

Eric Schwitzgebel

*Do You Have Constant Tactile
Experience of Your Feet in
Your Shoes?
Or Is Experience Limited to
What's in Attention?*

Abstract: According to *rich* views of consciousness (e.g., James, Searle), we have a constant, complex flow of experience (or 'phenomenology') in multiple modalities simultaneously. According to *thin* views (e.g., Dennett, Mack and Rock), conscious experience is limited to one or a few topics, regions, objects, or modalities at a time. Existing introspective and empirical arguments on this issue (including arguments from 'inattention blindness') generally beg the question. Participants in the present experiment wore beepers during everyday activity. When a beep sounded, they were to take note of the conscious experience, if any, they were having at the last undisturbed moment immediately prior to the beep. Some participants were asked to report any experience they could remember. Others were asked simply to report whether there was visual experience or not (and if so, what it was). Still others were asked about experience in the far right visual field, or tactile experience, or tactile experience in the left foot. A majority of participants in the full experience and the visual conditions reported visual experience in every single sample. Tactile and peripheral visual experience were reported less often. However, the proper interpretation of these results is uncertain.

Correspondence:

Eric Schwitzgebel, Department of Philosophy, University of California at Riverside, Riverside, CA 92521-0201, USA. Email eschwitz@ucr.edu

Keywords: Consciousness, attention, inattentional blindness, change blindness, introspection, visual experience, tactile experience, peripheral vision

1. ‘Rich’ and ‘Thin’ Views of Experience

Do you have constant tactile experience (or ‘consciousness’ or ‘phenomenology’) 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, no part of ‘what it’s like to be you’? We might think of consciousness as like a soup. Is it a *rich* soup, full of experience in a wide variety of modalities simultaneously — visual, auditory, tactile, olfactory, imagistic, proprioceptive, emotional — or is it a *thin* soup, limited to one or a few things at a time?

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 lot and seem to ‘wake up’ — ah, I’m at work already! — with virtually no explicit memory of having driven there. Now consider this: Did you have visual experience, visual consciousness, 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 regulatory influence on your behaviour. But maybe visual input can influence behaviour 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 (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). 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 impact 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 behavioural

responsiveness doesn't settle the question, at least not without further argument.

Ordinary people's intuitions differ. Researchers disagree. James (1890/1981) and Searle (1992) endorse the rich view, according to which the stream of experience involves both a centre of attention and a broad periphery of consciously experienced but unattended objects and background feelings. Jaynes (1976), Dennett (1991), and Mack and Rock (1998) endorse the thin view: Consciousness is limited to only one or a few objects, modalities, topics, or fields at a time. The unattended hum of the 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 very near the heart of consciousness studies. Although 'consciousness' is a treacherous word with no clear analytic or functional definition — no such definition, at least, that doesn't beg crucial questions — I don't think we need see the issue at hand as a mere vocabulary dispute. It seems to me at least, and I hope it seems to you, reader, that there's a real difference between thinking that conscious experience broadly outruns attention and thinking it doesn't. And the truth of the matter is not so patently obvious on its face that disagreement and uncertainty are impossible.

2. 'Consciousness'

Still, the words 'consciousness' and 'experience', used imprecisely, can muddy the debate, so let me make as plain as I can what I mean by them. Unfortunately, the two most respectable avenues for definition appear to be closed. No *analytic* definition is possible, I suspect, because 'consciousness' or 'experience' (I use the words interchangeably) is a foundationally simple concept indivisible into component parts. It doesn't seem to be like 'bachelor' (a marriageable but unmarried man) or 'quadrilateral' (a four-sided planar figure). Nor can we, at this stage of our knowledge, 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 exchange), since the functional role of consciousness is still very much in dispute. One characterization some people find helpful is this: Your conscious experience is whatever it is by virtue of which 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 (Nagel, 1974), it won’t help much. We can invoke other synonyms as well: ‘phenomenology’, ‘subjective experience’, Block’s (1995) ‘phenomenal consciousness’, Chalmers’ (1996) ‘qualia’.

We can characterize 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 playing 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. What’s at stake in the rich-vs.-thin debate is whether our processing of the unattended hum of the refrigerator is like attentive perception, dreaming, conscious imagery, inner speech, and felt emotion — whether it has that property, ‘consciousness’, that I’m trying to point to with these examples — or whether it’s rather more like subliminal perception or immune system response.

I’d like to ward against two ways of thinking about consciousness that risk hiding the substantive concern behind matters of definition. First, 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 behaviour in the right kind of way. To equate consciousness with this purely response-dependent sense of ‘awareness’ misses the central issue. 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 are merely guessing (e.g., Holender, 1986; Hannula *et al.*, 2005). Again, this begs the central question — conflating, it seems, an epistemic or detection sense of ‘awareness’ with a phenomenal, experiential sense (see, e.g., Paap’s 1986 critique of Holender). In fact, it’s probably best to avoid the word ‘awareness’ entirely because of its tendency to equivocate between those two senses.

We should also avoid using the term ‘conscious’ to mean acutely self-conscious or focally attentive. It *may* be true, if the thin view is correct, that conscious experience is limited to a few things in focal attention or about which we are especially self-attuned, but if so,

that's a *substantive* truth, not a matter to be settled analytically a priori by our definition of consciousness. The sense of 'consciousness' or 'experience' or 'phenomenality' or 'what-it's-like-ness' I'm using in this essay — the sense standard in most recent philosophical and psychological research on consciousness — is at least conceptually, if not empirically or introspectively, consistent with unattended input from your feet in your shoes being a part of your consciousness, and also with its not being so.

3. The Inadequacy of Intuitive Appeals on this Question

Advocates of the rich view, such as James and Searle, provide little positive argument. Typically, they simply 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, pp. 1066–7).

