Chapter Six

Do You Have Constant Tactile Experience of Your Feet in Your Shoes?

For example, up to this moment I have been focusing my attention on the philosophical problem of describing consciousness, and I have not been paying any attention to the feeling of the chair against my back, the tightness of my shoes, or the slight headache I have from drinking too much wine last night. Nonetheless, all of these phenomena are part of my conscious awareness.

– John Searle, *The Rediscovery of the Mind* (1992, p. 137-138)

We are constantly reacting to things without being conscious of them at the time. Sitting against a tree, I am always reacting to the tree and to the ground and to my own posture, since if I wish to walk, I will quite unconsciously stand up from the ground to do so. Immersed in the ideas of this first chapter, I am rarely conscious even of where I am.

– Julian Jaynes, *The Origin of Consciousness in the Breakdown of the Bicameral Mind* (1976), p. 22

i.

Do you have constant tactile experience of your feet in your shoes? Constant auditory experience of the hum of traffic in the background? Constant visual experience of the frame of your eyeglasses? Or, when you aren't attending to such matters, do they drop out of consciousness, so that they're in no way part of your stream of experience, your

phenomenology? Is consciousness *abundant,* the stream of experience bristling with phenomenology in a wide variety of modalities simultaneously – visual, auditory, tactile, olfactory, imagistic, proprioceptive, emotional – or is it *sparse*, limited to one or a few things at a time?[1]

Suppose you've driven the same route to work a thousand times. Today, you're deeply absorbed in thinking about an unpleasant interaction with your department head. Traffic is light; no dangerous situation occurs; you drive habitually. You arrive at the parking area and seem to "wake up" – ah, I'm at work already! – with virtually no memory of having driven there. Consider: Did you have visual experience while you were driving or not? You responded to events on the road: You stopped at the red light, you stayed in your lane. Visual input obviously had some influence on your behavior. But maybe visual input can influence behavior without the involvement of consciousness. Many psychologists believe that a very brief visual display, quickly masked and not consciously experienced, can shape one's later responses, for example in deciphering or choosing words that accord with the masked display.[2] In popular imagination – if not perhaps in actuality (see Bornstein 1989; Trappey 1996) – a single frame of the phrase "Drink Coke" inserted into a film may have no effect on your visual experience yet propel you to the soda machine at intermission. Although absent-minded driving isn't exactly like either of these cases, might you still have had *no conscious experience of the road* as you drove, or only very intermittent conscious experience? The mere fact of behavioral responsiveness doesn't settle the question, at least not without further argument.

Ordinary people's intuitions differ. Researchers disagree. William James (1890/1981) and John Searle (1992), for example, endorse the abundant view, according to which the stream of experience involves both a center of attention and a broad periphery of consciously experienced but unattended objects and background feelings. Julian Jaynes (1976), Dan Dennett (1969, 1991[3]), and Arien Mack and Irv Rock (1998) endorse the sparse view: Consciousness is limited to only one or a few objects, modalities, topics, or fields at a time. The unattended hum of traffic in the background is no part, not even a peripheral part, of your experience when you're sufficiently absorbed in other things.

Who's right? I hope you'll agree that this is a substantive question, and one absolutely central to our understanding of consciousness – that there's a huge difference between thinking that phenomenology abundantly outruns attention and thinking it doesn't. Does it seem obvious to you which view is right?

ii.

If it does seem obvious, you might be inclined to regard apparent disagreement as simply miscommunication. Maybe there are two different meanings of the word "consciousness" – a "consciousness$_1$" and a "consciousness$_2$" (with or without a further fact about which of these is conscious *proper*) – such that disputants are covertly employing different concepts. My own view, however, is that at least some of the disputants disagree in substance. In this chapter, I hope to convince you that it's not entirely obvious whether experience is sparse, abundant, or somewhere in between. And

if it isn't obvious where the truth lies, there's no compulsion to reinterpret people's apparent disagreements as merely terminological.

Still, some terminological clarification is in order. Throughout this book, I've been working with an intuitive notion of "consciousness" or "experience" or (as a stipulated synonym, less fraught with associations and secondary meanings) "phenomenology". I see no way to avoid resting on intuition in such matters, and I hope that you and I have both latched onto the same intuitive concept. If not, though, it's in this chapter, I think, that conceptual mismatch poses the most serious threat to the argument.

We're stuck with an intuitive notion of "consciousness" because the two best formal avenues for definition are closed. We can't define "consciousness" *analytically* because consciousness is a foundationally simple concept indivisible into component parts. It's not like "bachelor" (a marriageable but unmarried man) or "quadrilateral" (a closed planar figure with four straight sides). Nor can we characterize consciousness *functionally* by appeal to the role it plays in a system (a "heart" is an organ that pumps blood, "currency" is whatever physical tokens serve as the medium of economic exchange), since the functional role of consciousness, if any, is still very much in dispute. (Maybe the functional role of consciousness will never be clear: See the Preface and the end of this chapter.)

One characterization some people find helpful is this: Your conscious experience is whatever it is by virtue of which (in the terminology Nagel popularized in his 1974 article discussed in Chapter 4) there's "something it's like to be you", while there's nothing it's like (presumably!) to be a rock or a toy robot. But this characterization has

all the disadvantages of definition by synonymy. Unless you're already comfortable with "what it's like" talk, and mean by it the same thing that I do, it won't help much. We can invoke other synonyms as well: "subjective experience", "phenomenology" as used in contemporary Anglophone philosophy of mind, Ned Block's (1995) "phenomenal consciousness", David Chalmers's (1996) "qualia".

The best approach, I think, is to clarify by example and contrast: By "furniture" I mean tables, chairs, dressers, beds, that sort of thing, and not plates, doors, or toys; by "square" I mean these sorts of things and not these others. With enough positive and negative instances, hopefully one gets the idea. Relatively uncontentious examples of conscious experience include sensations of objects to which one is paying close attention, words uttered silently to oneself, dreams, deliberately formed visual images, thrills of emotion. Uncontentiously nonconscious are immune system response, dendritic growth, early visual processing, undetectable subliminal stimulation. I intend the term "consciousness" to pick out the most obvious difference between these two sorts of processes. In characterizing consciousness this way, of course, I have to trust that everyone regard the same property as the "obvious" one. Compare with the example and contrast approach to characterizing "square": There will many properties that all the demonstrated examples possess and all the contrasts lack – for example the property of *having four equal sides at right angles each less than two miles long* or the property of *being drawn by me at 1:01 p.m.* (supposing I drew all the examples but none of the contrasts at that time). But human cognition being what it is, I trust that those are not the properties you will latch on to, given those examples and contrasts.

In case of consciousness, though, could this be too optimistic? The term "consciousness" is multi-faceted. Might people latch on to different facets when presented with the examples and contrasts I've given, thus understanding different things by "conscious" and talking past each other – perhaps creating the appearance of substantive debate whether there's really only terminological confusion? For example, people use "conscious" sometimes to mean something like "awake". In this sense, dreaming people, no matter how vivid their dreams, are not conscious. Clearly, given that I've listed dreams as paradigm cases of conscious processes, this is not the sense of "consciousness" I intend. How about something like "reportable"? Perhaps the example cases are all reportable and the contrast cases not? Well, it depends on what you mean by "reportable": In some sense of "report", I can report that my immune system is right now repelling a viral sinus infection. But maybe there's some narrower sense of "reportable" – *introspective* reportability? – in which immune system response is not reportable. I'm inclined to think that this is *not* the most obvious difference between the examples and the contrasts: It's not clear how reportable dreams are (at best very partially, I think), and though "reportability" is something psychologists often care about, it doesn't have much salience as a property, I think, for most people.[4] Furthermore, introspective reportability in this sense may be parasitic on the more obvious and basic concept of consciousness.

