Complexity does matter. If it didn't, there would be a much shorter argument against strong AI: "Hey, look at this hand calculator. It doesn't understand Chinese, and any conceivable computer is just a giant hand calculator, so no computer could understand Chinese. Q.E.D." When we factor in the complexity, as we must, we really have to factor it in — and not just pretend to factor it in. That is hard to do, but until we do, any intuitions we have about what is "obviously" not present are not to be trusted. Like Frank Jackson's case of Mary the color scientist, Searle's thought experiment yields a strong, clear conviction only when we fail to follow instructions. These intuition pumps are defective; they do not enhance but mislead our imaginations.

But what, then, of my own intuition pumps? What of Shakey the robot, or the CADBLIND Mark II, or the biofeedback-trained blindsight patient, for instance? Are they not equally suspect, equally guilty of misleading the reader? I've certainly done my best in telling these tales to lead your imagination down certain paths, and to keep you from bogging down in complexities I deemed unnecessary to the point I was attempting to make. There is some asymmetry, however: My intuition pumps are, for the most part, intended to help you imagine new possibilities, not convince you that certain prospects are impossible. There are exceptions. My variation on the brain in the vat that opened the book was designed to impress on you the impossibility of certain sorts of deception, and some of the thought experiments in chapter 5 were intended to show that, unless there were a Cartesian Theater, there *could not* be a fact of the matter distinguishing Orwellian from Stalinesque content revisions. These thought experiments proceeded, however, by heightening the vividness for the "opposition"; the examples of the woman in the hat at the party and the long-haired woman with glasses, for instance, were designed to sharpen the very intuition I then sought to discredit by argument.

Still, let the reader beware: My intuition pumps, like anyone else's, are not the straightforward demonstrations they may seem to be; they are more art than science. (For further warnings about philosophers' thought experiments, see Wilkes, 1988.) If they help us conceive of new possibilities, which we can then confirm by more systematic methods, that is an achievement; if they lure us down the primrose path, that is a pity. Even good tools can be misused, and like any other workers, we will do better if we understand how our tools work.

2. WHAT IT IS LIKE TO BE A BAT

The most widely cited and influential thought experiment about consciousness is Thomas Nagel's "What Is It Like to Be a Bat?" (1974). He answers his title question by claiming that this is impossible for us to imagine. This claim is congenial to many, apparently; one sometimes sees his paper cited by scientists as if it were that rarity of rarities, a philosophical "result" — a received demonstration of a fact that any theory must subsequently accommodate.

Nagel chose his target creatures well. Bats, as fellow mammals, are enough like us to support the conviction that *of course* they are conscious. (If he had written "What Is It Like to Be a Spider?" many would be inclined to wonder what made him so sure it was like anything at all.) But thanks to their system of echolocation — bats can "see with their ears" — they are also different enough from us so that we can sense the vast gulf. Had he written a paper called "What Is It Like to Be a Chimpanzee?" or, more to the point, "What Is It Like to Be a Cat?" the opinion that his pessimistic conclusion was obvious would not be so close to unanimity. There are many people who are supremely confident that they know *just* what it's like to be a cat. (They are wrong, of course, unless they have supplemented all their loving and empathetic observation with vast amounts of physiological research, but they would be erring on the wrong side, from Nagel's point of view.)

For better or worse, most people seem quite cheerful about accepting Nagel's "result" regarding the inaccessibility to us of bat consciousness. Some philosophers have challenged it, however, and for good reason (Hofstadter, 1981; Hardin, 1988; Leiber, 1988; Akins, 1990). First we must be clear about just which result it is. It is not just the epistemological or evidential claim that even if someone succeeded ("by accident") in imagining what it is like to be a bat, we would never be able to confirm that this successful feat of imagination had occurred. It is rather that we human beings don't have and could never acquire the wherewithal, the representational machinery, to represent to ourselves what it is like to be a bat.