James invokes no further considerations in defence of the view than its intuitive appeal — here or (as far as I'm aware) anywhere else in his work. Siewert (1998), arguing for the richness of visual experience specifically, prepares the ground somewhat more carefully, clarifying what's at issue and what the rich view is *not* committed to. He emphasizes that every detail needn't be appreciated sharply or separately. But when it comes time for *defence* of the rich view, so clarified and qualified, Siewert gives us no more than James. It's as though he implicitly assumes that the only potential source of disagreement is misunderstanding, which once cleared up leaves the rich view simply evident to reflection.

The problem with this, of course, is that not everyone agrees with the rich view, even clearly stated. We don't all share James's and Searle's intuitions on the matter. Some people believe that the shirt on one's back or 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 consciousness may lapse entirely from time to time. This is not an obviously preposterous opinion. Others find themselves torn or uncertain, or inclined to a rich view in one modality and a thin view in another. And even if there were broad intuitive consensus favouring the rich view, that consensus might be mistaken.

Surely, then, it would be good to defend the rich view by something more than its natural charm.

Some advocates of the thin view likewise rely principally on folk intuition. Armstrong (1981), for example, appears to think it simply evident that we lack visual experience in the absent-minded driving case. Jaynes writes:

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 (1976, p. 22).

Jaynes compares the lack of consciousness here to the lack of consciousness in early visual processing and cortical blindness — cases he seems to regard as obviously and uncontroversially nonconscious. He invites us to agree based on our own sense of our experience but does not otherwise defend these claims.

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

4. The Inadequacy of Existing Empirical Arguments

We might then look for empirical arguments favouring one view over the other, arguments that go beyond mere appeal to the reader's intuitive sense of her own experience. Advocates of the rich view, 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 pre-attentive perceptual processing is *conscious*.)

Advocates of the thin view often offer empirical arguments for their position, but these arguments, too, are badly question-begging. A favourite 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 are '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 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 universally acknowledged on all sides of the debate. We are drawn to the unexpectedly approaching 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 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 that 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, Mack and Rock's experiments should no more trouble advocates of the rich view than does the obvious fact that someone deeply absorbed in something 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 the rich view suggest themselves. One might grant Mack and Rock's suggestion that the figure (or the shout) was not at all experienced, yet still hold that the uniform unattended background colour (or the traffic hum) *was* experienced: Perhaps the sensory systems fail to register anything of enough interest to do anything other than 'fill in' or represent the region in question (or the auditory environment) 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. Mack and Rock's 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 paper? Only conflicting folk psychological intuitions and badly question-begging arguments. In other words, we have essentially *no* evidence. Such is the absolute infancy of our understanding of consciousness.

5. Unpromising Avenues

How are we to approach the issue then? 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 consciousness? 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 metre of my ear', or 'the whole background didn't suddenly flick into a rerun of *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 rich-vs.-thin question; but we don't know it. In fact, perhaps, we will never *be able* to know it until we determine whether unattended stimuli are experienced or not — since (it seems) we need at least a rough understanding of what processes are conscious and not conscious prior to looking for a common neural basis among the conscious ones, and until we settle the rich-vs.-thin 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 rich or thin? This should make us nervous for several reasons. For one, subjects'

reports of their own mental states are often untrustworthy. Even introspecting in what would seem to be ideal conditions — carefully, patiently, regarding one’s own currently ongoing conscious experience — we often go badly wrong, or so I would argue (and have argued in, e.g., Schwitzgebel and Gordon, 2000; Schwitzgebel, 2002; 2004; in preparation). Indeed, scholars *have* tried addressing the question introspectively — James, Jaynes, and Searle, for example — and they arrive at different answers. (I assume this is not because James and Searle actually experience every morsel and modality as constantly alive while Jaynes lives largely in blankness!)

Advocates of the thin 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. Jaynes describes the problem 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 you that you have tactile experience of your feet in your shoes? Yes, it does now that you mention it. That you have auditory experience of the hum of the computer? Yes, now that I think of it, I seem to be experiencing that too. But of course we can’t conclude from such observations that we have constant experience of such things when we *aren’t* thinking about them. The mere fact of thinking about whether you have experience of your feet in your shoes may itself *create* that experience. What we want to know is whether you were experiencing your feet in your shoes *before* the matter came to mind. But that’s now in the irretrievable past; you’ve been thinking too much about introspection, about your feet; you’re corrupted.

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

6. Immediate Retrospection

Despite my considerable qualms about introspection (and the fact that my recent career is largely devoted to criticizing it), I can’t seem to resist the urge to search for a better, more empirically controlled, means of addressing this question introspectively. After all, the rich and thin views of experience posit radically different phenomenal

worlds. It seems that there should be *some* introspectively discoverable difference between them! (Doesn't it?)

To avoid the refrigerator-light error, we might try this: Give participants beepers to wear during their normal daily activities, beepers that go off 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 prior* to the beep, when (in most cases) they won't have been thinking about the richness or thinness 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 single, 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 preparedness. 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'm not convinced, I should confess, that this approach can definitively settle the question of the richness or thinness of experience. However, unless we are to completely disregard subjective reports of experience in studying this issue, some data of this sort must be collected and analysed. Knowing what people *say* about their experience when prompted to reflect on it is an essential starting point. Even if serious concerns can be raised about the accuracy of such reports — as I think they can — the attempt to make sense of them, to salvage what is legitimate in them, and to clarify exactly where and why they might go wrong, casts at least some initial light on the phenomena and starts to expose the theoretical and empirical possibilities.

7. Experimental Method

Participants

I recruited 21 participants, mostly through an email sent to philosophy graduate students and an announcement in an upper-division philosophy class at University of California, Riverside, but also through word of mouth. The recruiting announcement mentioned only that I sought volunteers to study conscious experience using a beeper worn during normal daily activity. Eleven graduate students in philosophy

participated and five undergraduates (not all philosophy majors). Other participants included a secretary, a K-12 teacher, a psychology graduate student, a psychology professor, and a statistics professor. Participants were paid \$20 per interview day, except the professors who participated gratis.

