

David Lewis

Experience the Best Teacher

They say that experience is the best teacher, and the classroom is no substitute for Real Life. There's truth to this. If you want to know what some new and different experience is like, you can learn it by going out and really *having* that experience. You can't learn it by being told about the experience, however thorough your lessons may be.

Does this prove much of anything about the metaphysics of mind and the limits of science? I think not.

Example: Skunks and Vegemite. I have smelled skunks, so I know what it's like to smell skunks. But skunks live only in some parts of the world, so you may never have smelled a skunk. If you haven't smelled a skunk, then you don't know what it's like. You never will, unless someday you smell a skunk for yourself. On the other hand, you may have tasted Vegemite, that famous Australian substance; and I never have. So you may know what it's like to taste Vegemite. I don't, and unless I taste Vegemite (what, and spoil a good example!), I never will. It won't help at all to take lessons on the chemical composition of skunk scent or Vegemite, the physiology of the nostrils or the taste-buds, and the neurophysiology of the sensory nerves and the brain.

Example: The Captive Scientist.¹ Mary, a brilliant scientist, has lived from birth in a cell where everything is black or white. (Even she herself is painted all over.) She views the world on black-and-white television. By television she reads books, she joins in discussion, she watches the results of experiments done under her direction. In this way she becomes the

world's leading expert on color and color vision and the brain states produced by exposure to colors. But she doesn't know what it's like to see color. And she never will, unless she escapes from her cell.

Example: The Bat.² The bat is an alien creature, with a sonar sense quite unlike any sense of ours. We can never have the experiences of a bat; because we could not become bat-like enough to have those experiences and still be ourselves. We will never know what it's like to be a bat. Not even if we come to know all the facts there are about the bat's behavior and behavioral dispositions, about the bat's physical structure and processes, about the bat's functional organization. Not even if we come to know all the same sort of physical facts about all the other bats, or about other creatures, or about ourselves. Not even if we come to possess all physical facts whatever. Not even if we become able to recognize all the mathematical and logical implications of all these facts, no matter how complicated and how far beyond the reach of finite deduction.

Experience is the best teacher, in this sense: having an experience is the best way or perhaps the only way, of coming to know what that experience is like. No amount of scientific information about the stimuli that produce that experience and the process that goes on in you when you have that experience will enable you to know what it's like to have the experience.

... But Not Necessarily

Having an experience is surely one good way, and surely the only practical way, of coming to know what that experience is like. Can we say, flatly, that it is the only *possible* way? Probably not. There is a change that takes place in you when you have the experience and thereby come to know what it's like. Perhaps the exact same change could in principle be produced in you by precise neurosurgery, very far beyond the limits of present-day technique. Or it could possibly be produced in you by magic. If we ignore the laws of nature, which are after all contingent, then there is no necessary connection between cause and effect: anything could cause anything. For instance, the casting of a spell could do to you exactly what your first smell of skunk would do. We might quibble about whether a state produced in this artificial fashion would deserve the *name* 'knowing what

it's like to smell a skunk', but we can imagine that so far as what goes on within you is concerned, it would differ not at all.³

Just as we can imagine that a spell might produce the same change as a smell, so likewise we can imagine that science lessons might cause that same change. Even that is possible, in the broadest sense of the word. If we ignored all we know about how the world really works, we could not say what might happen to someone if he were taught about the chemistry of scent and the physiology of the nose. There might have been a causal mechanism that transforms science lessons into whatever it is that experience gives us. But there isn't. It is not an absolutely necessary truth that experience is the best teacher about what a new experience is like. It's a contingent truth. But we have good reason to think it's true.

We have good reason to think that something of this kind is true, anyway, but less reason to be sure exactly what. Maybe some way of giving the lessons that hasn't yet been invented, and some way of taking them in that hasn't yet been practiced, could give us a big surprise. Consider sight-reading: a trained musician can read the score and know what it would be like to hear the music. If I'd never heard that some people can sight-read, I would never have thought it humanly possible. Of course the moral is that new music isn't altogether new—the big new experience is a rearrangement of lots of little old experiences. It just might turn out the same for new smells and tastes *vis-à-vis* old ones; or even for color vision *vis-à-vis* black and white;⁴ or even for sonar sense experience *vis-à-vis* the sort we enjoy. The thing we can say with some confidence is that we have no faculty for knowing on the basis of mere science lessons what some *new enough* experience would be like. But how new is "new enough"?—There, we just might be in for surprises.

Three Ways to Miss the Point

The First Way. A literalist might see the phrase 'know what it's like' and take that to mean: 'know what it resembles'. Then he might ask: what's so hard about that? Why can't you just be told which experiences resemble one another? You needn't have had the experiences—all you need, to be taught your lessons, is some way of referring to them. You could be told: the smell of skunk somewhat resembles the smell of burning rubber. I have

been told: the taste of Vegemite somewhat resembles that of Marmite. Black-and-white Mary might know more than most of us about the resemblances among color-experiences. She might know which ones are spontaneously called 'similar' by subjects who have them; which gradual changes from one to another tend to escape notice; which ones get conflated with which in memory; which ones involve roughly the same neurons firing in similar rhythms; and so forth. We could even know what the bat's sonar experiences resemble just by knowing that they do not at all resemble any experiences of humans, but do resemble—as it might be—certain experiences that occur in certain fish. This misses the point. *Pace* the literalist, 'know what it's like' does not mean 'know what it resembles'. The most that's true is that knowing what it resembles *may* help you to know what it's like. If you are taught that experience A resembles B and C closely, D less, E not at all, that will help you know what A is like—if you know already what B and C and D and E are like. Otherwise, it helps you not at all. I don't know any better what it's like to taste Vegemite when I'm told that it tastes like Marmite, because I don't know what Marmite tastes like either. (Nor do I know any better what Marmite tastes like for being told it tastes like Vegemite.) Maybe Mary knows enough to triangulate each color experience exactly in a network of resemblances, or in many networks of resemblance in different respects, while never knowing what any node of any network is like. Maybe we could do the same for bat experiences. But no amount of information about resemblances, just by itself, does anything to help us know what an experience is like.

