not, such properties have no role to play, in either physiological theories of experience, or in introspective theories.

Some may be inclined to argue this way: I can definitely imagine the experience of 'spectrum inversion' from the inside; after all, I have actually experienced temporary effects of the same type, such as the 'taste-displacement' effect of the maple syrup on the orange juice. What is imaginable, or actual, is possible. Therefore spectrum inversion or displacement (in all sensory modalities) is possible. But such phenomena just *are* the inversion or displacement of qualia, or intrinsic subjective properties. Therefore there must be qualia: intrinsic subjective properties.

This is fallacious. What one imagines and what one says one imagines may be two different things. To imagine visual field inversion, of the sort Stratton and Kohler's subjects experienced, is not necessarily to imagine the absolute inversion of a visual field (even if that is what it 'feels like' to the subjects). Less obviously, you imagining—as vividly as you like—a case of subjective colour-perception displacement is not necessarily you imagining what that phenomenon is typically called by philosophers: an inverted or displaced spectrum of *qualia*. In so far as that term carries the problematic implications scouted here, there is no support for its use arising simply from the vividness or naturalness of the imagined possibility.

If there are no such properties as qualia, does that mean that 'spectrum inversion' is impossible? Yes and no. Spectrum inversion as classically debated is impossible, but something like it is perfectly possible—something that is as like 'qualia inversion' as visual field inversion is like the impossible *absolute* visual image inversion we just dismissed.

Some Puzzling Real Cases

It is not enough to withhold our theoretical allegiances until the sunny day when the philosophers complete the tricky task of purifying the everyday concept of qualia. Unless we take active steps to shed this source concept, and replace it with better ideas, it will continue to cripple our imaginations and systematically distort our attempts to understand the phenomena already encountered.

What we find, if we look at the actual phenomena of anomalies of colour perception, for instance, amply bears out our suspicions about the inadequacy of the traditional notion of qualia. Several varieties of *cerebral achromatopsia* (brain-based impairment of colour vision) have been reported, and while there remains much that is unsettled about their analysis, there is little doubt that the philosophical thought ex-

periments have underestimated or overlooked the possibilities for counterintuitive collections of symptoms, as a few very brief excerpts from case histories will reveal.

Objects to the right of the vertical meridian appeared to be of normal hue, while to the left they were perceived only in shades of gray, though without distortions of form. . . . He was unable to recognize or name any color in any portion of the left field of either eye, including bright reds, blues, greens and yellows. As soon as any portion of the colored object crossed the vertical meridian, he was able to instantly recognize and accurately name its color (Damasio et al. 1980).

This patient would seem at first to be unproblematically describable as suffering a shift or loss of colour qualia in the left hemifield, but there is a problem of interpretation here, brought about by another case:

The patient failed in all tasks in which he was required to match the seen color with its spoken name. Thus, the patient failed to give the names of colors and failed to choose a color in response to its name. By contrast, he succeeded on all tasks where the matching was either purely verbal or purely nonverbal. Thus, he could give verbally the names of colors corresponding to named objects and vice versa. He could match seen colors to each other and to pictures of objects and could sort colors without error (Geschwind and Fusillo 1966).

This second patient was quite unaware of any deficit. He "never replied with a simple 'I don't know' to the demand for naming a colour" (Geschwind and Fusillo 1966, 140). There is a striking contrast between these two patients: both have impaired ability to name the colours of things in at least part of their visual field, but, whereas the former is acutely aware of his deficit, the latter is not. Does this difference make all the difference about qualia? If so, what on earth should we say about this third patient?

His other main complaint was that 'everything looked black or grey' and this caused him some difficulty in everyday life. . . . He had considerable difficulty recognizing and naming colours. He would, for example, usually describe bright red objects as either red or black, bright green objects as either green, blue or black, and bright blue objects as black. The difficulty appeared to be perceptual and he would make remarks suggesting this; for example when shown a bright red object he said 'a dirty smudgy red, not as red as you would normally see red.' Colours of lesser saturation or brightness were described in such terms as 'grey' 'off-white' or 'black,' but if told to guess at the colour, he would be correct on about 50 per cent of occasions, being notably less successful with blues and greens than reds (Meadows 1974).