The functional role or informational features of the example cases may systematically differ from the contrast cases in several ways, but if so those differences are not immediately obvious. Immune system response, for example, is informationally complex. Unrecalled dream experiences have little or no impact on behavior. Whatever systematic or functional differences there are between the examples and contrasts is not

obvious in a way that makes such differences the natural human choice as the target property. Perhaps a process is conscious just in case it plays a certain functional or informational role, but if so that's a fact to be discovered.

Okay, then, here's another possible competitor in the vicinity: *focally conscious*. I've specifically avoided listing cases like our sensory processing of the unattended refrigerator hum either among the examples or among the contrasts so as not to beg the question in favor of sparseness or abundance. As a result, perhaps there are two concepts with very different extensions that are picked out equally well by my examples and contrasts: focally conscious and conscious simpliciter (whether focal or not) or maybe (in a sparse vocabulary) conscious and "pre-conscious". Those who say consciousness is sparse may have the first concept in mind and those who say it is abundant may have the second, while not disagreeing in substance.

The possibility worries me. I believe it's a substantive question, not merely terminological, how sparse or abundant consciousness is, but I find it impossible to characterize consciousness in such a way as to entirely rule out the possibility that the whole dispute is just an efflux of terminological confusion. Favoring a substantive interpretation of the dispute, though, are two considerations: First, the property I've been calling consciousness, or experientiality, or phenomenality, or what-it's-like-ness, is I think so *very* obviously present in the examples and absent in the contrast cases that even if there's a nearby property that my examples share and my contrast cases lack, that other property is not the property people will naturally latch on to when presented with the examples and contrasts – unless, perhaps, those people have been thinking so much about information processing, psychological function, and the like, that properties other than

the target one start to seem to them like salient, obvious competitors. I urge such people to try, for the purposes of this exercise, to mimic the ordinary person's ignorance of such matters. We can all, if we want, mean the same thing by "consciousness". Second, I hope in this chapter to give the reader a feel for why the methodological and introspective issues in the territory are daunting enough to *enable* at least the possibility of substantive disagreement between people on questions in the vicinity. If such substantive disagreement is possible, then it's plausible to suppose that not every philosopher and psychologist sees the matter the same way, differing only in terminology.

Let me conclude this section by explicitly warding against one particular way of thinking about consciousness that risks hiding the substantive concern behind matters of definition. We mustn't equate consciousness with "awareness" if by "awareness" we mean sensitivity to outside stimuli. In absent-minded driving, I'm clearly "aware" of the road in the sense of being responsive to stop lights and turns. In some sense, I'm also "aware" of subliminal stimuli if they influence my behavior in the right way. To equate consciousness with this purely response-dependent sense of "awareness" misses the central issue, either begging the question in favor of abundance or shifting the topic from phenomenal consciousness to something else. Similarly, in the contemporary cognitive psychological literature on implicit perception, conscious "awareness" of a stimulus is sometimes assumed to be present if subjects can perform above chance on certain forced-choice tests regarding that stimulus, even if they feel that they're merely guessing (e.g., Holender 1986; Hannula et al. 2005). Again this begs the central question. It conflates the epistemic or detection sense of "awareness" with the phenomenal, experiential sense (see, e.g., Paap's 1986 critique of Holender and my further discussion in Chapter 7). In

fact, it's probably best to avoid the word "awareness" entirely because of its tendency to equivocate between those two senses.

The sense of "consciousness" or "experience" or "phenomenology" or "what-it's-like-ness" I'm using in this chapter and this book – the sense I take to be standard in most recent philosophical and psychological research on consciousness – is at least conceptually, if not empirically or introspectively, consistent with the unattended press of your feet against your shoes being part of your consciousness and also with its not being so. The aim of this chapter is to address that empirical or introspective question.

iii.

Those who see consciousness as abundant, such as James and Searle, generally provide little positive argument. They tend simply to state the position and expect the reader to agree. For example, James writes:

> The next thing to be noticed is this, that every one of the bodily changes, whatsoever it be, is felt, acutely or obscurely, the moment it occurs.... Our whole cubic capacity is sensibly alive; and each morsel of it contributes its pulsations of feeling, dim or sharp, pleasant, painful, or dubious, to that sense of personality that every one of us unfailingly carries with him (1890/1981, p. 1066-1067).

James invokes no further considerations in defense of the view than its intuitive appeal – here or (as far as I'm aware) anywhere else in his work.[5] Siewert (1998), arguing for the abundance of visual experience specifically, prepares the ground somewhat more

carefully, clarifying what's at issue and what the abundant view is *not* committed to. He emphasizes that every detail needn't be appreciated sharply or separately – an important qualification. But when it comes time for defense of abundance, so clarified and qualified, Siewert gives us no more than James or Searle. It's as though he implicitly assumes that the only potential source of disagreement is misunderstanding, which once cleared up leaves the abundance of visual experience simply evident to reflection.

The problem with this, of course, is that not everyone believes that consciousness is abundant, even when the view is stated clearly. We don't all share James's and Searle's intuitions on the matter. Some people believe that the shirt on one's back and the shoes on one's feet aren't experienced – not even vaguely, inarticulately, peripherally – at every moment of the day; they believe one's visual phenomenology may lapse entirely from time to time. This is not an obviously preposterous opinion. Others find themselves torn or uncertain, or inclined to see one modality as abundantly ever-present and another as experienced only sparsely and sporadically. And of course even if there were a broad intuitive consensus favoring the abundant view, that consensus might be mistaken. Surely, then, it would be good to defend abundance by something more than its natural charm.

Some advocates the sparse view likewise rely principally on folk intuition. David Armstrong (1981), for example, appears to think it simply evident that we lack visual experience in the absent-minded driving case. Julian Jaynes (1976) no more experiences, he says, the pressure of the unattended tree against his back than a person with no visual cortex visually experiences the light coming into her eyes. They invite us to agree based on our own sense of our experience but do not otherwise defend these claims.

A war of philosophical intuitions thus threatens. Never to my knowledge has such a war had a happy outcome.

iv.

We might then look for empirical arguments favoring one view over the other, arguments that go beyond mere appeal to the reader's intuitive sense of her own experience. Those who see consciousness as abundant, as far as I know, offer either no positive arguments or question-begging ones, such as Searle's (1993) bald assertion that our capacity to shift attention to previously unattended stimuli proves that we had pre-existing conscious experience of those stimuli. (What the capacity to shift attention shows, of course, is that we do some perceptual processing outside attention, not – at least not without considerable further argument – that that pre-attentive perceptual processing is *conscious.*)

Ned Block, Christof Koch, and others (Block 2007; Koch and Tsuchiya 2007; Srinivasan 2008) argue empirically against at least very sparse views. Block argues, partly introspectively and partly empirically, that the conscious experience of a briefly presented visual display (for example on a computer) extends considerably beyond the focally attended features that can be reported. Koch and Tsuchiya mention that people can report the gist of a peripherally presented visual scene (saying, for example, whether the scene contains an animal) even when their attention is largely consumed in a demanding task presented at the center of the display. But such discussions don't actually touch the issue at hand. As Koch and Tsuchiya stress, experimental results (or

phenomenological observations) in such paradigms at best show consciousness in the *near*-absence of attention: There may be some sense in which the entire visual display is attended, even if not centrally. What's more directly pertinent to the question of this chapter is whether subjects experience the frame of the computer monitor, or the picture on the wall behind the computer, or the pressure of their flesh against the chair, when they are concentrating intensely on a computer display. On such matters, Block is silent and Koch falls back on intuition (2004, p. 165).