The Second Way. Insofar as I don't know what it would be like to drive a steam locomotive fast on a cold, stormy night, part of my problem is just that I don't know what experiences I would have. The firebox puts out a lot of heat, especially when the fireman opens the door to throw on more coal; on the other hand, the cab is drafty and gives poor protection from the weather. Would I be too hot or too cold? Or both by turns? Or would it be chilled face and scorched legs? If I knew the answers to such questions, I'd know much better what it would be like to drive the locomotive. So maybe 'know what it's like' just means 'know what experiences one has'. Then again: what's the problem? Why can't you just be told what experiences you would have if, say, you tasted Vegemite? Again, you needn't have had the experiences—all you need, to be taught your lessons, is some way of

referring to them. We have ways to refer to experiences we haven't had. We can refer to them in terms of their causes: the experience one has upon tasting Vegemite, the experience one has upon tasting a substance of such-and-such chemical composition. Or we can refer to them in terms of their effects: the experience that just caused Fred to say "Yeeuch!" Or we can refer to them in terms of the physical states of the nervous system that mediate between those causes and effects: the experience one has when one's nerves are firing in such-and-such pattern. (According to some materialists, I myself for one, this means the experience which is identical with such-and-such firing pattern. According to other materialists it means the experience which is realized by such-and-such firing pattern. According to many dualists, it means the experience which is merely the lawful companion of such-and-such firing pattern. But whichever it is, we get a way of referring to the experience.) Black-and-white Mary is in a position to refer to color-experiences in all these ways. Therefore you should have no problem in telling her exactly what experiences one has upon seeing the colors. Or rather, your only problem is that you'd be telling her what she knows very well already! In general, to know what is the X is to know that the X is the Y, where it's not too obvious that the X is the Y. (Just knowing that the X is the X won't do, of course, because it is too obvious.) If Mary knows that the experience of seeing green is the experience associated with such-and-such pattern of nerve firings, then she knows the right sort of unobvious identity. So she knows what experience one has upon seeing green.

(Sometimes it's suggested that you need a 'rigid designator': you know what is the X by knowing that the X is the Y only if 'the Y' is a term whose referent does not depend on any contingent matter of fact. In the first place, this suggestion is false. You can know who is the man on the balcony by knowing that the man on the balcony is the Prime Minister even if neither 'the Prime Minister' nor any other phrase available to you rigidly designates the man who is, in fact, the Prime Minister. In the second place, according to one version of Materialism [the one I accept] a description of the form 'the state of having nerves firing in such-and-such a pattern' is a rigid designator, and what it designates is in fact an experience; and according to another version of Materialism, a description of the form 'having some or other state which occupies so-and-so functional role' is a rigid designator of an experience. So even if the false suggestion were granted,

still it hasn't been shown, without begging the question against Materialism, that Mary could not know what experience one has upon seeing red.)

Since Mary *does* know what experiences she would have if she saw the colors, but she *doesn't* know what it would be like to see the colors, we'd better conclude that 'know what it's like' does not after all mean 'know what experiences one has'. The locomotive example was misleading. Yes, by learning what experiences the driver would have, I can know what driving the locomotive would be like; but only because I already know what those experiences are like. (It matters that I know what they're like under the appropriate descriptions—as it might be, the description 'chilled face and scorched legs'. This is something we'll return to later.) Mary may know as well as I do that when the driver leans out into the storm to watch the signals, he will have the experience of seeing sometimes green lights and sometimes red. She knows better than I what experiences he has when signals come into view. She can give many more unobviously equivalent descriptions of those experiences than I can. But knowing what color-experiences the driver has won't help Mary to know what his job is like. It will help me.

The Third Way. Until Mary sees green, here is one thing she will never know: she will never know that she is seeing green. The reason why is just that until she sees green, it will never be true that she is seeing green. Some knowledge is irreducibly egocentric, or *de se*.⁵ It is not just knowledge about what goes on in the world; it is knowledge of who and when in the world one is. Knowledge of what goes on in the world will be true alike for all who live in that world; whereas egocentric knowledge may be true for one and false for another, or true for one at one time and false for the same one at another time. Maybe Mary knows in advance, as she plots her escape, that 9 A.M. on the 13th of May, 1997, is the moment when someone previously confined in a black-and-white cell sees color for the first time. But until that moment comes, she will never know that she herself is then seeing color—because she isn't. What isn't true isn't knowledge. This goes as much for egocentric knowledge as for the rest. So only those of whom an egocentric proposition is true can know it, and only at times when it is true of them can they know it. That one is then seeing color is an egocentric proposition. So we've found a proposition which Mary can never know until she sees color—which, as it happens, is the very moment when she

will first know what it's like to see color! Have we discovered the reason why experience is the best teacher? And not contingently after all, but as a necessary consequence of the logic of egocentric knowledge?

No; we have two separate phenomena here, and only some bewitchment about the 'first-person perspective' could make us miss the difference. In the first place, Mary will probably go on knowing what it's like to see green after she stops knowing the egocentric proposition that she's then seeing green. Since what isn't true isn't known she must stop knowing that proposition the moment she stops seeing green. (Does that only mean that we should have taken a different egocentric proposition: that one *has* seen green? No; for in that case Mary could go on knowing the proposition even after she forgets what it's like to see green, as might happen if she were soon recaptured.) In the second place, Mary might come to know what it's like to see green even if she didn't know the egocentric proposition. She might not have known in advance that her escape route would take her across a green meadow, and it might take her a little while to recognize grass by its shape. So at first she might know only that she was seeing some colors or other, and thereby finding out what some color-experiences or other were like, without being able to put a name either to the colors or to the experiences. She would then know what it was like to see green, though not under that description, indeed not under any description more useful than 'the color-experience I'm having now'; but she would not know the egocentric proposition that she is then seeing green, since she wouldn't know which color she was seeing. In the third place, the gaining of egocentric knowledge may have prerequisites that have nothing to do with experience. Just as Mary can't know she's seeing green until she *does* see green, she can't know she's turning 50 until she *does* turn 50. But—I hope!—turning 50 does not involve some special experience. In short, though indeed one can gain egocentric knowledge that one is in some situation only when one is in it, that is not the same as finding out what an experience is like only when one has that experience.

We've just rejected two suggestions that don't work separately, and we may note that they don't work any better when put together. One knows what is the X by knowing that the X is the Y, where the identity is not too obvious; and 'the Y' might be an egocentric description. So knowledge that the X is the Y might be irreducibly egocentric knowledge, therefore knowledge that cannot be had until it is true of one that the X is the Y. So one

way of knowing what is the X will remain unavailable until it comes true of one that the X is the Y. One way that I could gain an unobvious identity concerning the taste of Vegemite would be for it to come true that the taste of Vegemite was the taste I was having at that very moment—and that would come true at the very moment I tasted Vegemite and found out what it was like! Is this why experience is the best teacher?—No; cases of gaining an unobvious egocentric identity are a dime a dozen, and most of them do not result in finding out what an experience is like. Suppose I plan ahead that I will finally break down and taste Vegemite next Thursday noon. Then on Wednesday noon, if I watch the clock, I first gain the unobvious egocentric knowledge that the taste of Vegemite is the taste I shall be having in exactly 24 hours, and thereby I have a new way of knowing what is the taste of Vegemite. But on Wednesday noon I don't yet know what it's like. Another example: from time to time I find myself next to a Vegemite-taster. On those occasions, and only those, I know what is the taste of Vegemite by knowing that it is the taste being had by the person next to me. But on no such occasion has it ever yet happened that I knew what it was like to taste Vegemite.