Materials

I provided each participant with one ‘beeper’ on loan from Russ Hurlburt whose ‘Descriptive Experience Sampling’ methodology (Hurlburt, 1990; 1993; Hurlburt and Heavey, 2006; Hurlburt and Schwitzgebel, forthcoming) was the inspiration for this experiment (though I diverged from strict DES method). The beepers were designed to emit a beep anywhere from about one minute to one hour after being set. The beep came through an earpiece worn at all times, and I instructed participants to adjust the volume so that the tone would have a sharp and distinct, but not too startling, onset. Once a beep began, it continued until the participant pressed a conveniently located button to reset the timer.

Preliminary interview

I divided participants into one of 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. Participants in the full experience condition were not informed in advance of the specific purpose of the research. I told them that the aim was to explore everyday experience generally and we discussed issues of privacy.

Participants in the other four conditions were more fully informed of the nature of the research prior to obtaining consent. During a preliminary interview, I described the rich and thin views of experience, and I offered intuitive examples of cases in which people might disagree — usually the absent-minded driving case and the shirt-on-your-back case. I asked participants for their initial inclinations on the question, then encouraged them to set those inclinations aside and be open-minded about the research. I carefully explained the refrigerator light error and the appeal of using a beeper to avoid that error.

I attempted to attain some balance between rich-biased and thin-biased subjects within conditions, and also between philosophy graduate students and others. In some cases, this involved not making the final assignment, among the four informed conditions, until the end of the preliminary interview.

I instructed participants in the far right visual field condition to fixate on a point and showed them another point 18 degrees to the right of fixation to define the beginning of the far right visual field.

Instructions for collecting samples

I instructed participants to wear the beeper for 3–4 hours, at their convenience and not necessarily consecutively, within 24 hours of a scheduled interview, to collect approximately 6–8 samples during a diversity of activities. I instructed them to carry a sheet of paper and a pen, so that when a beep was emitted they could *immediately* record their experience. I emphasized the importance of immediacy: If they could not respond instantly, they were told to skip the sample entirely. I also encouraged them to skip samples they felt were too private to share in their entirety. No participant ever reported skipping more than two samples in a single interview day.

I instructed participants in the *full tactile* condition to write down *first* whether they were having a tactile experience or not; *second* what that tactile experience, if any, was; and *third* what they were doing in general and other background information. Proceeding in this order was crucial to the experiment, I emphasized. I also emphasized that it would be perfectly okay to find *no tactile experiences at all*, which would accord nicely with the thin view, or *tactile experience in every single sample*, which would accord nicely with the rich view, or some mixed result. I gave participants in the *full visual* condition the same instructions, except with reference to visual experience; similarly for the tactile left foot and far right visual field conditions. Participants in the *full experience* condition were simply asked to report what experience, if any, they had at the moment immediately prior to the beep.

Sampling interviews

I interviewed participants within twenty-four hours of collecting their samples, in most cases over the course of four separate interview days, with four separate groups of samples. All participants used their sampling notes as a reference, though I stressed that what mattered was what they said in the interview. (I collected the sampling notes at the end of each interview. They generally accorded well with the reports.) By drawing the participants into interviews about their experience, I hoped to (a) convey by my interest a serious attention to detail, (b) convey in the context of specific reports, as accurately as possible, what sort of phenomena the participants should be noting the presence or absence of, and (c) give the participants ample opportunity to clarify ambiguities in their reports, to resolve (or discover) confusions,

to express and respond to concerns, and to develop their own sense of the phenomena by explicitly comparing and contrasting with previous reports and responding to questions.

If a participant reported an experience, I asked follow-up questions about the nature of the experience. For example, if a participant reported a pain in her foot, I asked her to localize it as precisely as possible and to describe it in words as much as she is able, sometimes offering suggestions such as ‘did it seem to be static or changing?’, ‘was it more of a dull ache or more of a sharp prick, or neither of those?’ I also asked what was in the focus of attention in each sample, if there was a focus of attention (occasionally, participants reported having no focus of attention).

In discussing the first sample with participants in the *full experience* condition, I introduced the rich-vs.-thin debate as follows (the others had learned about the debate in the preliminary interview): After the participant had reported experiencing one or a few things and seemed to be slowing down or to have stopped, I asked whether she experienced anything else simultaneously with what was already reported, pointing out that some theorists would expect one to experience many things simultaneously while others would expect experience of only one or a few things at a time. I encouraged discussion of this point, to confirm that participants understood the issue and to gain a sense of their own predispositions. However, the rich-vs.-thin issue was not stressed to the exclusion of other issues. For example, depending on the participants’ reports, we also discussed (sometimes in depth) whether visual imagery need be located in space and whether it interferes with sensory visual experience, whether emotional experience is bodily or cognitive or a hybrid or something else, whether there is an experience of ‘willing’ a bodily motion, etc.

In general, I encouraged theoretical discussion. I discussed most of the topics raised in the first part of this paper with most participants, including clarifying as much as possible what is meant by ‘consciousness’ or ‘experience’, in hopes of avoiding question-begging interpretations. I regularly 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. Most participants heard, or themselves raised, at one point or another, all the doubts and concerns mentioned in Section 10 below and were asked for their take on them.

I recognize that this interview approach completely violates the ordinary methodological advice that subjects be as naive as possible. However, regarding the richness or thinness of experience, I think it’s

practically impossible to be naive, to have no initial inclinations or implicit assumptions. Given this, I thought it better to flood participants with competing alternatives and sources of scepticism than to leave them to the silent guidance of their own initial (or emerging) theories. By doing so, I hoped to make it seem not entirely obvious whether the rich or the thin view was correct, so that perhaps they would be guided more by data than by their presuppositions. I could not, of course, hope to be entirely successful in dislodging participants' biases, but my impression in the end (as we'll see) is that I was more successful than I would have expected. Detailed theoretical discussion during the interview also gave me the opportunity to convey as precisely as possible what was being asked and to correct misapprehensions about the task.