Advocates of the sparse view often offer empirical arguments for their position, but these arguments, like Searle's, are badly question-begging. A favorite argument is this: Absent attention, we fail to parse, respond to, notice, or remember what one might ordinarily think would be salient stimuli – a stream of speech coming in one ear (Cherry 1953; Moray 1959) or a woman in a gorilla suit walking through a ball game (Simons and Chabris 1999). Therefore, it's said, we're "blind" (or "deaf" or "numb") to these stimuli; we don't experience them (e.g., Dennett 1991; Mack and Rock 1998; Wright 2005). Here's the flaw in that argument: It's one thing (indeed a very important and interesting thing) to show that we don't do much processing of unattended stimuli; it's quite another to say that we have *no experience whatsoever* of those unattended objects. The conclusion simply doesn't follow (and many psychologists refrain from drawing it). We may not *parse* the speech semantically (very much) or *represent* the black blob in the middle of the crowd as an ersatz gorilla, but we may still experience that unattended speech and gorilla in some more inchoate way (Simons 2000). Furthermore, unless we really *are* blind, or deaf, or numb, we do process the unattended stimuli to some extent – as Searle points out, and as is all acknowledged on all sides of the debate. We are drawn

to the unexpectedly looming object, the unanticipated call of one's name, the familiar phone ring or doorbell that others can barely hear, the gentle tap on the shoulder. Such things must first register pre-attentively in some way to call our attention. The question is whether whatever limited processing or responsiveness or preparedness to respond we have prior to attention is enough to underwrite actual sensory consciousness. The present argument (and similar arguments involving "change blindness", e.g., Rensink 2000, 2004) doesn't address that question.

Some of Mack and Rock's experiments (e.g. Rock et al. 1992; Mack and Rock 1998) may give us pause. For example, subjects directed to attend carefully to a cross (presented for a fifth of a second, followed by a mask) will often fail to report some other stimulus (a dot, a triangle, etc.) unexpectedly presented in a nearby visual region, against an otherwise uniform background. Afterwards, they may say all they saw were the cross and the background. Mack and Rock describe these subjects as "inattentionally blind": They had no experience whatsoever of the unexpected figure. The conclusion is tempting.

But on reflection, the Mack and Rock experiments should no more trouble those who see experience as abundant than does the obvious fact that someone deeply absorbed may fail to notice a distant (or even not-so-distant) shout, saying afterward that she heard nothing or heard only the uniform buzz of traffic. Several interpretations consonant with abundance suggest themselves. One might accept that the figure (or the shout) was not at all experienced, yet still hold that the uniform unattended background color (or traffic hum) *was* experienced: Perhaps the sensory systems failed to register anything of enough interest to merit more than "filling in" or representing the unattended field as uniform; it

doesn't follow that there's no conscious experience of that uniformity. Or perhaps the figure contributed in an inchoate and unparsed way to an experience reported as uniform but actually an immemorable jumble – part of a stream of visual experience fluctuating not only with major changes in the display, measured in fractions of a second, but also with each saccade, blink, afterimage, accommodation, and glitch. The Mack and Rock experiments simply don't address these possibilities.

What evidence do we have, then, on the crucial, foundational question about consciousness posed at the beginning of this chapter? Only conflicting folk psychological intuitions and badly question-begging arguments. In other words, we have essentially *no* evidence. Such is the absolute infancy of consciousness studies.

v.

Assuming that there really is a substantive issue here, how are we to approach it? Further studies of the relationship between attention and successful report of stimuli won't, I think, help much. We already have the key data: People have some, but only a very limited, sensitivity to unattended stimuli. The question remains: Is that sensitivity (whatever it is) enough to underwrite phenomenology? At this point, the interpretive questions loom larger than the flat empirical ones. People will (sometimes) deny having seen, heard, felt, unattended things; but does that mean that those objects, or the fields containing those objects, or the entire unattended modality in which they can report nothing (other than perhaps "there wasn't a loud shout within a meter of my ear" or "the whole background didn't flick into a scene from *Gilligan's Island*") was entirely

unexperienced, rather than vaguely or inchoately or immemorably experienced? The typical attention-and-reportability study presupposes, more than it addresses, these larger interpretive issues, or else remains silent on them.

If we knew the neural basis of consciousness, we could perhaps use that knowledge to address the sparse-vs.-abundant question; but we don't know it. In fact – a point I'll return to later – we may never *be able* to know it until we determine whether experience is abundant, since (it seems) we need at least a rough understanding of what processes are conscious and not conscious before looking for a common neural basis among the conscious ones; and until we settle the sparse-vs.-abundant question we don't have even a rough understanding of what neural processes are the conscious ones.

Are we left, then, with introspection? – with simply *asking* ourselves, or experimental subjects, whether experience is sparse or abundant? If you take anything from this book, I hope it's a nervousness about relying too trustingly on introspective report in the study of consciousness. And indeed scholars *have* tried addressing the question introspectively – James, Jaynes, and Searle, for example – and they come to different conclusions. I assume that this is not because James and Searle actually experienced every morsel and modality while Jaynes lived largely in blankness.

Advocates of the sparse view have often remarked on another problem, too, that plagues attempts to address this question introspectively: the "refrigerator light" problem, so named after the mistaken impression a child might have that the refrigerator light is always on because it's on whenever he checks it (Thomas 1999; Block 2007). Jaynes puts the point nicely:

> It is like asking a flashlight in a dark room to search around for something that doesn't have any light shining on it. The flashlight, since there is light in whatever direction it turns, would have to conclude that there is light everywhere. And so consciousness can seem to pervade all mentality when actually it does not (1976, p. 23).

Does it seem to me that I have tactile experience of my feet in my shoes? Yes it does, now that I think of it. That I have auditory experience of the hum of the computer? Yes, I guess I seem to be experiencing that now too. But of course I can't conclude from such observations that I constantly experience such things when I'm *not* thinking about them. The mere fact of thinking about whether I experience my feet in my shoes may itself *create* that experience. What we want to know is whether I was experiencing my feet in my shoes *before* the matter came to mind. But that's now in the irretrievable past; I've been thinking too much about introspection, about my feet; I'm corrupted.

The question is thus a rather difficult one to study. The most obvious methods fail.

vi.

But perhaps we haven't introspected carefully enough. Maybe we can dodge the refrigerator light error and other major errors through better introspective method. After all, the sparse and abundant views posit radically different phenomenal worlds. It seems that there should be *some* introspectively discoverable difference between them.

To avoid the refrigerator light error, we might try this: Give participants beepers to wear during their normal daily activities, beepers that sound only at long intervals, when participants are likely to be immersed in other things. Instruct the participants to reflect, each time the beeper sounds, on what their experience was *immediately before* the beep, when (in most cases, presumably) they won't have been thinking about the sparseness or abundance of experience, or about their feet, or about the traffic in the background. Some participants might be asked to report everything in their experience; others whether or not they had visual experience; others simply whether they had tactile experience in the left foot.