This man's awareness of his deficit is problematic to say the least. It contrasts rather sharply with yet another case:

One morning in November 1977, upon awakening, she noted that although she was able to see details of objects and people, colors appeared 'drained out' and 'not true.' She had no other complaint . . . her vision was good, 20/20 in each eye. . . . The difficulty in color

perception persisted, and she had to seek the advice of her husband to choose what to wear. Eight weeks later she noted that she could no longer recognize the faces of her husband and daughter . . . [So in addition to achromatopsia, the patient has prosopagnosia, but her linguistic and cognitive performances were otherwise unaffected. The patient was able to tell her story cogently and to have remarkable insight about her defects (Damasio et al. 1980).

As Meadows notes, "Some patients thus complain that their vision for colours is defective while others have no spontaneous complaint but show striking abnormalities on testing."

What should one say in these cases? When no complaint is volunteered but the patient shows an impairment in colour vision, is this a sign that his qualia are unaffected? ("His capacities to discriminate are terribly impaired, but, luckily for him, his inner life is untouched by this merely public loss.") We could line up the qualia this way, but equally we could claim that the patient has simply not noticed the perhaps gradual draining away or inversion or merging of his qualia revealed by his poor performance. ("So slowly did his inner life lose its complexity and variety that he never noticed how impoverished it had become.") What if our last patient described her complaint just as she did above, but performed normally on testing? One hypothesis would be that her qualia had indeed, as she suggested, become washed out. Another would be that in the light of her sterling performance on the colour discrimination tests, her qualia were fine; she was suffering from some hysterical or depressive anomaly, a sort of colour-vision hypochondria that makes her complain about a loss of colour perception. Or perhaps one could claim that her qualia were untouched; her disorder was purely verbal: an anomalous understanding of the words she uses to describe her experience. (Other startlingly specific, colour-word disorders have been reported in the literature.)

The traditional concept leads us to overlook genuine possibilities. Once we have learned of the curious deficit reported by Geschwind and Fusillo (1966), for instance, we realize that our first patient was never tested to see if he could still sort colours seen on the left or pass other non-naming, non-verbal, colour-blindness tests. Those tests are by no means superfluous. Perhaps he would have passed them; perhaps, *in spite of what he says*, his qualia are as intact for the left field as for the right—if we take the capacity to pass such tests as 'critical'. Perhaps his problem is 'purely verbal'. If your reaction to this hypothesis is that this is impossible, that must mean you are making his verbal, reporting behaviour sovereign in settling the issue—but then you must rule out a priori the possibility of the condition I described as colour-vision hypochondria.

There is no prospect of *finding* the answers to these brain-teasers in our everyday usage or the intuitions it arouses, but it is of course open

to the philosopher to *create* an edifice of theory defending a particular set of interlocking proposals. The problem is that although normally a certain family of stimulus and bodily conditions yields a certain family of effects, any particular effect can be disconnected, and our intuitions do not tell us which effects are 'essential' to quale identity or qualia constancy (cf. Dennett 1978a, chap. 11.). It seems fairly obvious to me that none of the real problems of interpretation that face us in these curious cases are advanced by any analysis of how the concept of *qualia* is to be applied—unless we wish to propose a novel, technical sense for which the traditional term might be appropriated. But that would be at least a tactical error: the intuitions that surround and *purport* to anchor the current understanding of the term are revealed to be in utter disarray when confronted with these cases.

My informal sampling shows that some philosophers have strong opinions about each case and how it should be described in terms of qualia, but they find they are in strident (and ultimately comic) disagreement with other philosophers about how these 'obvious' descriptions should go. Other philosophers discover that they really do not know what to say—not because there are not enough facts presented in the descriptions of the cases, but because it begins to dawn on them that they have not really known what they were talking about over the years.

Filling the Vacuum

If qualia are such a bad idea, why have they seemed to be such a good idea? Why does it seem as if there are these intrinsic, ineffable, private, 'qualitative' properties in our experience? A review of the presumptive second-order properties of the properties of our conscious experiences will permit us to diagnose their attractiveness and find suitable substitutes (for a similar exercise, see Kitcher 1979).

Consider 'intrinsic' first. It is far from clear what an intrinsic property would be. Although the term has had a certain vogue in philosophy, and often seems to secure an important contrast, there has never been an accepted definition of the second-order property of intrinsicity. If even such a brilliant theory-monger as David Lewis can try and fail, by his own admission, to define the extrinsic/intrinsic distinction coherently, we can begin to wonder if the concept deserves our further attention after all. In fact Lewis (1983) begins his survey of versions of the distinction by listing as one option: "We could Quine the lot, give over the entire family as unintelligible and dispensable," but he dismisses the suggestion immediately: "That would be absurd" (p. 197). In the end, however, his effort to salvage the accounts of Chisholm (1976) and Kim (1982) are stymied, and he conjectures that "if we still want to break in we had best try another window" (p. 200).