Generally, the first time a participant in one of the visual conditions (the full experience, the 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) and blindness as genuine absence of visual experience, like the lack of visual experience you have behind your head. 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 did at that moment? I mentioned that, of course, a real blind person might differ from the participant in a number of ways, including in his potential behavioural responsiveness (the potential to react to a suddenly looming object, say), and in the quality of his auditory experience, in lacking visual imagery, etc. However, I asked the 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 in this thought experiment was just to ask as clearly as possible whether the participant completely lacked visual experience, or whether there was some vague, or peripheral, or inchoate visual experience. I stressed 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 when a participant reported no experience, if it seemed helpful. For participants in tactile conditions, I offered a corresponding 'numbness' thought experiment (with a distinction

between numbness positively felt and numbness as absence of tactile experience).

I tried always to be balanced and accepting in my interview style, to ask open-ended questions, to convey respect for the participants' opinions, but also frankly to express scepticism and sources of concern about the reports. If the participant seemed to show considerable bias toward one view or another — especially if she used inferential language in describing her experiences (e.g., 'the cabinet was to my right, so I must have had a visual experience of brownness') — I would stress the plausibility of the opposing view and discourage inferential answers.

All this discussion, of course, took a long time. I generally capped interviews at one hour, though they sometimes went longer. On the first interview day, especially, it was often impossible to get through more than one or two samples. If on the second day we still didn't get through more than two samples, I would ask the participant on the third day not to begin with the first sample collected but rather with the last or with some randomly selected sample.

At the end of the last interview, I asked participants to guess whether I was more inclined to the rich or the thin view. I also asked whether their opinions on the debate had changed over the course of sampling.

8. Results

I found most of the participants to be theoretically astute enough to comprehend the issues raised in the interviews, so that I felt they genuinely understood what was being asked and what the substantive and methodological issues were. I did, however, have doubts about the understanding of one of the undergraduates (S16), maybe two (S14), and concerns about the extent to which one of the graduate students (S13) may have allowed himself to get too caught up in challenging me on fundamental theoretical and methodological points (especially regarding whether the rich-vs.-thin dispute is 'merely linguistic'), to the detriment of his responses. I nonetheless include their data in the analysis. The mean number of samples collected was 17, ranging from 9 to 30 (excluding samples not discussed in the interviews).

I classified participants' answers into three categories: 'yes or leaning yes', 'undecided/don't know', 'no or leaning no' as summarized in Tables 1–4. Undecided responses were discarded from later analysis, as were samples in which the participant reported thinking about the experiment. This was a small percentage of the total samples in all but two cases: the disputatious graduate student, who was undecided in 5

Subject	Condition	Yes	Lean Yes	Don't Know	Lean No	No	Thinking of Exp	%	Bias
S4	full exp	9	0	(1)	0	0	0	100%	rich
S5	full exp	18	0	(1)	0	0	0	100%	rich
S6	full exp	9	0	0	0	0	0	100%	mod[erate] /thin
S16	full exp	5	0	0	2	2	0	56%	rich
S1	full visual	12	2	(1)	3	2	0	74%	thin
S2	full visual	22	2	0	0	0	0	100%	rich
S3	full visual	11	2	(1)	0	0	0	100%	mod/rich
S7	full visual	17	0	0	1	1	0	89%	rich
S17	full visual	16	1	0	0	3	0	85%	thin
S9	far right visual	15	1	0	0	0	0	100%	mod/rich
S12	far right visual	12	0	(2)	0	0	0	100%	mod/thin
S18	far right visual	7	1	(1)	0	3	0	73%	thin
S19	far right visual	11	1	(1)	0	0	0	100%	rich

median = 100%

TABLE 1: Were you having any visual experience immediately prior to the beep?

Subject	Condition	Yes	Lean Yes	Don't Know	Lean No	No	Thinking of Exp	%	Bias
S9	far right visual	7	0	(2)	1	6	0	50%	mod/rich
S12	far right visual	10	0	0	1	3	0	71%	mod/thin
S18	far right visual	4	2	(1)	0	5	0	55%	thin
S19	far right visual	7	1	(2)	1	0	0	89%	rich

median = 63%

TABLE 2: Were you having far right visual field experience immediately prior to the beep?

Subject	Condition	Yes	Lean Yes	Don't Know	Lean No	No	Thinking of Exp	%	Bias
S4	full exp	6	1	(1)	0	2	0	78%	rich
S5	full exp	15	0	0	0	4	0	79%	rich
S6	full exp	5	0	0	4	0	0	56%	mod/thin
S16	full exp	8	0	0	1	0	0	89%	rich
S8	tactile	10	2	(3)	1	1	0	86%	thin
S10	tactile	9	2	0	2	9	(1)	50%	thin
S13	tactile	0	6	(5)	2	0	0	75%	mod/rich?
S15	tactile	11	0	0	0	5	0	69%	rich

median = 76.5%

TABLE 3: Were you having any tactile experience immediately prior to the beep?

Subject	Condition	Yes	Lean Yes	Don't Know	Lean No	No	Thinking of Exp	%	Bias
S11	tactile left foot	2	2	(4)	4	14	(4)	18%	rich
S14	tactile left foot	9	3	(2)	1	2	0	80%	mod/thin?
S20	tactile left foot	2	1	0	4	12	(3)	16%	mod/thin
S21	tactile left foot	10	1	(1)	1	0	0	92%	thin

median = 49%

TABLE 4: Were you having tactile experience in your left foot immediately prior to the beep?

of 13 samples, and one participant in the tactile left foot condition who was undecided in 4 of 30 samples and thinking about the experiment in another 4. In analysing the results, I interpreted 'tactile' broadly to include any tactile, somatic, nociceptive, or proprioceptive experience, since participants had some difficulty distinguishing these in their reports. An alpha level of .05 was used for all statistical tests. The data were not assumed to be parametric.