The distinction is important. In chapter 12 we looked at the similar feat of imagining what it must have been like to be a Leipziger hearing one of Bach's cantatas for the first time. The epistemological problem is difficult, but straightforwardly addressable by the usual sorts of research. Figuring out just what sorts of experiences they would have had, and how these would differ from our experiences of Bach, is a

matter of historical, cultural, psychological, and, maybe, physiological investigation. We can figure out some of this quite readily, including some of the most striking differences from our own experience, but if we were to try to put ourselves into the very sequence of experiential states such a person would enjoy, we would face diminishing returns. The task would require us to subject ourselves to vast transformations — forgetting much of what we know, losing associations and habits, acquiring new habits and associations. We can use our "third-person" research to say what these transformations would be, but actually undergoing them would involve terrible costs of isolation from our contemporary culture — no listening to the radio, no reading about post-Bach political and social developments, and so forth. There is no need to go to those lengths to learn about Leipziger consciousness.

The same is true about imagining what it is like to be a bat. We should be interested in what we can know about the bat's consciousness (if any), not whether we can turn our minds temporarily or permanently into bat minds. In chapter 12, we undermined the presumption that there were "intrinsic" properties — qualia — that constitute *what it is like* to have one conscious experience or another, and as Akins (1990) points out, even if there *were* residual nondispositional, nonrelational properties of bat experiences, becoming intimately acquainted with them, while remaining ignorant of the researchable facts about the systematic structure of bat perception and behavior, would leave us ignorant of what it is like to be a bat. There is at least a lot that we can know about what it is like to be a bat, and neither Nagel nor anyone else has given us a good reason to believe there is anything interesting or theoretically important that is inaccessible to us.

Nagel claims that no amount of third-person knowledge could tell us what it is like to be a bat, and I flatly deny that claim. How might we resolve this dispute? By engaging in something that starts out as child's play — a game in which one person imagines what it is like to be x, and the other then tries to demonstrate that there is something wrong with that particular exercise of heterophenomenology.

Here are some simple warmup exercises:

A: Here's Pooh the teddy bear, thinking how nice it would be to have some honey for breakfast!

B: Wrong. The teddy bear has no provision for distinguishing honey from anything else. No operating sense organs, and not even a stomach. The teddy bear is filled with inert stuffing. It is not like anything to be a teddy bear.

A. Here's Bambi the deer, admiring the beautiful sunset, until the bright orange sky suddenly reminds him of the evil hunter's jacket!

B. Wrong. Deer are colorblind (well, they may have some sort of dichromatic vision). Whatever deer are conscious of (if anything), they don't distinguish colors such as orange.

A. Here's Billy the bat perceiving, in his special sonar sort of way, that the flying thing swooping down toward him was not his cousin Bob, but an eagle, with pinfeathers spread and talons poised for the kill!

B. Hang on — how far away did you say the eagle was? A bat's echolocation is only good for a few meters.

A. Um, well . . . And the eagle was already only two meters away!

B. Ah, now this is harder to say. Just what are the resolution limits of a bat's echolocation? Is it used to identify objects at all, or just as an alerter and tracker for capture? Would a bat be able to distinguish pinfeathers spread from pinfeathers closed just using echolocation? I doubt it, but we will have to design some experiments to see, and also, of course, some experiments to discover whether bats are capable of keeping track of, and reidentifying, their kin. Some mammals can, and others, we have good reason to believe, are utterly oblivious of such matters.

The sorts of investigation suggested by this exercise would take us a long way into an account of the structure of the bat's perceptual and behavioral world, so we could rank order heterophenomenological narratives for realism, discarding those that asserted or presupposed discriminatory talents, or reactive dispositions, demonstrably not provided for in the ecology and neurophysiology of the bat. For example, we would learn that bats would not be bothered by the loud squeaks they emit in order to produce their echoes, because they have a cleverly designed muscle that shuts down their ears in perfect timing with their squeaks, not unlike the timing devices that permit sensitive radar systems to avoid being blasted by their own outgoing signals. A lot of research has already been done on these issues, so we can already say much more, for instance about why bats use different frequency patterns for their squeaks, depending on whether they are scanning for prey, approaching a target, or homing in for the kill (Akins, 1989, 1990).

When we arrive at heterophenomenological narratives that no critic can find any positive grounds for rejecting, we should accept

them — tentatively, pending further discoveries — as accurate accounts of what it is like to be the creature in question. That, after all, is how we treat each other. In recommending that we treat bats and other candidates for interpretation the same way, I am not *shifting* the burden of proof but extending the normal, human, burden of proof to other entities.