The Hypothesis of Phenomenal Information

No amount of the physical information that black-and-white Mary gathers could help her know what it was like to see colors; no amount of the physical information that we might gather about bats could help us know what it's like to have their experiences; and likewise in other cases. There is a natural and tempting explanation of why physical information does not help. That is the hypothesis that besides physical information there is an irreducibly different kind of information to be had: *phenomenal information*. The two are independent. Two possible cases might be exactly alike physically, yet differ phenomenally. When we get physical information we narrow down the physical possibilities, and perhaps we narrow them down all the way to one, but we leave open a range of phenomenal possibilities. When we have an experience, on the other hand, we acquire phenomenal information; possibilities previously open are eliminated; and that is what it is to learn what the experience is like.

(Analogy. Suppose the question concerned the location of a point within a certain region of the x - y plane. We might be told that its x -coordinate lies

in certain intervals, and outside certain others. We might even get enough of this information to fix the x -coordinate exactly. But no amount of x -information would tell us anything about the y -coordinate; any amount of x -information leaves open all the y -possibilities. But when at last we make a y -measurement, we acquire a new kind of information; possibilities previously open are eliminated; and that is how we learn where the point is in the y -direction.)

What might the subject matter of phenomenal information be? *If* the hypothesis of phenomenal information is true, then you have an easy answer: it is information about experience. More specifically, it is information about a certain part or aspect or feature of experience. But if the hypothesis is false, then there is still experience (complete with all its parts and aspects and features) and yet no information about experience is phenomenal information. So it cannot be said in a neutral way, without presupposing the hypothesis, that information about experience is phenomenal information. For if the hypothesis is false and materialism is true, it may be that all the information there is about experience is physical information, and can very well be presented in lessons for the inexperienced.

It makes no difference to put some fashionable new phrase in place of 'experience'. If instead of 'experience' you say 'raw feel' (or just 'feeling'), or 'way it feels', or 'what it's like', then I submit that you mean nothing different. Is there anything it's like to be this robot? Does this robot have experiences?—I can tell no difference between the new question and the old. Does sunburn feel the same way to you that it does to me? Do we have the same raw feel? Do we have the same experience when sunburned?—Again, same question. 'Know the feeling', 'know what it's like'—interchangeable. (Except that the former may hint at an alternative to the hypothesis of phenomenal information.) So if the friend of phenomenal information says that its subject matter is raw feels, or ways to feel, or what it's like, then I respond just as I do if he says that the subject matter is experience. Maybe so, *if* the hypothesis of phenomenal information is true; but if the hypothesis is false and materialism is true, nevertheless there is still information about raw feels, ways to feel or what it's like; but in that case it is physical information and can be conveyed in lessons.

We might get a candidate for the subject matter of phenomenal information that is not just experience renamed, but is still tendentious. For instance, we might be told that phenomenal information concerns the

intrinsic character of experience. A friend of phenomenal information might indeed believe that it reveals certain special, nonphysical intrinsic properties of experience. He might even believe that it reveals the existence of some special nonphysical thing or process, *all* of whose intrinsic properties are nonphysical. But he is by no means alone in saying that experience has an intrinsic character. Plenty of us materialists say so too. We say that a certain color-experience is whatever state occupies a certain functional role. So if the occupant of that role (universally, or in the case of humans, or in the case of certain humans) is a certain pattern of neural firing, then that pattern of firing *is* the experience (in the case in question). Therefore the intrinsic character of the experience is the intrinsic character of the firing pattern. For instance, a frequency of firing is part of the intrinsic character of the experience. If we materialists are right about what experience is, then black-and-white Mary knows all about the intrinsic character of color-experience; whereas most people who know what color-experience is like remain totally ignorant about its intrinsic character.⁶

To say that phenomenal information concerns 'qualia' would be tendentious in much the same way. For how was this notion introduced? Often thus. We are told to imagine someone who, when he sees red things, has just the sort of experiences that we have when we see green things, and vice versa; and we are told to call this a case of 'inverted qualia'. And then we are told to imagine someone queerer still, who sees red and responds to it appropriately, and indeed has entirely the same functional organization of inner states as we do and yet has no experiences at all; and we are told to call this a case of 'absent qualia'. Now a friend of phenomenal information might well think that these deficiencies have something to do with the nonphysical subject matter of phenomenal information. But others can understand them otherwise. Some materialists will reject the cases outright, but others, and I for one, will make sense of them as best we can. Maybe the point is that the states that occupy the roles of experiences, and therefore *are* the experiences, in normal people are inverted or absent in victims of inverted or absent qualia. (This presupposes, what might be false, that most people are enough alike.) Experience of red—the state that occupies that role in normal people—occurs also in the victim of 'inverted qualia', but in him it occupies the role of experience of green; whereas the state that occupies in him the role of experience of red is the state that occupies in normal people the role of experience of green. Experience of red and

of green—that is, the occupants of those roles for normal people—do not occur at all in the victim of 'absent qualia'; the occupants of those roles for him are states that don't occur at all in the normal. Thus we make good sense of inverted and absent qualia; but in such a way that 'qualia' is just the word for role-occupying states taken *per se* rather than *qua* occupants of roles. Qualia, so understood, could not be the subject matter of phenomenal information. Mary knows all about them. We who have them mostly don't.⁷

It is best to rest content with an unhelpful name and a *via negativa*. Stipulate that 'the phenomenal aspect of the world' is to name whatever is the subject matter of phenomenal information, if there is any such thing; the phenomenal aspect, if such there be, is that which we can become informed about by having new experiences but never by taking lessons. Having said this, it will be safe to say that information about the phenomenal aspect of the world can only be phenomenal information. But all we really know, after thus closing the circle, is that phenomenal information is supposed to reveal the presence of some sort of nonphysical things or processes within experience, or else it is supposed to reveal that certain physical things or processes within experience have some sort of nonphysical properties.

The Knowledge Argument

If we invoke the hypothesis of phenomenal information to explain why no amount of physical information suffices to teach us what a new experience is like, then we have a powerful argument to refute any materialist theory of the mind. Frank Jackson (see note 1) calls it the "knowledge argument." Arguments against one materialist theory or another are never very conclusive. It is always possible to adjust the details. But the knowledge argument, if it worked, would directly refute the bare minimum that is common to *all* materialist theories.