A beeper is appealing because it has a sharp onset, targeting a specific moment of experience, and because participants can be told in advance what to reflect on in the targeted experience. No seconds-consuming verbal query is necessary. One can combine the advantages of surprise and preparation. With a little practice, the participant ideally could reflect on her naturally occurring experience within a second of each sampled event. (For more on the advantages of beepers in studying consciousness, see Hurlburt and Heavey 2004.)

I tried exactly this, and in the remainder of the chapter I'll discuss the procedure, the results, and the questions that arise. I'm not convinced that this approach can resolve the question of the sparseness or abundance of experience; but unless we plan to disregard subjective report completely, it seems that knowing what people *say* about their experience when prompted to reflect on it is an essential starting point. I offer this chapter in much the same spirit as I offered Chapter 5: If we want to be serious about consciousness, we must find some better methods for getting at it than casual observation

by untrained introspectors. But the problems that ensue reveal the long, difficult, perhaps impossible path before us.

vii.

I loaned beepers to 21 people, about half philosophy graduate students and about half miscellaneous well educated other folks. I divided them into five conditions: the *full experience* condition, the *full tactile* experience condition, the *full visual* experience condition, the *tactile left foot* condition, and the *far right visual field* condition. (I describe the experiment in more detail in Schwitzgebel 2007.) I told participants in the full experience condition only that our aim was to explore everyday experience generally; they did not know the specific purpose of the research. Participants in the other four conditions were fully informed and asked for their initial take on the sparse-vs.-abundant debate, which I explained using intuitive examples.

Participants wore the beeper for 3-4 hours a day at their convenience within 24 hours of a scheduled hour-long interview, collecting approximately 6-8 samples during a diversity of activities. In most cases, we collected four sets of samples over four separate sampling days. I instructed participants to carry a sheet of paper and a pen so that when a beep sounded they could immediately record their experiences. Participants in the four informed conditions (full visual, full tactile, far right visual field, and tactile left foot) were emphatically instructed *first* simply to note the presence or absence of the relevant experience (yes, no, leaning yes, leaning no, or not sure) and only *after* that to consider other aspects of their experience and environmental situation. I also emphasized that

participants should skip any sample to which they could not instantly respond. I also encouraged them to skip samples they felt were too private. No participant reported skipping more than two samples in a single interview day.

I interviewed participants at length about their experiences, including asking for everything they could remember of the details of the experience reported, as well as their environmental situation at the time, always focusing as precisely as possible on the last instant before the beep. My idea in asking for such detail was (a.) to communicate a serious interest in conscientious accuracy; (b.) to convey in the context of specific reports what sort of phenomena participants should be noting the presence or absence of; and (c.) to give the participants ample opportunity to clarify ambiguities in their reports, to resolve or discover confusions, to express and respond to concerns about the methods of the study, and to develop their own sense of the phenomena, through questions focused as carefully and neutrally as possible on the sparse-vs.-abundant issue. (For participants in the full experience condition, the sparse-vs.-abundant issue was discussed alongside other theoretical questions, such as the nature of emotional and imagery experiences, so as not to reveal the research hypothesis.) In general, I encouraged theoretical discussion. I clarified as much as possible what is meant by "consciousness" or "experience" in hopes of avoiding question-begging conceptions. I played devil's advocate, gently (I hope!) raising various potential doubts and concerns both about reports of experience and about reports of lack of experience. At some point in the interviews, most of the participants heard, or themselves raised, all the doubts and concerns mentioned in section x below.

Generally the first time a participant in one of the visual conditions (the full experience, full visual, and the far right visual field conditions) denied visual experience

in a sample, I introduced what I called the "phenomenal blindness" thought experiment. I explained phenomenal blindness as follows: There's a difference between blindness as pure blackness (like in the dark, though see Chapter 8 for more on this) and blindness as genuine absence of visual experience, like the lack of visual experience you have of what's behind your head. (Or does it seem to you like there's a curtain of blackness back there?) A phenomenally blind person is someone blind in the absence-of-experience sense. Once I felt the participant understood this distinction, I asked the following question: At the sampled moment, could a phenomenally blind person, a twin of you in all respects except lacking visual experience, have had the same conscious experience as you at that moment? I mentioned that, of course, a real blind person might differ in several ways, including in her potential to respond to a looming object, in the quality of her visual experience, in a lack of visual imagery, etc. However, I asked participants to disregard such differences if possible. A few participants rejected the thought experiment as faulty – I invited them to reject it, if they thought so – but most participants said they found the comparison helpful. I stressed that my aim was just to ask as clearly as possible whether she completely lacked visual experience. I emphasized that either answer was okay and that it was also okay to say "I don't know" or "I don't remember". I would occasionally return to this thought experiment if it seemed helpful. For participants in the tactile conditions I offered a corresponding "numbness" thought experiment, with a distinction between numbness positively felt and numbness as absence of tactile experience.

My interview approach of course completely violated the ordinary methodological advice – ordinary by contemporary standards, that is (see Chapter 5) –

that participants be as naive as possible. Regarding the sparseness or abundance of experience, I think it's practically impossible to be naive, to have no initial inclinations or implicit assumptions. Given that, I thought it better to flood my participants with competing alternatives and sources of skepticism than to leave them to the silent guidance of their own initial or emerging theories. I hoped to make it seem not entirely obvious to them whether experience is sparse or abundant, so that perhaps they would be guided more by the data than by their own presuppositions. Of course, I thus risked inflicting my own views too much on participants or forcing them into compromise views; I'll discuss these problems, as well as the effects of participant bias, in section x. The most obvious alternative method – simply giving participants standardized forms to complete – seemed to me too readily to invite misunderstanding and comfortable theory-driven patterns of response. Contemporary consciousness studies already has plenty of naive respondents and theory-driven armchair specialists. I wanted to see what ordinary intelligent folks would say about their experience when asked to think repeatedly and carefully about it, especially when faced with potential sources of doubt.

At the end of the last interview, I asked participants to guess whether I was more inclined toward a sparse or abundant view (or, in my labels of the time, a "thin" or a "rich" view). I also asked whether their opinions had changed over the course of sampling.

viii.

From each participant I collected 9 to 30 samples (with a mean of 17), excluding samples not discussed in the interviews. (For most participants, an hour was insufficient time to explore, carefully in detail, all the sampled experiences, especially in the first two interview days.) For analysis I classified participants' answers into three categories: "yes or leaning yes", "undecided or don't know", or "no or leaning no". I excluded undecided samples from analysis and also samples in which the participant reported thinking about the experiment. For most participants, such exclusions were a small minority.

Table 1 outlines the main results. In sum, the majority of participants in the three visual conditions – 8 out of 13 – reported visual experience in every single one of their samples. However, a significant minority, the other 5, reported no visual experience whatsoever in some of their samples, by the very stringent standard implied in the "phenomenal blindness" thought experiment described in section vii above. Although no one reported far right visual experience or tactile experience in every single sample, every participant reported such experiences in at least half their samples. Two participants reported occasional tactile left foot experiences (one reported it in 3 of 19 samples, another in 4 of 22) while two others reported such experiences very frequently (12 of 15 and 11 of 12).