We could use these investigations to dispel all sorts of overly romantic illusions about bat consciousness. We *know* that Randall Jarrell's delightful children's book, *The Bat-Poet* (1963), is fantasy, because we know that bats don't talk! Less obviously fantastical claims about their phenomenology succumb to less obvious, but still public, facts about their physiology and behavior. These investigations would show us a great deal about what a bat could and could not be conscious of under various conditions, by showing us what provisions there were in their nervous systems for representing this and that, and by checking experimentally to make sure the bat actually put the information to use in the modulation of its behavior. It is hard to imagine how much can be gleaned from this sort of research until you actually look into it. (For a surprisingly detailed preliminary investigation of what it is like to be a vervet monkey, for instance, see Cheney and Seyfarth, *How Monkeys See the World*, 1990.)

This invites an obvious objection: These investigations would show us a great deal about brain organization and information-processing in the bat, but they would show us only what bats are not conscious of, leaving entirely open what, if anything, bats *are* conscious of. As we know, much of the information-processing in nervous systems is entirely unconscious, so these methods of investigation will do nothing to rule out the hypothesis that bats are . . . flying zombies, creatures it is not like anything to be! (Wilkes, 1988, p. 224, wonders whether bat echolocation is a sort of blindsight, not like anything at all.)

Ah, the bat is out of the bag. This is indeed the ominous direction in which this discussion seems to be sliding, and we must head it off. Richard Dawkins (1986), in an illuminating discussion of the design of echolocation in horseshoe bats, gives us a clear version of the image that is lurking.

> The Doppler Effect is used in police radar speed-traps for motorists. . . . By comparing the outgoing frequency with the frequency of the returning echo the police, *or rather their automatic instrument* [my emphasis], can calculate the speed of each car. . . . By comparing the pitch of its cry with the pitch of the returning echo, therefore, the bat *(or rather its on-board computer in the brain)* [my emphasis] could, in theory, calculate how fast it was moving towards the tree. [pp. 30–31]

It is tempting to ask: Is there something in the bat that is situated relative to *its* "onboard computer" (which operates without a smidgen of consciousness) as the police are situated relative to their "automatic device"? The police don't have to calculate the Doppler shift consciously, but they do have to experience, consciously, the readout on their device that says, in bright red LED symbols: "75MPH." That is their cue for leaping on their motorcycles and starting up their sirens. We may plausibly suppose that the bat also does not consciously calculate the Doppler shift — its onboard computer takes care of that — but then isn't there a role left over, in the bat, for something like the experiencing cop, a witness to appreciate (consciously) the "output" of the bat's Doppler-effect-analysis computer? Note that we could easily enough replace the police officers with an automatic device that somehow recorded the registration number of the offending vehicle, looked up the operator's name and address and sent him or her a ticket. There is nothing special about the task the police are doing that shows it could not be done without any experiencing of anything. The same holds, it would seem, for the bat. A bat might be a zombie. It would be a zombie — so this line of reasoning suggests — unless there were an inner observer in it that reacts to an inner presentation in much the way the officers react to the flashing red lights on their instruments.

Don't fall in the trap. This is our old nemesis, the Audience in the Cartesian Theater. *Your* consciousness does not consist in the fact that your brain is inhabited by an inner agent to whom your brain presents displays, so our inability to find such a central agent in the bat's brain would not jeopardize its claim to consciousness, or our claim to be able to say what its consciousness was like. In order to understand a bat's consciousness, we must simply apply the same principles to the bat that we apply to ourselves.

But what could a bat do, then, that would be special enough to convince us that we were in the presence of genuine consciousness? It may seem that no matter what fancy output-users we situate behind the bat's Doppler-transducer, there could be no convincing, from-the-outside, "third-person" reason to grant the bat conscious experience. Not so. If the bat could talk, for instance, it would generate a text from which we could generate a heterophenomenological world, and that

would give us exactly the same grounds for granting it consciousness that serve for any person. But, as we just noted, bats can't talk. They can, however, behave in many nonverbal ways that can provide a clear basis for describing their heterophenomenological world, or, as the pioneer researcher von Uexküll (1909) called it, their *Umwelt* und *Innenwelt*, their Surroundworld and Innerworld.