It goes as follows. First in a simplified form; afterward we'll do it properly. Minimal materialism is a supervenience thesis: no difference without physical difference. That is: any two possibilities that are just alike physically are just alike *simpliciter*. If two possibilities are just alike physically, then no physical information can eliminate one but not both of them. If two possibilities are just alike *simpliciter* (if that is possible) then no information

whatsoever can eliminate one but not both of them. So if there is a kind of information—namely, phenomenal information—that can eliminate possibilities that any amount of physical information leaves open, then there must be possibilities that are just alike physically, but not just alike *simpliciter*. That is just what minimal materialism denies.

(Analogy. If two possible locations in our region agree in their *x*-coordinate, then no amount of *x*-information can eliminate one but not both. If, *per impossibile*, two possible locations agreed in all their coordinates, then no information whatsoever could eliminate one but not both. So if there is a kind of information—namely, *y*-information—that can eliminate locations that any amount of *x*-information leaves open, then there must be locations in the region that agree in their *x*-coordinate but not in all their coordinates.)

Now to remove the simplification. What we saw so far was the knowledge argument against materialism taken as a necessary truth, applying unrestrictedly to all possible worlds. But we materialists usually think that materialism is a contingent truth. We grant that there are spooky possible worlds where materialism is false, but we insist that our actual world isn't one of them. If so, then there might after all be two possibilities that are alike physically but not alike *simpliciter*; but one or both of the two would have to be possibilities where materialism was false. Spooky worlds could differ with respect to their spooks without differing physically. Our minimal materialism must be a *restricted* supervenience thesis: within a certain class of worlds, which includes our actual world, there is no difference without physical difference. Within that class, any two possibilities just alike physically are just alike *simpliciter*. But what delineates the relevant class? (It is trivial that our world belongs to *some* class wherein there is no difference without physical difference. That will be so however spooky our world may be. The unit class of our world is one such class, for instance. And so is any class that contains our world, and contains no two physical duplicates.) I think the relevant class should consist of the worlds that have nothing wholly alien to this world. The inhabitants of such a nonalien world could be made from the inhabitants of ours, so to speak, by a process of division and recombination. That will make no wholly different kinds of things, and no wholly different fundamental properties of things.⁸ Our restricted materialist supervenience thesis should go as follows: throughout the nonalien worlds, there is no difference without physical difference.

If the hypothesis of phenomenal information be granted, then the knowledge argument refutes this restricted supervenience nearly as decisively as it refutes the unrestricted version. Consider a possibility that is eliminated by phenomenal information, but not by any amount of physical information. There are two cases. Maybe this possibility has nothing that is alien to our world. In that case the argument goes as before: actuality and the eliminated possibility are just alike physically, they are not just alike *simpliciter*; furthermore, both of them fall within the restriction to nonalien worlds, so we have a counterexample even to restricted supervenience. Or maybe instead the eliminated possibility does have something *X* which is alien to this world—an alien kind of thing, or maybe an alien fundamental property of nonalien things. Then the phenomenal information gained by having a new experience has revealed something negative: at least in part, it is the information that *X* is *not* present. How can that be? If there is such a thing as phenomenal information, presumably what it reveals is positive: the presence of something hitherto unknown. Not, of course, something alien from actuality itself; but something alien from actuality as it is inadequately represented by the inexperienced and by the materialists. If Mary learns something when she finds out what it's like to see the colors, presumably she learns that there's *more* to the world than she knew before—not *less*. It's easy to think that phenomenal information might eliminate possibilities that are impoverished by comparison with actuality, but that would make a counterexample to the restricted supervenience thesis. To eliminate possibilities without making a counterexample, phenomenal information would have to eliminate possibilities less impoverished than actuality. And how can phenomenal information do that? Compare ordinary perceptual information. Maybe Jean-Paul can just *see* that Pierre is absent from the café, at least if it's a small café. But how can he just see that Pierre is absent from Paris, let alone from the whole of actuality?

(Is there a third case? What if the eliminated possibility is in one respect richer than actuality, in another respect poorer? Suppose the eliminated possibility has *X*, which is alien from actuality, but also it lacks *Y*. Then phenomenal information might eliminate it by revealing the actual presence of *Y*, without having to reveal the actual absence of *X*—but then I say there ought to be a third possibility, one with neither *X* nor *Y*, poorer and in no respect richer than actuality, and again without any physical

difference from actuality. For why should taking away X automatically restore Y? Why can't they vary independently?⁹ But this third possibility differs *simpliciter* from actuality without differing physically. Further, it has nothing alien from actuality. So we regain a counterexample to the restricted supervenience thesis.)

The knowledge argument works. There is no way to grant the hypothesis of phenomenal information and still uphold materialism. Therefore I deny the hypothesis. I cannot refute it outright. But later I shall argue, first, that it is more peculiar, and therefore less tempting, that it may at first seem; and, second, that we are not forced to accept it, since an alternative hypothesis does justice to the way experience best teaches us what it's like.

Three More Ways to Miss the Point

The hypothesis of phenomenal information characterizes information in terms of eliminated possibilities. But there are other conceptions of 'information'. Therefore the hypothesis has look-alikes: hypotheses which say that experience produces 'information' which could not be gained otherwise, but do not characterize this 'information' in terms of eliminated possibilities. These look-alikes do not work as premises for the knowledge argument. They do not say that phenomenal information eliminates possibilities that differ, but do not differ physically, from uneliminated possibilities. The look-alike hypotheses of phenomenal 'information' are consistent with materialism, and may very well be true. But they don't make the knowledge argument go away. Whatever harmless look-alikes may or may not be true, and whatever conception may or may not deserve the name 'information', the only way to save materialism is fix our attention squarely on the genuine hypothesis of phenomenal information, and deny it. To avert our eyes, and attend to something else, is no substitute for that denial.

Might a look-alike help at least to this extent: by giving us something true that we might have been confused with the genuine hypothesis, thereby explaining how we might have believed the hypothesis although it was false? I think not. Each of the look-alikes turns out to imply not only that experience can give us 'information' that no amount of lessons can give, but also that lessons in Russian can give us 'information' that no amount of lessons in English can give (and vice versa). I doubt that any

friend of phenomenal information ever thought that the special role of experience in teaching what it's like was on a par with the special role of Russian! I will have to say before I'm done that phenomenal information is an illusion, but I think I must look elsewhere for a credible hypothesis about what sort of illusion it might be.