As is evident from Table 1, a majority of participants in the full experience, full visual, and far right visual conditions (8 of 13) reported visual experience in 100% of the samples. No trend or significant difference between the three conditions is evident. The four participants in the far right visual field condition reported experience in that portion of their visual field considerably less often (Table 2). The difference in medians (100% vs. 63%) is significant ($p = .01$) by a one-tailed Mann-Whitney test (adjusted for ties).

Participants in the full experience and full tactile experience conditions reported tactile experience at a median rate of 76.5% (Table 3). (Reports of general tactile experience weren't collected in the tactile left foot condition, so as not to interfere with the focus on the foot.) This

median is significantly lower than the 100% median in the visual conditions ($p = .01$) using a two-tailed Mann-Whitney test. Results in the tactile left foot condition were highly variable and suggestive of the possibility of an underlying bimodal distribution: 18%, 80%, 16%, and 92% (Table 4). The difference in medians between that condition and the full tactile condition (49% vs. 76.5%) is not statistically significant (Mann-Whitney one-tailed, $p = .37$). Participants in the full experience and full tactile experience conditions rarely reported tactile experience in the left foot, but they were not specifically asked to reflect on the potential presence or absence of tactile-left-foot experience in response to the beep.

I classified participants as either rich- or thin-biased, based on the preliminary interview (in most conditions) or the first day discussion (in the full experience condition). In one case (S14), it seemed likely from later discussion that the participant had mischaracterized his initial bias, but for the most part discussion bore out, or at least did not plainly contradict, participants' assessments of their initial predispositions. Eleven participants were classified as rich-biased — as thinking (*prior* to collecting data) that, generally, we have a broad range of experiences in different modalities at any one time — and ten were classified as thin-biased, as thinking that, generally, experience is quite limited (or completely absent) outside a narrow range of attention. This suggests that intuitions regarding the rich/thin debate are roughly evenly divided in the population studied.

One might expect participants with an initial rich bias to report more experience than participants with a thin bias. To check this, I classified participants into two groups: those with a rich bias and results at or above median or a thin bias and below median results, on the one hand, and those with a rich bias and below median results or a thin bias and results at or above median, on the other hand. Participants with scores in two categories were classified twice, once for each category. (Thus, participants in the full experience condition were classified for both their visual and tactile data, and participants in the far right visual field condition were classified for both their general visual and their far right field data.) In 17 cases participants' results accorded with their bias and in 12 cases the results went against the bias. Thus, there was only a weak trend for rich-biased participants to report more experience, not statistically significant ($p = .31$ on the one-tailed Fisher exact test). This comports with my general sense that participants were impressively open-minded and often changed their views and expressed surprise over the course of the experiment.

Finally, although I incline toward the rich view, participants were unable to discern my bias. At the end of the last interview day, given a forced choice, 7 guessed I favoured the rich view, while 13 guessed I favoured the thin view. (One refused to guess.) Most spontaneously said I had been even-handed. So also did an outside observer with a thin bias who listened to one of the audiotapes.

9. Toward a Moderate View?

If we assume participants' responses accurately reflected their experience, then the data seem to support *neither* the rich view *nor* the thin view as normally characterized. Advocates of the rich view typically assume that we have *constant*, or very nearly constant, visual and tactile experience — probably even constant tactile experience in the left foot (recall James's statement about every 'morsel' being 'sensibly alive'). The tactile data, if interpreted at face value, plainly do not support so strong a claim. Neither do some of the data from the visual condition: Peripheral visual experience was often denied and some participants denied visual experience altogether in a substantial minority of samples. And they did so by a very demanding test: In denying visual (or tactile) experience, they agreed, for the most part, that a phenomenally blind (or numb) twin of them — someone with absolutely no visual (or tactile) conscious phenomenology at all — could have had exactly the same experience they did at the last undisturbed moment prior to the beep.

Against the thin 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 remedy or prevent; but every subject except the disputatious graduate student was plainly and unambivalently confident, at least once, of having had positive conscious experience without attention — including subjects who began the experiment with a seemingly strong commitment to a radically thin view of experience. Even if we discount self-reports of attention, it seems highly unlikely that participants were attending to events in their far right visual field 63% of the time or to tactile or somatic 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 radically rich and thin 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 beyond the field of attention but does not occur in every major modality 100% of the time. Typically, participants expressed some degree of what seemed to be genuine surprise at their results — thin-biased participants surprised to (seemingly) find experience where they thought there'd be none, rich-biased participants 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.

10. Overreporting or Underreporting?

I would love to be able to agree with the consensus of my participants. But I have — as the reader no doubt also has — a considerable list of concerns about the accuracy and representativeness of these reports. Still, despite these concerns, I believe some tentative and conditional conclusions are possible. To see this, it's helpful to classify the concerns into two groups: those that suggest my participants *overreported* their experience (and thus that our conscious lives are relatively thin) and those that tend to suggest that my participants *underreported* their experience (and thus that our conscious lives are relatively rich).

But first let me mention a concern that falls into neither category — a concern about the *moderate* view specifically. Why do theorists tend to split between the radically rich and the radically thin, with little middle ground? 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, questions about non-sensory experiential richness.) We already have good theoretical reason, 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 thin view, the latter on the rich. 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 epiphenomenal, it ought to have some 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. Perhaps there's a kind of especially *diffuse* attention, distinct from focal attention, which is capable of being

spread broadly across multiple modalities and objects (but still not across all major modalities all of 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 behavioural, non-introspective evidence of mechanisms of this sort. Conversely, however, if we *can* find cognitive-behavioural evidence of such 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! (See also Lamme, 2003.)