Heterophenomenology without a text is not impossible, just difficult (Dennett, 1988a, 1988b, 1989a, 1989b). One branch of animal heterophenomenology is known as cognitive ethology, the attempt to model animals' minds by studying — and experimenting on — their behavior in the field. The possibilities and difficulties of this sort of investigation are well represented in Cheney and Seyfarth (1990), Whiten and Byrne (1988), and in Ristau (1991), a festschrift dedicated to Donald Griffin, the pioneer investigator of bat echolocation and the creator of the field of cognitive ethology. One of the frustrating difficulties encountered by these investigators is that many of the experiments one dreams of running turn out to be utterly impractical in the absence of language; one simply cannot *set* up subjects (and know that one has set them up) in the ways these experiments would require without conversing with the subjects (Dennett, 1988a).

This is not just an epistemological problem for the heterophenomenologist; the very difficulty of creating the requisite experimental circumstances in the natural environment demonstrates something more fundamental about the minds of languageless creatures. It shows that the ecological situations of these animals have never provided them with *opportunities* for the development (by evolution, by learning, or by both) of many of the advanced mental activities that shape our minds, and so we can be quite sure they have never developed them. For instance, consider the concept of a *secret*. A secret is not just something you know that others don't know. For you to have a secret you need to know that the others don't know it, and you have to be able to control that fact. (If you are the first to see the approaching stampede, you may know something the others don't know, but not for long; you can't keep this bit of privileged information secret.) The behavioral ecology of a species has to be rather specially structured for there to be any role for secrets at all. Antelopes, in their herds, have no secrets and no way of getting any. So an antelope is probably no more capable of hatching a secret plan than it is capable of counting to a hundred or enjoying the colors of a sunset. Bats, who engage in relatively solitary forays during which they might be able to recognize that very isolation from their rivals, meet one of the necessary conditions for having secrets. Do they

also have interests that might be noticeably well served by exploiting secrets? (What could a clam do with a secret? Just sit there in the mud, chuckling to itself?) Do bats also have habits of stealth or deception in hunting that might be adapted for more elaborate secret-keeping activity? There are in fact many questions of this sort that, once raised, suggest further investigations and experiments. The structure of a bat's mind is just as accessible as the structure of a bat's digestive system; the way to investigate either one is to go back and forth systematically between an assay of its contents and an assay of the world from which its contents were derived, paying attention to the methods and goals of the derivation.

Wittgenstein once said, "If a lion could talk, we could not understand him" (1958, p. 223). I think, on the contrary, that if a lion could talk, that lion would have a mind so different from the general run of lion minds, that although we could understand him just fine, we would learn little about ordinary lions from him. Language, as we saw in earlier chapters, plays an enormous role in the structuring of a human mind, and the mind of a creature lacking language — and having really no need for language — should not be supposed to be structured in these ways. Does this mean that languageless animals "are not conscious at all" (as Descartes insisted)? This question always arises at this moment as a sort of incredulous challenge, but we shouldn't feel obliged to answer it as it stands. Notice that it presupposes something we have worked hard to escape: the assumption that consciousness is a special all-or-nothing property that sunders the universe into two vastly different categories: the things that have it (the things that it is like something to be, as Nagel would put it) and the things that lack it. Even in our own case, we cannot draw the line separating our conscious mental states from our unconscious mental states. The theory of consciousness we have sketched allows for many variations of functional architecture, and while the presence of language marks a particularly dramatic increase in imaginative range, versatility, and self-control (to mention a few of the more obvious powers of the Joycean virtual machine), these powers do not have the *further* power of turning on some special inner light that would otherwise be off.

When we imagine what it is like to be a languageless creature, we start, naturally, from our own experience, and most of what then springs to mind has to be adjusted (mainly downward). The sort of consciousness such animals enjoy is dramatically truncated, compared to ours. A bat, for instance, not only can't wonder whether it's Friday; it can't even wonder whether it's a bat; there is no role for wondering to play

in its cognitive structure. While a bat, like even the lowly lobster, has a biological self, it has no selfy self to speak of — no Center of Narrative Gravity, or at most a negligible one. No words-on-the-tip-of-its-tongue, but also no regrets, no complex yearnings, no nostalgic reminiscences, no grand schemes, no reflections on what it is like to be a cat, or even on what it is like to be a bat. This list of dismissals would be cheap skepticism if we didn't have a positive empirical theory on which to base it. Am I claiming to have proven that bats *could not* have these mental states? Well, no, but I also can't prove that mushrooms *could not* be intergalactic spaceships spying on us.