The Fourth Way. If a hidden camera takes photographs of a room, the film ends up bearing traces of what went on in the room. The traces are distinctive: that is, the details of the traces depend on the details of what went on, and if what went on had been different in any of many ways, the traces would have been correspondingly different. So we can say that the traces bear information, and that he who has the film has the information. That might be said because the traces, plus the way they depend on what went on, suffice to eliminate possibilities; but instead we might say 'information' and just mean 'distinctive traces'. If so, it's certainly true that new experience imparts 'information' unlike any that can be gained from lessons. Experience and lessons leave different kinds of traces. That is so whether or not the experience eliminates possibilities that the lessons leave open. It is equally true, of course, that lessons in Russian leave traces unlike any that are left by lessons in English, regardless of whether the lessons cover the same ground and eliminate the same possibilities.

The Fifth Way. When we speak of transmission of 'information', we often mean transmission of text. Repositories of 'information', such as libraries, are storehouses of text. Whether the text is empty verbiage or highly informative is beside the point. Maybe we too contain information by being storehouses of text. Maybe there is a language of thought, and maybe the way we believe things is to store sentences of this language in some special way, or in some special part of our brains. In that case, we could say that storing away a new sentence was storing away a new piece of 'information', whether or not that new piece eliminated any possibilities not already eliminated by the sentences stored previously. Maybe, also, the language of thought is not fixed once and for all, but can gain new words. Maybe, for instance, it borrows words from public language. And maybe, when one has a new experience, that causes one's language of thought to gain a new word which denotes that experience—a word which could not have been added to the language by any other means. If all this is so, then when Mary

sees colors, her language of thought gains new words, allowing her to store away new sentences and thereby gain 'information'. All this about the language of thought, the storing of sentences, and the gaining of words is speculation. But it is plausible speculation, even if no longer the only game in town. If it is all true, then we have another look-alike hypothesis of phenomenal 'information'. When Mary gains new words and stores new sentences, that is 'information' that she never had before, regardless of whether it eliminates any possibilities that she had not eliminated already.

But again, the special role of experience turns out to be on a par with the special role of Russian. If the language of thought picks up new words by borrowing from public language, then lessons in Russian add new words, and result in the storing of new sentences, and thereby impart "information" that never could have been had from lessons in English. (You might say that the new Russian words are mere synonyms of old words, or at least old phrases, that were there already; and synonyms don't count. But no reason has been given why the new inner words created by experience may not also be synonyms of old phrases, perhaps of long descriptions in the language of neurophysiology.)

The Sixth Way. A philosopher who is skeptical about possibility, as so many are, may wish to replace possibilities themselves with linguistic *ersatz* possibilities: maximal consistent sets of sentences. And he may be content to take 'consistent' in a narrowly logical sense, so that a set with 'Fred is married' and 'Fred is a bachelor' may count as consistent, and only an overt contradiction like 'Fred is married' and 'Fred is not married' will be ruled out.¹⁰ The *ersatz* possibilities might also be taken as sets of sentences of the language of thought, if the philosopher believes in it. Then if someone's language of thought gains new words, whether as a result of new experience or as a result of being taught in Russian, the *ersatz* possibilities become richer and more numerous. The sets of sentences that were maximal before are no longer maximal after new words are added. So when Mary sees colors and her language of thought gains new words, there are new *ersatz* possibilities; and she can straightway eliminate some of them. Suppose she knows beforehand that she is about to see green, and that the experience of seeing green is associated with neural firing pattern F. So when she sees green and gains the new word G for her experience, then straightway there are new, enriched *ersatz* possibilities with sentences saying that she has G

without F, and straightway she knows enough to eliminate these *ersatz* possibilities. (Even if she does not know beforehand what she is about to see, straightway she can eliminate at least those of her new-found *ersatz* possibilities with sentences denying that she then has G.) Just as we can characterize information in terms of elimination of possibilities, so we can characterize *ersatz* 'information' in terms of elimination of *ersatz* 'possibilities'. So here we have the closest look-alike hypothesis of all, provided that language-of-thoughtism is true. But we still do not have the genuine hypothesis of phenomenal information, since the eliminated *ersatz* possibility of G without F may not have been a genuine possibility at all. It may have been like the *ersatz* possibility of married bachelors.

Curiouser and Curiouser

The hypothesis of phenomenal information is more peculiar than it may at first seem. For one thing, because it is opposed to more than just materialism. Some of you may have welcomed the knowledge argument because you thought all along that physical information was inadequate to explain the phenomena of mind. You may have been convinced all along that the mind could do things that no physical system could do: bend spoons, invent new jokes, demonstrate the consistency of arithmetic, reduce the wave packet, or what have you. You may have been convinced that the full causal story of how the deeds of mind are accomplished involves the causal interactions not only of material bodies but also of astral bodies; not only the vibrations of the electromagnetic field but also the good or bad vibes of the psionic field; not only protoplasm but ectoplasm. I doubt it, but never mind. It's irrelevant to our topic. The knowledge argument is targeted against you no less than it is against materialism itself.

Let *parapsychology* be the science of all the nonphysical things, properties, causal processes, laws of nature, and so forth that may be required to explain the things we do. Let us suppose that we learn ever so much parapsychology. It will make no difference. Black-and-white Mary may study all the parapsychology as well as all the psychophysics of color vision, but she still won't know what it's like. Lessons on the aura of Vegemite will do no more for us than lessons on its chemical composition. And so it goes. Our intuitive starting point wasn't just that *physics* lessons couldn't help the inexperienced to know what it's like. It was that *lessons* couldn't help. If

there is such a thing as phenomenal information, it isn't just independent of physical information. It's independent of every sort of information that could be served up in lessons for the inexperienced. For it is supposed to eliminate possibilities that any amount of lessons leave open. Therefore phenomenal information is not just parapsychological information, if such there be. It's something very much stranger.

The genuine hypothesis of phenomenal information, as distinguished from its look-alikes, treats information in terms of the elimination of possibilities. When we lack information, several alternative possibilities are open, when we get the information some of the alternatives are excluded. But a second peculiar thing about phenomenal information is that it resists this treatment. (So does logical or mathematical 'information'. However, phenomenal information cannot be logical or mathematical, because lessons in logic and mathematics no more teach us what a new experience is like than lessons in physics or parapsychology do.) When someone doesn't know what it's like to have an experience, where are the alternative open possibilities? I cannot present to myself in thought a range of alternative possibilities about what it might be like to taste Vegemite. That is because I cannot imagine either what it *is* like to taste Vegemite, or any alternative way that it *might* be like but in fact isn't. (I could perfectly well imagine that Vegemite tastes just like peanut butter, or something else familiar to me, but let's suppose I've been told authoritatively that this isn't so.) I can't even pose the question that phenomenal information is supposed to answer: is it this way or that? It seems that the alternative possibilities must be unthinkable beforehand; and afterward too, except for the one that turns out to be actualized. I don't say there's anything altogether impossible about a range of unthinkable alternatives; only something peculiar. But it's peculiar enough to suggest that we may somehow have gone astray.