TABLE 1

type of experience	conditions included	# of people reporting that	median reported	distribution of reported rates (one percentage

		type of experience in 100% of their samples	rate of that type of experience	rate for each participant)*
any visual experience	full experience, full visual, far right visual field	8 out of 13	100%	56%, 73%, 74%, 85%, 89%, 100%, 100%, 100%, 100%, 100%, 100%, 100%
far right visual experience	far right visual field	0 out of 4	63%	50%, 55%, 71%, 89%
any tactile experience	full experience, full tactile†	0 out of 8	76.5%	50%, 56%, 69%, 75%, 78%, 79%, 86%, 89%
tactile left foot experience	tactile left foot	0 out of 4	49%	16%, 18%, 80%, 92%

Taken at face value, these results appear to conflict with *both* the sparse view and the abundant view. Advocates of abundance typically assume that we have *constant,* or

* The percentages in the full visual condition statistically differ from all other conditions (Mann-Whitney, one-tailed, $p < .05$).

† Participants in the tactile left foot condition were not asked to report their tactile experience in general so as not to interfere with their focus on the foot. In contrast, participants in the far right visual condition found it quite natural to discuss the general presence or absence of visual experience.

very nearly constant, visual and tactile experience – probably even constant tactile experience in the left foot (recall James's statement that every "morsel" is "sensibly alive"). The tactile data appear to contradict that claim. So also do some of the data from the visual conditions: Participants often denied peripheral visual experience, and some of them denied having any visual experience whatsover in a substantial minority of samples.

Against the sparse view, every participant reported experience of unattended objects or in unattended modalities in some samples. I haven't attempted to quantify this since the self-report of attention is fraught with perils and confusions I didn't even attempt to prevent or remedy; but I did explicitly ask participants about it from time to time. Every participant but one was unambivalently confident, at least once, of having had a conscious experience without attention – including those who began the experiment seemingly with a strong commitment to a very sparse view of experience. Even if we disregard self-reports of attention, it seems unlikely that participants were attending to events in their far right visual field 63% of the time or to tactile events in their left foot 49%, or even 16%, of the time during the course of several hours of normal activities, unless wearing the beeper dramatically altered their run of experience.

Thus one might read the data as supporting a moderate view, a view somewhere between the very sparse and very abundant views normally espoused by those who write on this topic. My participants universally exited the experiment with a moderate view of some sort, thinking that experience extends well beyond the field of attention but does not include the entire field of every major modality anything like 100% of the time. Typically, they expressed some degree of what seemed to be genuine surprise at their results – those initially inclined to think of consciousness as sparse (10 of the 21

participants, based on the preliminary interview) surprised to seemingly find experience where they thought there'd be none, and those initially inclined toward abundance (11 of the 21) surprised at what they took to be the absence of experience in some cases. Most reported moderating their view by the end of the experiment.

I would love to be able to agree with the consensus of my participants. Unfortunately, I find myself overwhelmed with qualms. Most of my concerns are not particular to this experiment, though. Jointly I think they cast pretty serious doubt on the scientific tractability of the sparse-vs.-abundant dispute – and consequently on any issues that turn on its resolution, such as (I'll argue) a general theory of consciousness.

ix.

Before discussing the methodological worries, however, let me mention a concern about the theoretical viability of a moderate view. Theorists tend to split between seeing consciousness as very sparse or as very abundant. Here's why (I think): To make sense, theoretically, of the moderate view, one must introduce an extra moving part into one's theory of sensory consciousness. (I'm bracketing here, as I have throughout the chapter, questions about non-sensory experience.) We already have good theoretical reasons, independent of any specific commitments about consciousness, to allow into our perceptual theories the phenomena of attention and supraliminal perceptual attunement or responsiveness (I won't say "awareness", for reasons discussed earlier). It's easy and natural to suppose that conscious experience co-occurs with one of these – the former on the sparse view, the latter on the abundant. There isn't as natural a theoretical space,

however, for something that might explain why, if we accept a moderate view, some unattended sensory stimuli are consciously experienced and others aren't. If perceptual consciousness isn't causally inert (or even, on some theories, if it is), it ought to have some important, fundamental, cognitive-functional correlate. But what could that be, on the moderate view? Nothing of broad currency in contemporary psychology seems to play quite the right role.

We can suggest things. Maybe there's a kind of *diffuse* attention, distinct from focal attention, which is capable of being spread broadly across multiple sensory modalities and objects (but still not across all major modalities all the time). Or maybe intense concentration pulls enough resources away from non-focal sensory processing to prevent unattended stimuli from entering consciousness, while less intense concentration permits those stimuli to be (peripherally) experienced. Such views may be plausible. But still, we should be leery of embracing them without the corroborating support of behavioral, non-introspective evidence of such mechanisms. Conversely, however, if we *can* find cognitive-behavioral evidence of diffuse forms of attention or intense forms of concentration, perhaps they'll map nicely onto introspective reports of the presence or absence of peripheral experience and so support a moderate view. That would be pretty nifty.

Advocates of moderate views such as Victor Lamme (2003, 2005) and Koch and Tsuchiya (2007) have also suggested other candidate functional roles that decouple consciousness and attention but do not imply that conscious experience is abundant. Yet given the methodological concerns I'm about to raise, I'm not too optimistic about such possibilities.

x.

I'll divide the concerns into two groups, concerns about *overreporting* and about *underreporting* – the first, of course, suggesting that experience is sparser than participants say, the second that it's more abundant. (This section is long. If you find yourself inclined to skim, the concerns with the larger numbers seem to me the more troublesome.)

Overreporting concern #1: The effects of wearing the beeper. Many participants expressed concern that participating in the experiment would cause them to reflect more about the relevant modality or region and thus experience it more, distorting their results toward the abundant side. I grant that it's likely that there's some effect of this sort. However, since the experiment is not concerned with small differences, only a very large effect of this type would invalidate the general results – only a pervasive transformation of experience, moment to moment, for hours at a time. That doesn't seem especially likely. Besides its a priori implausibility (about which I'm sure people can disagree), it seems reasonable to suppose that if there *were* such a large effect we'd see either rising rates of experience (as participants were trained to think about that modality or region) or declining rates (as participants grew more accustomed to the beeper and let it affect them less), depending on the mechanism guiding the supposed beeper-caused transformation of experience. But average rates of reported experience were stable between the first and last days (8 participants reported more experience on the last day, 8 less, and 13 [mostly at ceiling] reported the same amount[6]). And of course I excluded from analysis samples

in which participants reported thinking about the experiment – only eight of the total samples, less than half a sample per participant. Most participants reported quickly becoming absorbed in their ordinary activities, largely forgetting about the beeper until it sounded.