Even if we are as loath as Lewis is to abandon the distinction, should we not be suspicious of the following curious fact? If challenged to explain the idea of an intrinsic property to a neophyte, many people would hit on the following sort of example: consider Tom's ball; it has many properties, such as its being made of rubber from India, its belonging to Tom, its having spent the last week in the closet, and its redness. All but the last of these are clearly *relational* or *extrinsic* properties of the ball. Its redness, however, is an intrinsic property. Except that this is not so. Ever since Boyle and Locke we have known better. Redness—public redness—is a quintessentially relational property, as many thought experiments about 'secondary qualities' show. [One of the first was Berkeley's [1713] pail of lukewarm water, and one of the best is Bennett's (1965) phenol-thio-urea.] The seductive step, on learning that public redness (like public bitterness, etc.) is a relational property after all, is to cling to intrinsicity ("*something* has to be intrinsic") and move it into the subject's head. It is often thought, in fact, that if we take a Lockean, relational position on objective bitterness, redness, etc., we *must* complete our account of the relations in question by appeal to non-relational, intrinsic properties. If what it is to be objectively bitter is to produce a certain effect in the members of the class of normal observers, we must be able to specify that effect and distinguish it from the effect produced by objective sourness and so forth.

What else could distinguish this effect but some intrinsic property? Why not another relational or extrinsic property? The relational treatment of monetary value does not require, for its completion, the supposition of items of intrinsic value (value independent of the valuers' dispositions to react behaviourally). The claim that certain perceptual properties are different is, in the absence of any supporting argument, just question-begging. It will not do to say that it is just obvious that they are intrinsic. It may have seemed obvious to some, but the considerations raised by Chase's quandary show that it is far from obvious that any intrinsic property (whatever that comes to) could play the role for the Lockean, relational treatment of the public perceptual properties.

Why not give up intrinsicity as a second-order property altogether, at least pending resolution of the disarray of philosophical opinion about what intrinsicity might be? Until such time the insistence that qualia are the intrinsic properties of experience is an empty gesture at best; no one could claim that it provides a clear, coherent, understood prerequisite for theory.¹¹

What, then, of ineffability? Why does it seem that our conscious experiences have ineffable properties? Because they do have *practically* ineffable properties. Suppose, in *intuition pump 13: the osprey cry*, that I have never heard the cry of an osprey, even in a recording, but know roughly, from reading my bird books, what to listen for "a series of short, sharp, cheeping whistles, *cheep, cheep* or *chewk chewk*, etc; sounds

annoyed" (Peterson 1947) (or words to that effect or better). The verbal description gives me a partial confinement of the logical space of possible bird cries. On its basis I can rule out many bird calls I have heard or might hear, but there is still a broad range of discriminable-by-me possibilities within which the actuality lies hidden from me like a needle in a haystack.

Then one day, armed with both my verbal description and my binoculars, I identify an osprey visually, and then hear its cry. "So *that's* what it sounds like," I say to myself, ostending—it seems—a particular mental complex of intrinsic, ineffable qualia. I dub the complex 'S' (*pace* Wittgenstein), rehearse it in short-term memory, check it against the bird book descriptions, and see that, while the verbal descriptions are true, accurate, and even poetically evocative—I decide I could not do better with a thousand words—they still fall short of *capturing* the qualia complex I have called S. In fact, that is why I need the neologism, 'S', to refer directly to the ineffable property I cannot pick out by description. My perceptual experience has pin-pointed for me the location of the osprey cry in the logical space of possibilities in a way verbal description could not.

But tempting as this view of matters is, it is overstated. First of all, it is obvious that from a single experience of this sort I do not—cannot—know how to generalize to other osprey calls. Would a cry that differed only in being half an octave higher also be an osprey call? That is an empirical, ornithological question for which my experience provides scant evidence. But moreover—and this is a psychological, not ornithological, matter—I do not and cannot know, from a single such experience, which physical variations and constancies in stimuli would produce an indistinguishable experience in me. Nor can I know whether I would react the same (have the same experience) if I were presented with what was, by all physical measures, a re-stimulation identical to the first. I cannot know the modulating effect, if any, of variations in my body (or psyche).