Isn't this an awfully anthropocentric prejudice? Besides, what about deaf-mutes? Aren't they conscious? Of course they are — but let's not jump to extravagant conclusions about their consciousness, out of misguided sympathy. When a deaf-mute acquires language (in particular, Sign language, the most natural language a deaf-mute can learn), a full-fledged human mind is born, clearly different in discoverable ways from the mind of a hearing person, but capable of all the reflective intricacy and generative power — perhaps more. But without a natural language, a deaf-mute's mind is terribly stunted. (See Sacks, 1989, especially the annotated bibliography.) As the philosopher Ian Hacking (1990) notes in a review of Sacks's book, "It takes a vivid imagination even to have a sense of what a deaf child is missing." One does not do deaf-mutes a favor by imagining that in the absence of language they enjoy all the mental delights we hearing human beings enjoy, and one does not do a favor to nonhuman animals by trying to obscure the available facts about the limitations of their minds.

And this, as many of you are aching to point out, is a subtext that has been struggling to get to the surface for quite a while: Many people are afraid to see consciousness explained because they fear that if we succeed in explaining it, we will lose our moral bearings. Maybe we *can* imagine a conscious computer (or the consciousness of a bat) but we *shouldn't try*, they think. If we get into that bad habit, we will start treating animals as if they were wind-up toys, babies and deaf-mutes as if they were teddy bears, and — just to add insult to injury — robots as if they were real people.

3. MINDING AND MATTERING

I take the title of this section from an article by Marian Stamp Dawkins (1987), who has done careful investigations of the moral implications of animal heterophenomenology. (Her early work is reported in her book *Animal Suffering: The Science of Animal Welfare*, 1980.) As she notes, our moral attitudes towards other animals are full of inconsistencies.

> We have only to think of various different sorts of animals to show up our inconsistencies. There are demonstrations against killing baby harp seals, but there are no comparable campaigns to stop the killing of rats. Many people are quite happy to eat pigs or sheep but horrified by the idea of eating dogs or horses. [p. 150]

Dawkins points out that there are two main strands to this tangle: the ability to reason and the ability to suffer. Descartes made much of the inability of nonhuman animals to reason (at least the way human beings reason), which provoked a famous response from the British utilitarian philosopher Jeremy Bentham: "a full-grown horse or dog is beyond comparison a more rational, as well as a more conversible animal than an infant of a day or a week, or even a month old. But suppose they were otherwise, what would it avail? The question is not, Can they *reason*? nor, Can they *talk*, but Can they *suffer*?" (Bentham, 1789) These usually appear to be opposing benchmarks of moral standing, but as Dawkins argues, "giving ethical value to the ability to suffer will in the end lead us to value animals that are clever. Even if we start out by rejecting Descartes' reasoning criterion, it is the reasoning animals that are the ones most likely to possess the capacity to suffer" (p. 153).

The reasons for this are implicit in the theory of consciousness we have developed. Suffering is not a matter of being visited by some ineffable but intrinsically awful state, but of having one's life hopes, life plans, life projects blighted by circumstances imposed on one's desires, thwarting one's intentions — whatever they are. The idea of suffering being somehow explicable as the presence of some intrinsic property — horribility, let's say — is as hopeless as the idea of amusement being somehow explicable as the presence of intrinsic hilarity. So the presumed inaccessibility, the ultimate unknowability, of another's suffering is just as misleading as the other fantasies about intrinsic qualia we have unmasked, though more obviously pernicious. It follows — and this does strike an intuitive chord — that the capacity to suffer is a function of the capacity to have articulated, wide-ranging, highly discriminative desires, expectations, and other sophisticated mental states.

Human beings are not the only creatues smart enough to suffer; Bentham's horse and dog show by their behavior that they have enough