Overreporting concern #2: Bias. Experimenters must always hope that their own bias isn't driving their results, but rarely are they in a position to evaluate the likelihood of that. Experimenter bias effects are, of course, pervasive in psychological research (as Robert Rosenthal in particular has emphasized, from Rosenthal and Fode 1963 through Rosenow and Rosenthal 1997; see also the brief discussion in Chapter 3). They seem especially likely to play a role to consciousness research and in research involving open-ended interviews. I admit that I came into the experiment not neutral between the theses but thinking that experience was probably abundant – though I found my conviction somewhat shaken as the experiment progressed. So I sympathize with the advocate of the sparse view who suspects that it's mainly my bias driving the results and that the subjects of a more even-handed researcher would have reported sparser experience. One test of this, of course, would be for a sparse-minded researcher to replicate my methods. To assess my own bias, I did two things. First, at the end of the experiment I had subjects guess which view I favored. Only 7 of the 21 subjects guessed than I favored the abundant view, while 13 guessed that I favored the sparse view (one refused to guess). Second, I asked Russ Hurlburt to review some of my interview tapes. I sought Hurlburt's opinion in particular because he is the world's leading practitioner of open interviews on randomly sampled experiences, because he repeatedly emphasizes the importance of being open-minded in interviewing style (e.g., in Hurlburt and Heavey 2006; Hurlburt

and Schwitzgebel 2007), and because he appears to favor a sparse view. He said he thought my interviews were even-handed. I conclude that at least my *overt* bias was not extreme. It's of course still possible that my bias affected responses covertly.

I originally also thought participant bias would be a major factor in the results. However I was fortunate to have participants evenly balanced in their initial inclinations toward sparseness or abundance (10 vs. 11, in an absolute, not a comparative, assessment). Also there was a surprisingly weak relationship between participants' initial inclinations and their final results. Participants inclined toward a sparse view were only slightly more likely to report sparse results than were participants with an abundant view. Counting the full experience and far right visual field groups twice (once for each type of recorded experience [any visual and any tactile for the full experience subjects, any visual and far right visual for the far right visual field subjects]), I found that in 17 cases participants' results tended in the direction of their initial inclinations, relative to the results of the group as a whole, and in 12 cases the results went against their bias – not a statistically significant tendency. For example, one subject who initially expressed a strong inclination toward the abundant view nonetheless reported visual experience in only five out of nine samples. I don't think either of these facts – the equal distribution of initial bias or the weak relationship between initial bias and final results – should entirely dispel concerns about participant bias, but I do think it's reasonable to be optimistic that it isn't *mainly* participant bias driving the results.

Overreporting concern #3: Timing errors. Participants might have reported more experience than they actually had either if they reached too far into the past, gathering up the last conscious experience, whenever it was, in the relevant modality or region, or if

they inadvertently reported on their experience *after* the beep, experience that may have been created *by* the beep. Certainly people are subject to illusions of timing.[7] In response I can say two things, neither entirely adequate, I know: (1.) I repeatedly stressed the importance of trying to home in as accurately as possible on the *last undisturbed moment before* the beep. For what it's worth, the participants all felt they could do this, most of the time, with reasonable if not perfect accuracy. (2.) Unless the timing illusions are worse for central visual experience than for peripheral and tactile experience – and I see no reason to suppose this – the asymmetry of the results, and thus the main effect, still stands.

A version of this concern, trickier to defuse I think, turns on the idea that the target experience may be affected by the beep even if the participant accurately reports the beep as being experienced only *after* the target experience. This needn't be as exotic a matter as backwards causation, of course. The beep is presumably experienced only some time after it stimulates the ear – only after, perhaps, some tens of milliseconds of neural processing. It could be that that preliminary processing, before the beep is actually experienced, affects whatever is experienced as having happened immediately before the beep. Maybe, even, experienced temporal order, when events are near enough together, is partly an after-the-fact reconstruction (see Dennett 1991). We could hardly expect participants to be able to assess such matters introspectively.

Overreporting concern #4: Stimulus error and confabulation. Participants may have reported on states of the world rather than on their *experience* of the world, leading to overly abundant or otherwise erroneous reports; or they may have reported on what seemed *plausibly* to have been their experience rather than on the actual experience itself.

For example, a participant asked to report on visual experience in the far right visual field may simply have reported on visual *objects* to her right – what it seems must have been in her field of view (and thus, on the abundant view, experienced by her) – regardless of whether those objects were actually experienced immediately before the beep. Wearing the beeper in piloting this experiment, I sometimes had the following reaction: The beep sounds, I close my eyes (some participants did this, some didn't, I left it up to them), and I attempt to recall my immediately prior visual experience. There was a black street in front of me, green trees to my left. But am I simply now recalling the *objects* that I remember to have been before me, or am I recalling my *experience* of those objects? The two judgments are different, as is evident from cases of remembered illusion, but in veridical cases it's no trivial task to pull them apart. (The issue here is similar to the issue about "transparency" and the extent to which we introspect combination tones, discussed in Chapter 5.) Titchener and Boring call the mistake of confusing reportage of outward objects with reportage of the experience of those objects "stimulus error" (or "R-error"). Their suggestions for how to avoid stimulus error, however, are sketchy and not especially helpful in this particular context (see Titchener 1901-1905, 1910/1915, 1912; Boring 1921).

The following four facts are at best a *partial* response to this concern. (1.) Every participant (except perhaps one) appears to have understood the distinction between reporting experience and reporting remembered objects, at least on a superficial level; and each at least once denied experience of something in her sensory environment that *could* have been experienced (e.g., an object in her field of view, the contact between flesh and shoe). (2.) If participants generally reported on what was in their environment,

that is, on what the abundant view would predict they would be experiencing, then we would see near ceiling results in every condition, which we don't. (3.) In the far right visual field condition, participants quite readily reported blurriness or vagueness in their experience – properties, of course, of the experience itself, not of objects in the outside world. (4.) Participants favoring a sparse view should, it seems, have been less prone than others to confabulate experience in accord with an expectation that everything in their sensory fields would be experienced. Presumably they had no such expectation. They should have been quite ready to recognize a difference between knowing an object is nearby and having sensory experience of that object, since their view demands that the two often come apart. And yet their results looked very much the same as those of the participants favoring the abundant view (see overreporting concern #2 above).

Underreporting concern #1: Preference for mixed reports. Experimental subjects often prefer moderate or mixed responses to extreme ones. I attempted to counter this preference by stating explicitly in the interviews (in the four informed conditions) that it would be okay to respond entirely with yesses or entirely with noes and that in fact that would be interesting as support for the theoretical views at stake – but I doubt I was entirely successful. Since participants gave nearly uniform "yes" answers to general visual experience questions, pressure to mix it up can't explain the entire pattern of the results, of course; but a friend of abundance might suggest that it's only in the "obvious" central vision case that participants will have had the self-confidence to present an extreme pattern of data. (No parallel argument is available to the friend of sparseness. She'll need something other than the middle-of-the-scale bias to explain the visual results that run counter to her position.)

If I browbeat people into changing their report from sparse to abundant and vice versa, that could also generate an overall pattern of intermediate data. There was, indeed, some risk of this, since I tried to counteract participants' biases by occasionally pointing out the plausibility of the alternative view. Fortunately, an analysis of first-day vs. final-day results shows no evidence of massive browbeaten conversions. Twenty subjects stayed either above or below the median report rates for the type of experience in question, and nine crossed median (again counting full experience and far right visual field participants twice).

Underreporting concern #2: Subtle experience. In the tactile left foot condition, one participant – a philosophy graduate student who reported tactile left foot experience in 11 out of 12 samples – typically said he had a *general* sense of the position and disposition of his body, its posture and contact with things. He usually claimed not to have experienced his left foot separately and distinctly, but only as a small and subtle part of this holistic bodily sense. This pattern of reporting apparently surprised him: He initially expressed an inclination to the sparse view. Indeed, within the full tactile and tactile left foot conditions, four participants (three initially sparse-biased, one initially abundant-biased) reported discovering a holistic bodily sense of this sort, and all had above-median results. Is this just a compelling theoretical idea that, once entertained, inclined these participants to invent experience to match it (see overreporting concern #4 above)? Or did this idea reflect a discovery of, and allow them to report, a subtle sort of background experience that others might easily miss?