But before we (or is it just me?) get too ebullient, let's look at concerns suggesting that participants may have considerably *overreported* their experience in this study.

1. Effects of wearing the beeper

Many participants expressed concern that participating in the experiment might cause them to reflect more about the relevant modality or region, and thus experience it more, distorting the results toward the rich side. I grant that it's likely there's some effect of this sort. However, since this 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 one supposes to underlie the beeper-caused transformation of experience. But average rates of reported experience were stable between the first and last days (8 participants increased, 8 decreased, and 13 [mostly ceiling] remained the same [counting full experience and far right visual field participants separately for each measure, as described above]). And of course I excluded from analysis samples in which participants reported thinking about the experiment — only a small minority of samples, as is evident from the tables. Most participants reported quickly becoming absorbed in their ordinary activities, largely forgetting about the beeper until it went off.

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, and seem especially likely to play a role in consciousness research and research involving relatively open interviews. I confess that I came into the experiment drawn toward the rich view (though I found my conviction somewhat shaken as the experiment proceeded). Consequently, I sympathize with the advocate of the thin view who suspects that it's primarily my bias behind the results — that a more even-handed researcher would have found results more in accord with the thin view. The surest test of this, of course, is to give it a try (and I hope someone will). I do think my *overt* bias, however, could not have been extreme, given the failure of the majority of my participants to detect it and the positive assessment of my even-handedness by a thin-biased outside observer.

I originally also thought participant bias would be a major factor in the results. However I was fortunate to have participants evenly balanced (11 vs. 10) between rich and thin bias (in an absolute, not a comparative assessment), and was both surprised and delighted by the weak relationship between their expectations and their results. While I think neither of these facts should entirely dispel concerns about participant bias, I do think it's reasonable to be optimistic that it isn't mainly participant bias driving the present results.

3. *Timing errors*

Participants may have erred toward excessively rich reports 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 human beings are subject to illusions of timing (e.g., Libet, 1985; Nijhawan, 1994; Spence *et al.*, 2001). 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.

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 rich or otherwise erroneous reports; or they may have reported on what seems *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 rich view, experienced by her) — regardless of whether those objects were actually experienced immediately prior to the beep. Wearing the beeper in piloting this experiment, I sometimes had the following type of reaction: The beep goes off, 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? Is there, perhaps, no difference between those two judgments? (That last thought seems wrong to me, though perhaps it would bode well for the task if it were true.) It's no trivial matter to sort out such issues (though Titchener, 1901–5; 1910/1915; 1912, and Boring, 1921, make — I think — some promising suggestions).

Four facts partly reassure me on this head: (1) Every participant (except perhaps the confused undergraduate, S16) appears to have understood the distinction between reporting experience and reporting remembered objects, at least on a superficial level, and 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 pressure of her shoe). (2) If participants generally reported on what was in their environment, or what (on a rich view) one might *expect* to have experienced, 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 and/or vagueness in their experience — properties, of course, of the experience itself, not of objects in the outside world. (4) Thin-biased participants should, it seems have been less prone than others to confabulate experience in accord with an expectation that everything, say, in their field of view 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 rich-biased participants.

Finally, three concerns suggesting that participants may have *underreported* their experience:

1. *Preference for mixed results*

Experimental subjects often prefer moderate or mixed responses to extreme ones. I attempted partly to counter this inclination in the interviews, for example by stating explicitly (in the four informed conditions) that all yesses or noes would be okay and indeed interesting as support for the rich or thin views — 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 results, of course; but an advocate of the rich view 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 advocate of the thin view: 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 reports from rich to thin 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 (9 crossing median, 20 staying either above or below [again counting full experience and far right visual participants twice]).

2. *Subtle experience*

In the tactile left foot condition, one participant — a philosophy graduate student who reported tactile left foot experience in 92% of his 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 came into the experiment with a thin bias. Indeed, within the full tactile and tactile left foot conditions, four participants (three thin-biased, one rich-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 the confabulation objection 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?

3. *Memory error*

I don't think we should be too concerned about *long-term* memory error. In the four informed conditions, the basic data point is very simple, and participants may 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 prior to 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 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. It's often difficult to detect that change, even when the stimuli are presented repeatedly. 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 in different clothing, with a different build. Many people fail to notice the change. Experiments like this (along with older experiments on eyewitness testimony [Münsterberg, 1927; Loftus, 1979; Haber and Haber, 2000] and the forgetting of mundane everyday details [Sanford, 1917/1982; Nickerson and Adams, 1979]) suggest that we may fail to encode or remember surprisingly large aspects of our perceived external environment.

Now whether we likewise might 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 the rich view is correct. 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.

11. Conclusion

The phenomenological difference between the rich and the thin views is vast. On the first view, our stream of conscious experience is aswarm with detail in many modalities at once, both inside and outside the field of attention; on the second, the stream of experience is limited to one or a few attention-occupying activities or perceptions at a time. On the first view, unconscious perception exists only on 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 to decide than 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 theoretically 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 I doubt we can justifiably comfort ourselves with that solution. The positions, framed even in a single vocabulary, seem substantively distinct; neither seems preposterous; they seem both *different* and *live*. They remain live despite the phenomenological gulf between them because the refrigerator light error renders concurrent introspection extremely problematic as a source of data, and because stimulus error, memory error, and the potential subtlety of the target experiences make even carefully collected retrospective reports difficult to interpret.

What then are we left with? The present study doesn't escape such concerns. Have I simply wasted your time with a crummy experiment?