From Phenomenal to Epiphenomenal

A third peculiar thing about phenomenal information is that it is strangely isolated from all other sorts of information; and this is so regardless of whether the mind works on physical or parapsychological principles. The phenomenal aspect of the world has nothing to do with explaining why people seemingly talk about the phenomenal aspect of the world. For instance, it plays no part in explaining the movements of the pens of phi-

losophers writing treatises about phenomenal information and the way experience has provided them with it.

When Mary gets out of her black-and-white cell, her jaw drops. She says "At last! So this is what it's like to see colors!" Afterward she does things she couldn't do before, such as recognizing a new sample of the first color she ever saw. She may also do other things she didn't do before: unfortunate things, like writing about phenomenal information and the poverty of materialism. One might think she said what she said and did what she did because she came to know what it's like to see colors. Not so, if the hypothesis of phenomenal information is right. For suppose the phenomenal aspect of the world had been otherwise, so that she gained different phenomenal information. Or suppose the phenomenal aspect of the world had been absent altogether, as we materialists think it is. Would that have made the slightest difference to what she did or said then or later? I think not. Making a difference to what she does or says means, at least in part, making a difference to the motions of the particles of which she is composed. (Or better: making a difference to the spatiotemporal shape of the wavefunction of those particles. But let that pass.) For how could she do or say anything different, if none of her particles moved any differently? But if something nonphysical sometimes makes a difference to the motions of physical particles, then physics as we know it is wrong. Not just silent, not just incomplete—wrong. Either the particles are caused to change their motion without benefit of any force, or else there is some extra force that works very differently from the usual four. To believe in the phenomenal aspect of the world, but deny that it is epiphenomenal, is to bet against the truth of physics. Given the success of physics hitherto, and even with due allowance for the foundational ailments of quantum mechanics, such betting is rash! A friend of the phenomenal aspect would be safer to join Jackson in defense of *epiphenomenal* qualia.

But there is more to the case than just an empirical bet in favor of physics. Suppose there is a phenomenal aspect of the world, and suppose it does make some difference to the motions of Mary's jaw or the noises out of her mouth. Then we can describe the phenomenal aspect, if we know enough, in terms of its physical effects. It is that on which physical phenomena depend in such-and-such way. This descriptive handle will enable us to give lessons on it to the inexperienced. But insofar as we can give lessons on it, what we have is just parapsychology. That whereof we cannot learn except

by having the experience still eludes us. I do not argue that *everything* about the alleged distinctive subject matter of phenomenal information must be epiphenomenal. Part of it may be parapsychological instead. But I insist that *some* aspect of it must be epiphenomenal.

Suppose that the hypothesis of phenomenal information is true and suppose that V_1 and V_2 are all of the maximally specific phenomenal possibilities concerning what it's like to taste Vegemite; anyone who tastes Vegemite will find out which one obtains, and no one else can. And suppose that P_1 and P_2 are all the maximally specific physical possibilities. (Of course we really need far more than two Ps, and maybe a friend of phenomenal information would want more than two Vs, but absurdly small numbers will do for an example.) Then we have four alternative hypotheses about the causal independence or dependence of the Ps on the Vs. Each one can be expressed as a pair of counterfactual conditionals. Two hypotheses are patterns of dependence.

K_1 : if V_1 then P_1 , if V_2 then P_2

K_2 : if V_1 then P_2 , if V_2 then P_1

The other two are patterns of independence.

K_3 : if V_1 then P_1 , if V_2 then P_1

K_4 : if V_1 then P_2 , if V_2 then P_2

These dependency hypotheses are, I take it, contingent propositions. They are made true, if they are, by some contingent feature of the world, though it's indeed a vexed question what sort of feature it is.¹¹ Now we have eight joint possibilities.

$K_1V_1P_1$ $K_3V_1P_1$ $K_3V_2P_1$ $K_2V_2P_1$

$K_2V_1P_2$ $K_4V_1P_2$ $K_4V_2P_2$ $K_1V_2P_2$

Between the four on the top row and the four on the bottom row, there is the physical difference between P_1 and P_2 . Between the four on the left and the four on the right, there is the phenomenal difference between V_1 and V_2 . And between the four on the edges and the four in the middle there is a parapsychological difference. It is the difference between dependence and independence of the physical on the phenomenal; between efficacy and epiphenomenalism, so far as this one example is concerned. There's nothing ineffable about that. Whether or not you've tasted Vegemite, and whether or not you can conceive of the alleged difference between V_1 and V_2 , you

can still be told whether the physical difference between P_1 and P_2 does or doesn't depend on some part of the phenomenal aspect of the world.

Lessons can teach the inexperienced which parapsychological possibility obtains, dependence or independence. Let it be dependence: we have either K_1 or K_2 . For if we had independence, then already we would have found our epiphenomenal difference: namely, the difference between V_1 and V_2 . And lessons can teach the inexperienced which of the two physical possibilities obtains. Without loss of generality let it be P_1 . Now two of our original eight joint possibilities remain open: $K_1V_1P_1$ and $K_2V_2P_1$. The difference between those is not at all physical, and not at all parapsychological: it's P_1 , and it's dependence, in both cases. The difference is entirely phenomenal. And also it is entirely epiphenomenal. Nothing physical, and nothing parapsychological, depends on the difference between $K_1V_1P_1$ and $K_2V_2P_1$. We have the same sort of pattern of dependence either way; it's just that the phenomenal possibilities have been swapped. Whether it's independence or whether it's dependence, therefore, we have found an epiphenomenal part of the phenomenal aspect of the world. It is the residue left behind when we remove the parapsychological part.

Suppose that someday I taste Vegemite, and hold forth about how I know at last what it's like. The sound of my holding forth is a physical effect, part of the realized physical possibility P_1 . This physical effect is exactly the same whether it's part of the joint possibility $K_1V_1P_1$ or part of its alternative $K_2V_2P_1$. It may be caused by V_1 in accordance with K_1 , or it may instead be caused by V_2 in accordance with K_2 , but it's the same either way. So it does not occur because we have K_1V_1 rather than K_2V_2 or vice versa. The alleged difference between these two possibilities does nothing to explain the alleged physical manifestation of my finding out which one of them is realized. It is in that way that the difference is epiphenomenal. That makes it very queer, and repugnant to good sense.