This inscrutability of projection is surely one of the sources of plausibility of Wittgenstein's scepticism regarding the possibility of a private language:

Wittgenstein emphasizes that ostensive definitions are always in principle capable of being misunderstood, even the ostensive definition of a color word such as 'sepia'. How someone understands the word is exhibited in the way someone goes on, 'the use that he makes of the word defined'. One may go on in the right way given a purely minimal explanation, while on the other hand one may go on in another way no matter how many clarifications are added, since these too can be misunderstood (Kripke 1982, 83; see also 40–46).

But what is inscrutable in a single glance, and somewhat ambiguous after limited testing, can come to be justifiably seen as the deliverance

of a highly specific, reliable, and projectible property detector, once it has been field-tested under a suitably wide variety of circumstances.

In other words, when first I hear the osprey cry, I may have identified a property detector in myself, but I have no idea (yet) what property my newfound property detector detects. It might seem then that I know nothing new at all—that my novel experience has not improved my epistemic predicament in the slightest. But of course this is not so. I may not be able to describe the property or identify it relative to any readily usable public landmarks (yet), but I am acquainted with it in a modest way: I can refer to the property I detected: it is the property I detected in *that* event. My experience of the osprey cry has given me a new way of thinking about osprey cries (an unavoidably inflated way of saying something very simple) which is practically ineffable both because it has (as yet for me) an untested profile in response to perceptual circumstances, and because it is—as the poverty of the bird book description attests—such a highly informative way of thinking: a deliverance of an informationally very sensitive portion of my nervous system.

In this instance I mean information in the formal information theory sense of the term. Consider (*intuition pump 14: the Jello box*) the old spy trick, most famously encountered in the case of Julius and Ethel Rosenberg, of improving on a password system by tearing something in two (a Jello box, in the Rosenberg's case), and giving half to each of the two parties who must be careful about identifying each other. Why does it work? Because tearing the paper in two produces an edge of such informational complexity that it would be virtually impossible to reproduce by deliberate construction. (Cutting the Jello box along a straight edge with a razor would entirely defeat the purpose.) The particular jagged edge of one piece becomes a *practically* unique pattern-recognition device for its mate; it is an apparatus for detecting the shape property *M*, where *M* is uniquely instantiated by its mate. It is of the essence of the trick that we cannot replace our dummy predicate '*M*' with a longer, more complex, but accurate and exhaustive description of the property, for, if we could, we could use the description as a recipe or feasible algorithm for producing another instance of *M* or another *M* detector. The only *readily available* way of saying what property *M* is is just to point to our *M* detector and say that *M* is the shape property detected by this thing here.

And that is just what we do when we seem to ostend, with the mental finger of inner intention, a quale or qualia complex in our experience. We refer to a property—a public property of uncharted boundaries—via reference to our personal and idiosyncratic capacity to respond to it. That idiosyncrasy is the extent of our privacy. If I wonder whether your blue is my blue, your middle C is my middle C, I can coherently be wondering whether our discrimination profiles over a

wide variation in conditions will be approximately the same. And they may not be; people experience the world quite differently. But that is empirically discoverable by all the usual objective testing procedures.¹²

Peter Bieri has pointed out to me that there is a natural way of exploiting Dretske's (1981) sense of information in a reformulation of my first three second-order properties of qualia: intrinsicity, ineffability, and privacy. (There are problems with Dretske's attempt to harness information theory in this way—see my discussion in Dennett 1987a, chapter 8—but they are not relevant to this point.) We could speak of what Bieri would call 'phenomenal information properties' of psychological events. Consider the information—what Dretske would call the *natural meaning*—that a type of internal perceptual event might carry. That it carries that information is an objective (and hence, in a loose sense, intrinsic) matter since it is independent of what information (if any) the subject *takes* the event type to carry. Exactly what information is carried is (practically) ineffable, for the reasons just given. And it is private in the sense just given: proprietary and potentially idiosyncratic.

Consider how Bieri's proposed 'phenomenal information properties' (let us call them *pips*) would apply in the case of Chase and Sanborn. Both Chase and Sanborn ought to wonder whether their pips have changed. Chase's speech shows that he is under the impression that his pips are unchanged (under normal circumstances—all bets are off if he has just eaten horse-radish). He believes that the same objective things in the world—in particular, chemically identical caffeine-rich fluids—give rise to his particular types of taste experiences now as six years ago.