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 probably have wasted your time — or at most given a helpful illustration of what *not* to do in consciousness research. But if it's impossible to dodge these concerns — or if attempts to dodge them invite other problems as bad or worse — then I beg you to stave off any impatience you may have with the flaws in my methods, and treat them instead as forcing a choice. Either

embrace a broad-ranging scepticism about the scientific study of consciousness, or accept the methodological disarray of the field and attempt to balance, evaluate, and triangulate the types of concerns that inevitably trouble research in this area. *If* methodological concerns in this field are inevitable, one can either be a purist and do without consciousness (or operationalize the term behaviouristically) or one can do one's best to muddle forward through doubt and ambiguity. If this paper raises this choice clearly, forces a decision on this matter, then I think it has done a service, even if the decision is in the end to reject all research of this sort.

But back to the antecedent: Could research on this topic have been done appreciably better, avoiding the present concerns without falling into others? I think not. Let me stress, first, how deeply wrong I think it is to suppose that 'objective' methods — methods that do away with introspective (or immediately retrospective) reports — will be a panacea. Either they merely operationalize 'consciousness' — equating it by definitional fiat with some behavioural or cognitive pattern — or they simply beg the question. Does mere behavioural responsiveness, for example, or above-chance responding on forced-choice questions about the presence of stimuli, demonstrate 'consciousness' of those stimuli? Not in any sense that should move an advocate of the thin 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 the rich view. The fact is, we simply do not know enough, yet, about the relationships between cognition and consciousness to take *any* 'objective' measure of consciousness as valid without begging the question at hand.

We must rely, at this beginning stage of our inquiry into consciousness, at least partly upon participants' reports. Yet these are inherently troublesome. Any concurrent introspective method — any method that asks the participant to report on her experience as it occurs — will be polluted by the introspective act itself. And any retrospective method will invite concerns about short-term memory, at least. Both methods will inevitably invite charges of participant bias, confabulation, conceptual confusion, potential insensitivity to subtle experiences. Different methods will, perhaps, involve trade-offs between these potential failings — stressing the possibility of subtle experiences, perhaps, will 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, dodge entirely the sorts of concerns raised here.

So this returns us to our choice: Err on the side of caution, and entirely discard attempts to study this foundational question about (phenomenal) consciousness — which I think would be tantamount to abandoning hope in any scientifically precise theory of consciousness — or attempt to weigh and adjudicate the risks and imperfections.

Let's assume we take the latter course: Where does the present research leave us? Perhaps the reader will share my sense that the most troublesome and intractable of the concerns involve stimulus error, subtle experience, and short-term memory error. The other sources of error are all at least partially addressed by the study and amenable to further checking (e.g., by replicating with a different experimenter, by exploring more fully the range and nature of temporal illusions). The subtle experience and short-term memory errors suggest that experience is *at least* as rich as my participants report (possibly even radically rich); and I find myself drawn toward that conclusion — albeit in a tentative and uncertain way, pending further evidence. But first I need to shake the feeling that maybe even my thin-biased participants sometimes confused, for example, the *visual memory* of a peripherally-seen object with a *memory of having visually experienced* that object immediately prior to the beep. Further interpretation of these results turns in large part, I think, on successfully addressing that objection. I seem to have misplaced my brilliant response — probably it's somewhere in this stack of papers over here

Acknowledgements

Many thanks to the participants in this study — as much collaborators as subjects — who spent hours in conversation with me about these topics and doubtless shaped my thinking more than I know. Thanks also to the audience at the 2005 meeting of the Association for the Scientific Study of Consciousness and especially to Nelson Cowan, Dan Dennett, and Russ Hurlburt for helpful discussion and criticism. This research was supported in part by an Academic Senate grant from UC Riverside.

References

- Armstrong, D.M. (1981), *The Nature of Mind* (Ithaca, NY: Cornell University Press).
- Block, N. (1995), 'On a confusion about the function of consciousness', *Behavioral & Brain Sciences*, **18**, pp. 227–87.
- Boring, E.G. (1921), 'The stimulus error', *American Journal of Psychology*, **32**, pp. 449–71.