Underreporting concern #3: Memory error. I don't think we should be too concerned about *long-term* memory error. In the four informed conditions (that is, all but

the full experience condition), the basic data point is very simple, and I permitted participants to consult their notes during the interview. If a participant gets it right in the first few seconds after the beep, it seems unlikely she'll misreport later. (There's obviously much more room for long-term memory error in the full experience condition, but fortunately those results harmonize with the results in the other conditions.) The bigger issue is this: What's the likelihood of failing to remember the targeted moment of experience, or non-experience, between the time of its occurrence immediately before the beep and the act of judgment shortly after the beep?

It's noteworthy how much we fail to remember even over very short intervals, if our attention is not upon it as it occurs. Perhaps the most striking recent experiments on this topic are the "change blindness" experiments of Ronald Rensink and others (e.g., Rensink et al. 1997, 2000; Simons and Levin 1998). You look at a picture. It flickers and is replaced by a very similar picture, with one major change. For example, a large railing substantially changes position, or a large jet engine near the middle of the picture disappears and then, after another flicker, it reappears. It's often difficult to detect that change, even when the stimuli are presented repeatedly. (There are many change blindness demonstration videos on the internet if you're not familiar with the phenomenon.) Or: You're having a conversation with a stranger. In the middle of the conversation, two people carrying a door briefly walk between you, and the stranger is surreptitiously replaced by another person in different clothing, with a different build. Many people fail to notice the change. Experiments like this – along with older experiments on the unreliability of eyewitness testimony (Münsterberg 1927; Loftus 1979; Haber and Haber 2000) and on the forgetting of mundane everyday details like the

direction of Lincoln's face on the penny (Sanford 1917/1982; Nickerson and Adams 1979) – suggest that we may fail to encode or remember surprisingly large aspects of our perceived external environment.

Whether we likewise fail to encode or remember large tracts of our *stream of conscious experience,* as distinguished from our outward environment, is an open question, but I see no reason to suppose it merits a different answer, especially if experience is abundant. If sensory experience is a complex, massively detailed flux, it may be at least as expensive and pointless to retain as are the unimportant or readily available environmental details we so easily forget. The beeper method brings to a practical minimum the delay between experience and reflection, but the experience and reflection still aren't simultaneous – they can't be, if we're to avoid the refrigerator light error – and that very non-simultaneity may be enough to guarantee the forgetting of substantial portions of experience that are never recorded even in short-term memory.

How should we assess these various shortcomings? Three strike me as fairly tractable, perhaps with further experiments: the effects of wearing the beeper, bias, and the preference for mixed reports. The timing error issue seems a bit trickier. Perhaps it could be partly addressed by asking participants to focus on their experience not *immediately* before the beep but rather, say, one full second before the beep; but that seems likely to aggravate at least the memory issues, perhaps also the confabulation and bias issues. The subtle experience error and the confabulation/stimulus error seem to me trickier still. The subtle experience error raises some of the same tangled issues as the combination tones case from Chapter 5: How do we know when someone is introspecting

well enough that we can trust her assertion that she lacks some subtle experience? Though I was able to say a few things against the confabulation and stimulus error worry, it seems to me that such reassurances should rightly leave an advocate of sparseness largely unmoved. Perhaps even the sparse-leaning participants confused, for example, the visual memory of a peripherally seen object with the memory of having visually experienced that object. In veridical cases, these do seem awfully hard to tease apart. And the short term memory concern seems to me absolutely intractable: How could anyone introspectively discern whether an experience, however recently past, never occurred or instead occurred but was never encoded into memory?

If all the concerns pointed the same direction, we could perhaps nonetheless reach some rough conclusions. For example, if the only really troubling concerns suggested underreporting, we might conclude that experience was *at least* as abundant as participants suggest. That would still leave us no basis for deciding between a moderate and a radically abundant view – no basis in a retrospective self-report experiment like this one, at least. Whether there might be other bases for deciding, I'll address shortly. Unfortunately, there are daunting concerns on both sides. Experience might still, for all we know, be anywhere from radically sparse to radically abundant.

xi.

The phenomenological difference between sparseness and abundance is vast. If defenders of abundance are right, then our stream of experience is aswarm with detail in many modalities at once, both inside and outside the field of attention; if defenders of

sparseness are right, experience is limited to one or a few attention-occupying activities or objects at a time. On the first view, unconscious perception exists only in the margins if it exists at all; on the second, *most* of our perception is unconscious. On the first, we always have a complex flow of visual experience; on the second we may quite often have no visual experience at all. What, it seems, could be easier than to decide between these two views? Shouldn't a moment's reflection settle the matter incontrovertibly?

The fact that it doesn't is striking, and interesting – and methodologically very important. One might take the apparent evasiveness of what seems like it should be an obvious issue to suggest some merely linguistic or communicative trouble, a problem of speaking past each other, of disagreement or inconsistency in the use of words. Yet as I suggested earlier, I doubt we can justifiably comfort ourselves with that thought. Forget about the interpretation of James and Jaynes and Searle, and just consider the issue in a single vocabulary. However you think of "consciousness", it's an open question exactly how far it spreads. Maybe you think there's only a narrow range of plausible views about that, but if the methodological problems are as serious as I've suggested, then there should actually be a broad range of substantively distinct plausible views – all the way from very sparse to very abundant. Such views remain live despite the phenomenological gulf between them because the refrigerator light error hobbles concurrent introspection and because stimulus error, memory error, and the potential subtlety of the target experiences make even carefully collected retrospective reports difficult to interpret.

Here's the key question I'd like you to consider: Could this experiment, or another on the same topic, have been done appreciably better, so as to avoid these concerns? If so, I'm afraid I've probably wasted your time with a crummy experiment,

or at best given an illustration of what *not* to do in consciousness research. So can we construct, even if only in fantasy, a better experiment? Let's consider self-report methods first. Any concurrent self-report method – any method that asks the participant to report on his experience as it occurs – will be polluted by the introspective act itself. Any retrospective method will invite concerns about short-term memory at least. Both approaches will likely accrue charges of participant bias, confabulation, conceptual confusion, potential insensitivity to subtle experiences. Different approaches will, perhaps, involve trade-offs between these potential failings – stressing the possibility of subtle experience may increase sensitivity to those experiences but also increase the risk of confabulation; selecting unsophisticated subjects may reduce certain sorts of theoretical bias but raise the risk of conceptual confusion; etc. But no self-report method can, I think, effectively avoid the sorts of concerns raised here. My experiment, however crummy, is not *contingently* crummy.

How about more objective or theoretical methods? Could we do away with subjective reports – at least subjective reports on this particular issue – and simply, say, look at the brain or at patterns of behavior? Cognitive-behavioral approaches without the aid of self-report will not, I think, solve our problem. Either they operationalize "consciousness", equating it by definitional fiat with some behavioral or cognitive pattern, or they simply beg the question. Does mere behavioral responsiveness, for example, or above-chance responding on forced-choice questions about the presence of stimuli, demonstrate "consciousness" of those stimuli? Not in any way that should move an advocate of the sparse view. Does failure to report stimuli outside of attention show that they weren't consciously experienced? Not in any way that should satisfy an

advocate of abundance. The problem is, we simply do not know enough yet about the relationships between cognition and phenomenology to take *any* objective cognitive-behavioral measure of consciousness as valid without begging the question at hand.