Sanborn is under the impression that his pips are different. He thinks his objective property detectors are deranged. He no longer has confidence that their deliverances today inform him of what they did six years ago. And what, exactly, did they inform him of then? If Sanborn were an ordinary person, we would not expect him to have an explicit answer, since most of us treat our taste detectors as mere *M* detectors, detecting whatever it is that they detect. (There are good reasons for this, analysed by Akins 1987.) But professional coffee-tasters are probably different. They probably have some pretty good idea of what kind of chemical-analysis transduction machinery they have in their mouths and nervous systems.

So far, so good. We could reinterpret Chase's and Sanborn's speeches as hypotheses about the constancies or change in the outputs of their perceptual information-processing apparatus, and just the sort of empirical testing we imagined before would tend to confirm or disconfirm their opinions thus interpreted. But what would justify calling such an information-bearing property 'phenomenal'?

Such a pip has, as the testimony of Chase and Sanborn reveals, the power to provoke in Chase and Sanborn acts of (apparent) re-identifi-

cation or recognition. This power is of course a Lockean, dispositional property on a par with the power of bitter things to provoke a certain reaction in people. It is this power alone, however it might be realized in the brain, that gives Chase and Sanborn 'access' to the deliverances of their individual property detectors.

We may 'point inwardly' to one of the deliverances of our idiosyncratic, proprietary property detectors, but when we do, what are we pointing *at*? What does that deliverance itself *consist of*? Or what are its consciously apprehensible properties, if not just our banished friends the qualia? We must be careful here, for if we invoke an inner perceptual process in which we observe the deliverance with some inner eye and thereby discern its properties, we will be stepping back into the frying pan of the view according to which qualia are just ordinary properties of our inner states.

But nothing requires us to make such an invocation. We do not have to know how we identify or re-identify or gain access to such internal response types in order to be able so to identify them. This is a point that was forcefully made by the pioneer functionalists and materialists, and has never been rebutted (Farrell 1950; Smart 1959). The properties of the 'thing experienced' are not to be confused with the properties of the event that realizes the experiencing. To put the matter vividly, the physical difference between someone's imagining a purple cow and imagining a green cow *might* be nothing more than the presence or absence of a particular zero or one in one of the brain's 'registers'. Such a brute physical presence is all that it would take to anchor the sorts of dispositional differences between imagining a purple cow and imagining a green cow that could then flow, causally, from that 'intrinsic' fact. (I doubt that this is what the friends of qualia have had in mind when they have insisted that qualia are intrinsic properties.)

Moreover, it is our very inability to expand on, or modify, these brute dispositions so to identify or recognize such states that creates the doctrinal illusion of 'homogeneity' or 'atomicity to analysis' or 'grainlessness' that characterizes the qualia of philosophical tradition.

This putative grainlessness, I hypothesize, is nothing but a sort of functional invariability: it is close kin to what Pylyshyn (1980, 1984) calls *cognitive impenetrability*. Moreover, this functional invariability or impenetrability is not absolute but itself plastic over time. Just as on the efferent side of the nervous system, *basic actions*—in the sense of Danto (1963, 1965) and others (see Goldman 1970)—have been discovered to be variable, and subject under training to decomposition (one can learn with the help of 'biofeedback' to will the firing of a particular motor neuron 'directly'), so what counts for an individual as the simple or atomic properties of experienced items is subject to variation with training.¹³

Consider the results of 'educating' the palate of a wine-taster, or 'ear training' for musicians. What had been 'atomic' or 'unanalysable' be-

comes noticeably compound and describable; pairs that had been indistinguishable become distinguishable, and when this happens we say *the experience changes*. A swift and striking example of this is illustrated in *intuition pump 15: the guitar string*. Pluck the bass or low E string open and listen carefully to the sound. Does it have describable parts or is it one and whole and ineffably guitarish? Many will opt for the latter way of talking. Now pluck the open string again and carefully bring a finger down lightly over the octave fret to create a high 'harmonic'. Suddenly a *new* sound is heard: 'purer' somehow and of course an octave higher. Some people insist that this is an entirely novel sound, while others will describe the experience by saying "the bottom fell out of the note"—leaving just the top. But then on a third open plucking one can hear, with surprising distinctness, the harmonic overtone that was isolated in the second plucking. The homogeneity and ineffability of the first experience is gone, replaced by a duality as 'directly apprehensible' and clearly describable as that of any chord.