- Bornstein, R.F. (1989), 'Subliminal techniques as propaganda tools: Review and critique', *Journal of Mind and Behavior*, **10**, pp. 231–62.
- Chalmers, D.J. (1996), *The Conscious Mind* (New York: Oxford University Press).
- Cherry, E.C. (1953), 'Some experiments on the recognition of speech, with one and two ears', *Journal of the Acoustical Society of America*, **25**, pp. 975–9.
- Dennett, D.C. (1991), *Consciousness Explained* (Boston, MA: Little, Brown and Company).
- Haber, R.N. and Haber, L. (2000), 'Experiencing, remembering, and reporting events', *Psychology, Public Policy, and Law*, **6**, pp. 1057–97.
- Hannula, D.E., Simons, D.J. and Cohen, N.J. (2005), 'Imaging implicit perception: Promise and pitfalls', *Nature Reviews Neuroscience*, **6**, pp. 247–55.
- Holender, D. (1986), 'Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal', *Behavioral and Brain Sciences*, **9**, pp. 1–66.
- Holender, D. and Duscherer, K. (2004), 'Unconscious perception: The need for a paradigm shift', *Perception & Psychophysics*, **66**, pp. 872–81.
- Hurlburt, R.T. (1990), *Sampling Normal and Schizophrenic Inner Experience* (New York: Plenum).
- Hurlburt, R.T. (1993), *Sampling Inner Experience In Disturbed Affect* (New York: Plenum).
- Hurlburt, R.T. and Heavey, C.L. (2004), 'To beep or not to beep: Obtaining accurate reports about awareness', *Journal of Consciousness Studies*, **11** (7–8), pp. 113–28.
- Hurlburt, R.T. and Heavey, C.L. (2006), *Exploring Inner Experience* (Amsterdam: John Benjamins).
- Hurlburt, R.T. and Schwitzgebel, E. (forthcoming), *Describing Inner Experience?* (Cambridge, MA: MIT Press).
- James, W. (1890/1981), *The Principles of Psychology* (Cambridge, MA: Harvard University Press).
- Jaynes, J. (1976), *The Origin of Consciousness In the Breakdown of the Bicameral Mind* (Boston, MA: Houghton Mifflin).
- Lamme, V.A.F. (2003), 'Why visual attention and awareness are different', *Trends in Cognitive Sciences*, **7**, pp. 12–18.
- Libet, B. (1985), 'Unconscious cerebral initiative and the role of conscious will in voluntary action', *Behavioral & Brain Sciences*, **8**, pp. 529–66.
- Loftus, E.F. (1979), *Eyewitness Testimony* (Cambridge, MA: Harvard University Press).
- Mack, A., and Rock, I. (1998), *Inattention Blindness* (Cambridge, MA: MIT).
- Marcel, A.J. (1980), 'Conscious and preconscious recognition of polysemous words: Locating the selective effects of prior verbal context', in *Attention & Performance*, ed. R.S. Nickerson, **8**, pp. 435–57.
- Merikle, P.M., Smilek, D. and Eastwood, J.D. (2001), 'Perception without awareness: Perspectives from cognitive psychology', *Cognition*, **79**, pp. 115–34.
- Moray, N. (1959), 'Attention and dichotic listening: Affective cues and the influence of instructions', *Quarterly Journal of Experimental Psychology*, **11**, pp. 56–60.
- Münsterberg, H. (1927), *On the Witness Stand* (New York: Clark Boardman).
- Nagel, T. (1974), 'What is it like to be a bat?', *Philosophical Review*, **83**, pp. 435–50.
- Nickerson, R.S. and Adams, M.J. (1979), 'Long-term memory for a common object', *Cognitive Psychology*, **11**, pp. 287–307.
- Nijhawan, R. (1994), 'Motion extrapolation in catching', *Nature*, **370**, pp. 256–7.
- O'Brien, G. and Opie, J. (1999), 'A connectionist theory of phenomenal experience', *Behavioral and Brain Sciences*, **22**, pp. 127–96.

- Paap, K.R. (1986), 'The pilfering of awareness and guilt by association', *Behavioral and Brain Sciences*, **9**, pp. 45–6.
- Rensink, R.A. (2000), 'When good observers go bad: Change blindness, inattention blindness, and visual experience', *Psyche*, **6** (9).
- Rensink, R.A. (2004), 'Visual sensing without seeing', *Psychological Science*, **15**, pp. 27–32.
- Rensink, R.A., O'Regan, J.K. and Clark, J.J. (1997), 'To see or not to see: The need for attention to perceive changes in scenes', *Psychological Science*, **8**, pp. 368–73.
- Rensink, R.A., O'Regan, J.K. and Clark, J.J. (2000), 'On the failure to detect changes in scenes across brief interruptions', *Visual Cognition*, **7**, pp. 127–45.
- Rock, I., Linnett, C.M., Grant, P. and Mack, A. (1992), 'Perception without attention: Results of a new method', *Cognitive Psychology*, **24**, pp. 502–34.
- Sanford, E.C. (1917/1982), 'Professor Sanford's morning prayer', Reproduced in *Memory Observed*, ed. U. Neisser (San Francisco, CA: W.H. Freeman).
- Schwitzgebel, E. (2002), 'How well do we know our own conscious experience? The case of visual imagery', *Journal of Consciousness Studies*, **9** (5–6), pp. 35–53.
- Schwitzgebel, E. (2004), 'Introspective training apprehensively defended: Reflections on Titchener's lab manual', *Journal of Consciousness Studies*, **11** (7–8), pp. 58–76.
- Schwitzgebel, E. (in preparation), 'The unreliability of naive introspection'.
- Schwitzgebel, E. & Gordon, M.S. (2000), 'How well do we know our own conscious experience? The case of human echolocation', *Philosophical Topics*, **28**, pp. 235–46.
- Searle, J.R. (1992). *The Rediscovery of the Mind* (Cambridge, MA: MIT).
- Searle, J.R. (1993), 'The problem of consciousness', *Consciousness & Cognition*, **2**, pp. 310–19.
- Siewert, C.P. (1998), *The Significance of Consciousness* (Princeton: Princeton University Press).
- Simons, D.J. (2000), 'Attentional capture and inattention blindness', *Trends in Cognitive Sciences*, **4**, pp. 147–55.
- Simons, D.J. and Chabris, C.F. (1999), 'Gorillas in our midst: Sustained inattention blindness for dynamic events', *Perception*, **28**, pp. 1059–74.
- Simons, D.J. and Levin, D.T. (1998), 'Failure to detect changes to people during a real-world interaction', *Psychonomic Bulletin & Review*, **5**, pp. 644–9.
- Spence, C., Shore, D.I. and Klein, R.M. (2001), 'Multisensory prior entry', *Journal of Experimental Psychology: General*, **130**, pp. 799–832.
- Snodgrass, M., Bernat, E. and Shevrin, H. (2004), 'Unconscious perception: A model-based approach to method and evidence', *Perception & Psychophysics*, **66**, pp. 846–67.
- Titchener, E.B. (1901–1905), *Experimental Psychology* (New York: Macmillan).
- Titchener, E.B. (1910/1915), *A Text-Book of Psychology* (New York: Macmillan).
- Titchener, E.B. (1912), 'The schema of introspection', *American Journal of Psychology*, **23**, pp. 485–508.
- Trappey, C. (1996), 'A meta-analysis of consumer choice and subliminal advertising', *Psychology & Marketing*, **13**, pp. 517–30.
- Wright, W. (2005), 'Distracted drivers and unattended experience', *Synthese*, **144**, pp. 41–68.

Paper received March 2006