The Ability Hypothesis

So the hypothesis of phenomenal information turns out to be very peculiar indeed. It would be nice, and not only for materialists, if we could reject it. For materialists, it is essential to reject it. And we can. There is an alternative hypothesis about what it is to learn what an experience is like: the *ability hypothesis*. Laurence Nemirow summarizes it thus:

some modes of understanding consist, not in the grasping of facts, but in the acquisition of abilities. . . . As for understanding an experience, we may construe that as an ability to place oneself, at will, in a state representative of the experience. I understand the experience of seeing red if I can at will visualize red. Now it is perfectly clear why there must be a special connection between the ability to place oneself in a state representative of a given experience and the point of view of experienter: exercising the ability just *is* what we call 'adopting the point of view of experienter'. . . . We can, then, come to terms with the subjectivity of our understanding of experience without positing subjective facts as the objects of our understanding. This account explains, incidentally, the linguistic incommunicability of our subjective understanding of experience (a phenomenon which might seem to support the hypothesis of subjective facts). The latter is explained as a special case of the linguistic incommunicability of abilities to place oneself at will in a given state, such as the state of having lowered blood pressure, and the state of having wiggling ears.¹²

If you have a new experience, you gain abilities to remember and to imagine. After you taste Vegemite, and you learn what it's like, you can afterward remember the experience you had. By remembering how it once was, you can afterward imagine such an experience. Indeed, even if you eventually forget the occasion itself, you will very likely retain your ability to imagine such an experience.

Further, you gain an ability to recognize the same experience if it comes again. If you taste Vegemite on another day, you will probably know that you have met the taste once before. And if, while tasting Vegemite, you know that it is Vegemite you are tasting, then you will be able to put the name to the experience if you have it again. Or if you are told nothing at the time, but later you somehow know that it is Vegemite that you are then remembering or imagining tasting, again you can put the name to the experience, or to the memory, or to the experience of imagining, if it comes again. Here, the ability you gain is an ability to gain information if given other information. Nevertheless, the information gained is not phenomenal, and the ability to gain information is not the same thing as information itself.

Earlier, I mentioned 'knowing what an experience is like under a description'. Now I can say that what I meant by this was having the ability to remember or imagine an experience while also knowing the egocentric proposition that what one is then imagining is the experience of such-and-such description. One might well know what an experience is like under one description, but not under another. One might even know what some

experience is like, but not under any description whatever—unless it be some rather trivial description like 'that queer taste that I'm imagining right now'. That is what would happen if you slipped a dab of Vegemite into my food without telling me what it was: afterward, I would know what it was like to taste Vegemite, but not under that description, and not under any other nontrivial description. It might be suggested that 'knowing what it's like to taste Vegemite' really means what I'd call 'knowing what it's like to taste Vegemite under the description 'tasting Vegemite''; and if so, knowing what it's like would involve both ability and information. I disagree. For surely it would make sense to say: "I know this experience well, I've long known what it's like, but only today have I found out that it's the experience of tasting Vegemite." But this verbal question is unimportant. For the information involved in knowing what it's like under a description, and allegedly involved in knowing what it's like, is anyhow not the queer phenomenal information that needs rejecting.

(Is there a problem here for the friend of phenomenal information? Suppose he says that knowing what it's like to taste Vegemite means knowing that the taste of Vegemite has a certain 'phenomenal character'. This requires putting the name to the taste, so clearly it corresponds to our notion of knowing what it's like to taste Vegemite under the description 'tasting Vegemite'. But we also have our notion of knowing what it's like *simpliciter*, and what can he offer that corresponds to that? Perhaps he should answer by appeal to a trivial description, as follows: knowing what it's like *simpliciter* means knowing what it's like under the trivial description 'taste I'm imagining now', and that means knowing that the taste one is imagining now has a certain phenomenal character.)

As well as gaining the ability to remember and imagine the experience you had, you also gain the ability to imagine related experiences that you never had. After tasting Vegemite, you might for instance become able to imagine tasting Vegemite ice cream. By performing imaginative experiments, you can predict with some confidence what you would do in circumstances that have never arisen—whether you'd ask for a second helping of Vegemite ice cream, for example.

These abilities to remember and imagine and recognize are abilities you cannot gain (unless by super-neurosurgery, or by magic) except by tasting Vegemite and learning what it's like. You can't get them by taking lessons on the physics or the parapsychology of the experience, or even by taking

comprehensive lessons that cover the whole of physics and parapsychology. The ability hypothesis says that knowing what an experience is like just *is* the possession of these abilities to remember, imagine, and recognize. It isn't the possession of any kind of information, ordinary or peculiar. It isn't knowing that certain possibilities aren't actualized. It isn't knowing-that. It's knowing-how. Therefore it should be no surprise that lessons won't teach you what an experience is like. Lessons impart information; ability is something else. Knowledge-that does not automatically provide know-how.

There are parallel cases. Some know how to wiggle their ears; others don't. If you can't do it, no amount of information will help. Some know how to eat with chopsticks, others don't. Information will help up to a point—for instance, if your trouble is that you hold one chopstick in each hand—but no amount of information, by itself, will bring you to a very high level of know-how. Some know how to recognize a C-38 locomotive by sight, others don't. If you don't, it won't much help if you memorize a detailed geometrical description of its shape, even though that does all the eliminating of possibilities that there is to be done. (Conversely, knowing the shape by sight doesn't enable you to write down the geometrical description.) Information very often contributes to know-how, but often it doesn't contribute enough. That's why music students have to practice.

Know-how is ability. But of course some aspects of ability are in no sense knowledge: strength, sufficient funds. Other aspects of ability are, purely and simply, a matter of information. If you want to know how to open the combination lock on the bank vault, information is all you need. It remains that there are aspects of ability that do *not* consist simply of possession of information, and that we *do* call knowledge. The ability hypothesis holds that knowing what an experience is like is that sort of knowledge.

If the ability hypothesis is the correct analysis of knowing what an experience is like, then phenomenal information is an illusion. We ought to explain that illusion. It would be feeble, I think, just to say that we're fooled by the ambiguity of the word 'know': we confuse ability with information because we confuse knowledge in the sense of knowing-how with knowledge in the sense of knowing-that. There may be two senses of the word 'know', but they are well and truly entangled. They mark the two pure endpoints of a range of mixed cases. The usual thing is that we gain information and ability together. If so, it should be no surprise if we apply

to pure cases of gaining ability, or to pure cases of gaining information, the same word 'know' that we apply to all the mixed cases.