The difference in experience is striking, but the complexity apprehended on the third plucking was *there* all along (being responded to or discriminated). After all, it was by the complex pattern of overtones that you were able to recognize the sound as that of a guitar rather than of a lute or harpsichord. In other words, although the subjective experience has changed dramatically, the *pip* has not changed; you are still responding, as before, to a complex property so highly informative that it practically defies verbal description.

There is nothing to stop further refinement of one's capacity to describe this heretofore ineffable complexity. At any time, of course, there is one's current horizon of distinguishability—and that horizon is what sets, if anything does, what we should call the primary or atomic properties of what one consciously experiences (Farrell 1950). But it would be a mistake to transform the fact that inevitably there is a limit to our capacity to describe things we experience into the supposition that there are absolutely indescribable properties in our experience.

So when we look one last time at our original characterization of qualia, as ineffable, intrinsic, private, directly apprehensible properties of experience, we find that there is nothing to fill the bill. In their place are relatively or practically ineffable public properties we can refer to indirectly via reference to our private property detectors—private only in the sense of idiosyncratic. And in so far as we wish to cling to our subjective authority about the occurrence within us of states of certain types or with certain properties, we can have some authority—not infallibility or incorrigibility, but something better than sheer guessing—but only if we restrict ourselves to relational, extrinsic properties like the power of certain internal states of ours to provoke acts of apparent re-identification. So contrary to what seems obvious at first blush, there simply are no qualia at all.

Acknowledgments

The first version of this paper was presented at University College, London, in November 1978, and in various revisions at a dozen other universities in 1979 and 1980. It was never published, but was circulated widely as Tufts University Cognitive Science Working Paper 7, December 1979. A second version was presented at the Universities of Adelaide and Sydney in 1984, and in 1985 to psychology department colloquia at Harvard and Brown under the title "Properties of conscious experience." The second version was the basis for my presentation at the workshop from which this book arises, and was circulated in pre-print in 1985, again under the title "Quining qualia." The present version, the fourth, is a substantial revision, thanks to the helpful comments of many people, including Kathleen Akins, Ned Block, Alan Cowey, Sydney Shoemaker, Peter Bieri, William Lycan, Paul Churchland, Gilbert Harman, and the participants at Villa Olmo.

Notes

1. A representative sample of the most recent literature on qualia would include Block 1980; Shoemaker 1981, 1982; Davis 1982; White 1985; Armstrong and Malcolm 1984; Churchland 1985; and Conee 1985.

2. The difference between 'eliminative materialism'—of which my position on qualia is an instance—and a 'reductive' materialism that takes on the burden of identifying the problematic item in terms of the foundational materialistic theory is thus often best seen not so much as a doctrinal issue but as a tactical issue: how might we most gracefully or effectively enlighten the confused in this instance? See my discussion of 'fatigues' in the Introduction to *Brainstorms* (Dennett 1978a) and, earlier, my discussion of what the enlightened ought to say about the metaphysical status of *sakes* and *voices* in *Content and consciousness* (Dennett 1969, chap. 1).

3. The plausibility of this concession depends less on a high regard for the technology than on a proper scepticism about human powers, now documented in a fascinating study by Lehrer (1983).

4. Shoemaker (1984, 356) seems to be moving reluctantly towards agreement with this conclusion: 'So unless we can find some grounds on which we can deny the possibility of the sort of situation envisaged . . . we must apparently choose between rejecting the functionalist account of qualitative similarity and rejecting the standard conception of qualia. I would prefer not to have to make this choice; but if I am forced to make it, I reject the standard conception of qualia'.

5. Shoemaker (1982) attributes a view to Wittgenstein (acknowledging that 'it is none too clear' that this is actually what Wittgenstein held) which is very close to the view I defend here. But to Shoemaker, 'it would seem offhand that Wittgenstein was mistaken' (p. 360), a claim Shoemaker supports with a far from offhand thought experiment—which Shoemaker misanalyses if the present paper is correct. (There is no good reason, contrary to Shoemaker's declaration, to believe that his subject's *experience* is systematically different from what it was before the inversion.) Smart (1959) expresses guarded and partial

approval of Wittgenstein's hard line, but cannot see his way clear to as uncompromising an eliminativism as I maintain here.