The same holds for purely neural approaches. One might think we could resolve the issue by discovering what neural features indisputably nonconscious mental episodes share and what neural features indisputably conscious mental episodes share and then see whether our brain's response to unattended stimuli looks more like the former or more like the latter. Unfortunately, the gap between the indisputably conscious and the indisputably nonconscious is too wide to be bridged in this way. Early visual processing and early lexical processing are indisputably nonconscious (at least in mainstream opinion); focal visual attention and deliberate episodes of inner speech are indisputably conscious. There will be many neural and cognitive features the latter share that the former lack, and some of those features will be shared with tactile processing of an unattended foot. But which of those features are essential for consciousness? We don't know.

Furthermore, we may *never* know until we resolve the sparse-vs.-abundant dispute. As I mentioned in section v, the search for neural correlates of consciousness (if there even are such) makes no sense unless we have in advance at least a *rough* idea of the sorts of mental states that are conscious, and we don't have even a rough idea of the sorts of mental states that are conscious until we settle the sparse-vs.-abundant question. Suppose we find a neural state that occurs when and only when a sensory process involves focal attention. Is that a neural correlate of consciousness? Not if experience is abundant; it might just be a correlate of attention. Suppose we find a neural state that

occurs whenever there is sensory responsiveness of any but the most minimal sort. Is that a neural correlate of consciousness? Not if experience is sparse; it might just be a correlate of sensory sensitivity. Suppose we find a neural state that correlates with something like *diffuse* attention. To declare that to be the neural correlate of consciousness, thus adopting a moderate view, begs the question against both sparse and abundant views, unless further arguments can be marshaled; but those further arguments must either turn on self-report, which as we saw appears to be highly problematic in this domain, or on further behavioral, cognitive, or neural measures, which will be equally question begging.

xii.

If there were a theory of consciousness so elegant and so nicely articulated with the data that it compelled acceptance independent of the sparse-vs.-abundant question, we could perhaps turn to it to resolve the dispute. But there is no such theory. One challenge in constructing such a theory is that the sparseness or abundance of experience appears to be a part of the basic data *in light of which* a theory of consciousness must be constructed. Arguably, we must know whether experience is sparse or abundant *before* we can justifiably embrace a general theory of consciousness.

Consider some actual theories. Bernard Baars (1988, 1997) has advanced a "global workspace" theory of consciousness, according to which sensory content is conscious just in case it's in the narrow theater of working memory, where only a small amount of attended content can be manipulated at a time and broadcast across the

cognitive system. Francis Crick (1994) argues that the neural correlate of consciousness is synchronized 40 hertz neural oscillations in the subset of neurons corresponding to an attended object. David Chalmers (1996) argues (tentatively) that consciousness is present wherever there is information processing. Built right into these views from the start – explicitly for Baars and Crick, and very near the surface for Chalmers – is a commitment to a sparse (Baars, Crick) or abundant (Chalmers) view. Such theories will not and should not seem plausible to researchers with different antecedent opinions on the sparse-vs.-abundant question. Those of different inclination may quite reasonably regard such theories as, at best, theories of something *else,* not consciousness per se – perhaps focal attention, perhaps information processing.

I have trouble imagining a theory of consciousness ever arising that is both compelling independently of its take on the sparse-vs.-abundant issue and specific enough in its implications to resolve that issue. To settle the one matter, it seems we must first settle the other. There's a Catch-22 here, I'm inclined to think, *unless* we find a way to generate trustworthy subjective reports on the topic; but we've just seen reason to think that such subjective reports do not merit our trust. Parallel concerns arise regarding the distribution of consciousness across the range of animals and machines, at least for those theories of consciousness that purport to treat consciousness in general and not just human consciousness: Such theories will build in from the start commitments on the distribution of consciousness, commitments it may be impossible to ground independently of the theory, given the difficulty in obtaining and evaluating non-human reports. Problems of this sort may be simply insurmountable. And if so, they suggest that a methodologically well justified scientific consensus on a theory of consciousness –

even just on a theory of the neural or cognitive correlates of consciousness (setting aside the "hard problem" of how *anything* material could give rise to consciousness) – is beyond our reach.

This could all just reflect the limits of my imagination. Five centuries ago, no doubt, it would have seemed inconceivable that we could reach a well justified scientific consensus on the first few seconds of the history of the universe. The ingenuity of scientists almost always in the long run embarrasses naysayers about science. So maybe I shouldn't be so pessimistic. Maybe I should say only this: The obstacles are formidable. Not that I counsel abandoning academic research into the issue. Even if science does fail, there's always philosophy, where the insolubility of a problem is no objection its contemplation – where we can, if necessary, bat questions around endlessly and count it as success if we give structure to our confusion.

[1] In earlier work (Hurlburt and Schwitzgebel 2007; Schwitzgebel 2007), I used the terms "rich" and "thin" instead of "abundant" and "sparse". But people complained to me that "thin" and "rich" and "thin" don't sound like opposites (consider Paris Hilton); and the terms have been used widely by philosophers to mark other distinctions (even by me in Chapter 5, where I described auditory experience as "rich" if it has many aspects). So I feel compelled try out a new terminology.

[2] E.g. Marcel 1980; Merikle et al. 2001; Snodgrass et al. 2004. But for critiques of this literature see Holender 1986; O'Brien and Opie 1999; Holender and Duscherer 2004.

[3] Usually Dennett seems to espouse the sparse view, and he has confirmed this in conversation. However on p. 137 of his 1991 book, he seems to tilt the other way on the absent-minded driving case, and on the following page he seems to express the view that at least in some cases there may be no fact of the matter. In general, I don't see how all of Dennett's statements about consciousness can be reconciled, an issue I explore in Schwitzgebel 2007b. (Dennett responds to these concerns in the same journal issue. How adequately, the reader may judge for herself.)

[4] Similar remarks apply, I think, to Block's (1995) notion of "access consciousness".

[5] However, on p. 430 of the same work James writes: "The pressure of our clothes and shoes, the beating of our hearts and arteries, our breathing, certain steadfast bodily pains, habitual odors, tastes in the mouth, etc., are examples from other senses, of the same lapse into unconsciousness of any too unchanging content". This seems to either

qualify or contradict the passage on p. 1066-1067. (Or could it be a comment on a different topic, "consciousness", perhaps, as opposed to what is "felt"? Unfortunately for that attempt at resolution on p. 432 James seems to say that ignored stimuli are also "unfelt"). Further evidence is on p. 402-403, where James seems to acknowledge that we have "sensations" and "sense-impressions" without attention. Fortunately, James interpretation is not the aim of this chapter, so if I've erred in attributing an abundant view to him, plenty of other psychologists and philosophers could serve in his stead.

[6] The total of 29 reflects the fact that I counted the full experience and the far right visual field participants twice each – the full experience participants for both their visual and tactile reports and the far right visual field participants for both their full visual and their far right visual field reports. See also the second note to Table 1.

[7] See Bessel 1822/1823; James 1890/1981; Geldard and Sherrick 1972; Libet 1985; Dennett 1991; Nijhawan 1994; Spence et al. 2001.