Along with information and ability, acquaintance is a third element of the mixture. If Lloyd George died too soon, there's a sense in which Father never can know him. Information won't do it, even if Father is a most thorough biographer and the archives are very complete. (And the trouble isn't that there's some very special information about someone that you can only get by being in his presence.) Know-how won't do it either, no matter how good Father may be at imagining Lloyd George, seemingly remembering him, and recognizing him. (Father may be able to recognize Lloyd George even if there's no longer any Lloyd George to recognize—if *per impossibile* he did turn up, Father could tell it was him.) Again, what we have is not just a third separate sense of 'know'. Meeting someone, gaining a lot of information about him that would be hard to gain otherwise, and gaining abilities regarding him usually go together. The pure cases are exceptions.

A friend of phenomenal information will agree, of course, that when we learn what an experience is like, we gain abilities to remember, imagine, and recognize. But he will say that it is because we gain phenomenal information that we gain the abilities. He might even say the same about other cases of gaining know-how: you can recognize the C-38 when you have phenomenal information about what it's like to see that shape, you can eat with chopsticks or wiggle your ears when you gain phenomenal information about the experience of doing so, and so on. What should friends of the ability hypothesis make of this? Is he offering a conjecture, which we must reject, about the causal origin of abilities? I think not. He thinks, as we do, that experiences leave distinctive traces in people, and that these traces enable us to do things. Likewise being taught to recognize a C-38 or to eat with chopsticks, or whatever happens on first wiggling the ears, leave traces that enable us to do things afterward. That much is common ground. He also interprets these enabling traces as representations that bear information about their causes. (If the same traces had been caused in some deviant way they might perhaps have carried misinformation.) We might even be able to accept that too. The time for us to quarrel comes only when he says that these traces represent special phenomenal facts, facts which cannot be represented in any other way, and therefore which cannot be taught in physics lessons or even in parapsychology

lessons. That is the part, and the *only* part, which we must reject. But that is no part of his psychological story about how we gain abilities. It is just a gratuitous metaphysical gloss on that story.

We say that learning what an experience is like means gaining certain abilities. If the causal basis for those abilities turns out also to be a special kind of representation of some sort of information, so be it. We need only deny that it represents a special kind of information about a special subject matter. Apart from that it's up for grabs what, if anything, it may represent. The details of stimuli: the chemical composition of Vegemite, reflectances of surfaces, the motions of well-handled chopsticks or of ears? The details of inner states produced by those stimuli: patterns of firings of nerves? We could agree to either, so long as we did not confuse 'having information' represented in this special way with having the same information in the form of knowledge or belief. Or we could disagree. Treating the ability-conferring trace as a representation is optional. What's essential is that when we learn what an experience is like by having it, we gain abilities to remember, imagine, and recognize.

Acknowledgment

Part of this paper derives from a lecture at LaTrobe University in 1981. I thank LaTrobe for support in 1981, Harvard University for support under a Santayana Fellowship in 1988, and Frank Jackson for very helpful discussion.

Notes

1. See Frank Jackson, "Epiphenomenal qualia," *Philosophical Quarterly* 32 (1982), pp. 127–136, and reprinted in this volume; "What Mary didn't know," *Journal of Philosophy* 83 (1986), pp. 291–295.
2. See B. A. Farrell, "Experience," *Mind* 59 (1950), pp. 170–198; and Thomas Nagel, "What is it like to be a bat?" *Philosophical Review* 83 (1974), pp. 435–450, also in Thomas Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979).
3. See Peter Unger, "On experience and the development of the understanding," *American Philosophical Quarterly* 3 (1966), pp. 1–9.
4. For such speculation, see Paul M. Churchland, "Reduction, qualia, and the direct introspection of brain states," *Journal of Philosophy* 82 (1985), pp. 8–28.

5. See my "Attitudes *de dicto* and *de se*," *Philosophical Review* 88 (1979), pp. 513–543, also in my *Philosophical Papers*, vol. 1 (New York: Oxford University Press, 1983); and Roderick Chisholm, *The First Person: An Essay on Reference and Intentionality* (Minneapolis: University of Minnesota Press, 1981).
 6. See Gilbert Harman, "The intrinsic quality of experience," *Philosophical Perspectives* 4 (1990).
 7. See Ned Block and Jerry A. Fodor, "What psychological states are not," *Philosophical Review* 81 (1972), pp. 159–181, also in Ned Block (ed.), *Readings in Philosophy of Psychology*, vol. 1 (Cambridge, Mass.: Harvard University Press, 1980); and my "Mad pain and Martian pain," in *Readings in Philosophy of Psychology*, vol. 1, and in my *Philosophical Papers*, vol. 1.
 8. See my "New work for a theory of universals," *Australasian Journal of Philosophy* 61 (1983), pp. 343–377, especially pp. 361–364. For a different view about how to state minimal materialism, see Terence Horgan, "Supervenience and microphysics," *Pacific Philosophical Quarterly* 63 (1982), pp. 29–43.
 9. On recombination of possibilities, see my *On the Plurality of Worlds* (Oxford: Blackwell Publishers, 1986), pp. 87–92. The present argument may call for a principle that also allows recombination of properties; I now think that would not necessarily require treating properties as nonspatiotemporal parts of their instances. On recombination of properties, see also D. M. Armstrong, *A Combinatorial Theory of Possibility* (Cambridge: Cambridge University Press, 1989).
 10. See *On the Plurality of Worlds*, pp. 142–165, on linguistic *ersatz* possibilities.
 11. On dependency hypotheses, see my "Causal decision theory," *Australasian Journal of Philosophy* 59 (1981), pp. 5–30, reprinted in my *Philosophical Papers*, vol. 11 (New York: Oxford University Press, 1986).
 12. Laurence Nemirow, review of Nagel's *Mortal Questions*, *Philosophical Review* 89 (1980), pp. 475–476. For a fuller statement, see Nemirow, "Physicalism and the cognitive role of acquaintance"; and *Functionalism and the Subjective Quality of Experience* (doctoral dissertation, Stanford, 1979). See also Michael Tye, "The subjective qualities of experience," *Mind* 95 (1986), pp. 1–17.
- I should record a disagreement with Nemirow on one very small point. We agree that the phrase 'what experience E is like' does not denote some 'subjective quality' of E, something which supposedly would be part of the subject matter of the phenomenal information gained by having E. But whereas I have taken the phrase to denote E itself, Nemirow takes it to be a syncategorematic part of the expression 'know what experience E is like'. See "Physicalism and the cognitive role of acquaintance," section III.