6. In 1979, I read an earlier version of this paper in Oxford, with a commentary by John Foster, who defended qualia to the last breath, which was: 'qualia should not be quined but fostered!' Symmetry demands, of course, the following definition for the most recent edition of *The philosophical lexicon* (Dennett 1987b): 'foster, v. To acclaim resolutely the existence or importance of something chimerical or insignificant.

7. This example first appeared in print in my reflections on Smullyan in *The Mind's I* (Hofstadter and Dennett 1981, 427–428).

8. Kripke (1982, 40) comes close when he asks rhetorically: 'Do I not know, directly, and with a fair degree of certainty, that I mean plus [by the function I call "plus"]?' [my emphasis]. Kripke does not tell us what is implied by a 'fair degree of certainty', but presumably he means by this remark to declare his allegiance to what Millikan (1984) attacks under the name of 'meaning rationalism'.

9. We can save the traditional claim by ignoring presumably private or subjective qualia and talking always of public tastes—such as the public taste of Maxwell House coffee that both Chase and Sanborn agree has remained constant. Individuals can be said to acquire a taste for such a public taste.

10. 'I am not so wild as to deny that my sensation of red today is like my sensation of red yesterday. I only say that the similarity can consist only in the physiological force behind consciousness—which leads me to say, I recognize this feeling the same as the former one, and so does not consist in a community of sensation.'—(Peirce, *Collected Works*, vol. V, 172 fn 2).

11. A heroic (and, to me, baffling) refusal to abandon intrinsicality is Sellars' (1981) contemplation over the years of his famous pink ice cube, which leads him to postulate a revolution in microphysics, restoring objective 'absolute sensory processes' in the face of Boyle and Locke and almost everybody since them (also see my commentary in Dennett 1981).

12. Stich (1983) discusses the implications for psychological theory of incommensurability problems that can arise from such differences in discrimination profiles (see, especially, chaps. 4 and 5).

13. See Churchland (1979, especially chap. 2) for supporting observations on the variability of perceptual properties, and for novel arguments against the use of 'intrinsic properties' as determiners of the meaning of perceptual predicates. See also Churchland (1985) for further arguments and observations in support of the position sketched here.

References

Akins, K. (1987). Information and organisms: or why Nature doesn't build epistemic engines. Unpublished PhD thesis, University of Michigan, Ann Arbor.

Armstrong, D. and N. Malcolm, ed. (1984). *Consciousness and causality*. Oxford: Blackwell Scientific.

- Bennett, J. (1965). Substance, reality and primary qualities. *American Philosophical Quarterly* 2, 1–17.
- Berkeley, G. (1713). *Three dialogues between Hylas and Philonous*. London.
- Block, N. (1978). Troubles with functionalism. In C. W. Savage, ed., *Perception and cognition: Issues in the foundations of psychology*. Minneapolis: University of Minnesota Press.
- Block, N. (1980). Are absent qualia impossible? *The Philosophical Review* 89, 257–274.
- Block, N., and J. Fodor (1972). What psychological states are not. *The Philosophical Review* 81, 159–181.
- Chisholm, R. (1976). *Person and object*. La Salle, IL: Open Court Press.
- Churchland, P. M. (1979). *Scientific realism and the plasticity of mind*. Cambridge: Cambridge University Press.
- Churchland, P. M. (1985). Reduction, qualia and the direct inspection of brain states. *Journal of Philosophy* 82, 8–28.
- Conee, E. (1985). The possibility of absent qualia. *The Philosophical Review* 94 (3), 345–366.
- Damasio, A., Yamada, T., Damasio, H., Corbett, J., and McKee, J. (1980). Central achromatopsia: behavioral, anatomic, and physiological aspects. *Neurology* 30, 1064–1071.
- Danto, A. (1963). What we can do. *Journal of Philosophy* 60, 435–445.
- Danto, A. (1965). Basic actions. *American Philosophical Quarterly* 60, 141–148.
- Davis, L. (1982). Functionalism and absent qualia. *Philosophical Studies* 41 (2), 231–251.
- Dennett, D. C. (1969). *Content and consciousness*. Andover: Routledge and Kegan Paul.
- Dennett, D. C. (1976). Are dreams experiences? *The Philosophical Review* 73, 151–171.
- Dennett, D. C. (1978a). *Brainstorms*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1978b). Two approaches to mental images. In D. C. Dennett, *Brainstorms*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1979). On the absence of phenomenology. In D. F. Gustafson and B. L. Tapscott, eds., *Body, mind, and method*. Dordrecht: D. Reidel.
- Dennett, D. C. (1981). Wondering where the yellow went. *The Monist* 64, 102–108.
- Dennett, D. C. (1982). How to study human consciousness empirically: Or nothing comes to mind. *Synthese* 53, 159–180.
- Dennett, D. C. (1987a). *The intentional stance*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1987b). *The philosophical lexicon*, (8th ed.). Copy available from the American Philosophical Association, University of Delaware, Newark, DE.

- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Elster, J. (1985). *Making sense of Marx*. Cambridge: Cambridge University Press.
- Farrell, B. (1950). Experience. *Mind* 59, 170–198.
- Gert, B. (1965). Imagination and verifiability. *Philosophical Studies* 16, 44–47.
- Geschwind, N. and Fusillo, M. (1966). Color-naming defects in association with alexia. *Archives of Neurology* 15, 137–146.
- Goldman, A. (1970). *A theory of human action*. Englewood Cliffs, NJ: Prentice-Hall.
- Gregory, R. (1977). *Eye and brain* (3rd ed.). London: Weidenfeld and Nicolson.
- Hofstadter, D., and D. C. Dennett (1981). *The mind's I: Fantasies and reflections on self and soul*. New York: Basic Books.
- Kim, J. (1982). Psychophysical supervenience. *Philosophical Studies* 41, 51–70.
- Kitcher, P. (1979). Phenomenal qualities. *American Philosophical Quarterly* 16, 123–129.
- Kripke, S. (1982). *Wittgenstein on rules and private language*. Cambridge, MA: Harvard University Press.
- Lehrer, A. (1983). *Wine and conversation*. Bloomington: Indiana University Press.
- Lewis, D. (1983). Extrinsic properties. *Philosophical Studies* 44, 197–200.
- Locke, J. (1959). *An essay concerning human understanding*, ed. A. C. Fraser. New York: Dover.
- Lycan, W. (1973). Inverted spectrum. *Ratio* 15, 315–319.
- Malcolm, N. (1956). Dreaming and skepticism. *The Philosophical Review*, 65, 14–37.
- Malcolm, N. (1959). *Dreaming*. Andover: Routledge and Kegan Paul.
- Meadows, J. C. (1974). Disturbed perception of colours associated with localized cerebral lesions. *Brain* 97, 615–632.
- Millikan, R. (1984). *Language, thought and other biological categories*. Cambridge, MA: MIT Press.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83, 435–451.
- Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.
- Peirce, C. (1931–58). *Collected Works*, vol. V, ed., C. Hartshorne and P. Weiss. Cambridge MA: Harvard University Press.
- Peterson, R. T. (1947). *A field guide to the birds*. Boston: Houghton Mifflin.

- Putnam, H. (1965). Brains and behavior. In J. Butler, ed., *Analytical philosophy*, second series, Oxford: Blackwell Scientific.
- Pylyshyn, Z. (1980). Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences* 3, 111–132.
- Pylyshyn, Z. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
- Sellars, W. (1981). Foundations for a metaphysics of pure process (the Carus Lectures). *The Monist* 64, 3–90.
- Shoemaker, S. (1969). Time without change. *Journal of Philosophy* 66, 363–381.
- Shoemaker, S. (1975). Functionalism and qualia. *Philosophical Studies* 27, 291–315.
- Shoemaker, S. (1981). Absent qualia are impossible—A reply to Block. *The Philosophical Review* 90, 581–599.
- Shoemaker, S. (1982). The inverted spectrum. *Journal of Philosophy* 79, 357–381.
- Shoemaker, S. (1984). *Identity, cause, and mind*. Cambridge: Cambridge University Press.
- Smart, J. C. (1959). Sensations and brain processes. *The Philosophical Review* 68, 141–156.
- Stich, S. (1983). *From folk psychology to cognitive science: The case against belief*. Cambridge, MA: MIT Press.
- Taylor, D. M. (1966). The incommunicability of content. *Mind* 75, 527–541
- Von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron response. *Science* 224, 1260–1262.
- Walzer, M. (1985). What's left of Marx. *The New York Review of Books*, 21 November, 43–46.
- White, S. (1985). Professor Shoemaker and so-called 'qualia' of experience. *Philosophical Studies* 47, 369–383.
- Wittgenstein, L. (1958). *Philosophical investigations*. Oxford: Blackwell Scientific.