

**Describing Inner Experience?
Proponent Meets Skeptic**

**Russell T. Hurlburt
University of Nevada, Las Vegas**

**Eric Schwitzgebel
University of California, Riverside**

**In Press
The MIT Press
Anticipated Summer 2007**

**Prepublication version
Do not quote without permission of the authors.**

**Part One
Proponent Meets Skeptic**

Chapter One

Introduction

On a remarkably thin base of evidence – largely the spectral analysis of points of light – astronomers possess, or appear to possess, an abundance of knowledge about the structure and history of the universe. We likewise know more than might even have been imagined a few centuries ago about the nature of physical matter, about the mechanisms of life, about the ancient past. Enormous theoretical and methodological ingenuity has been required to obtain such knowledge; it does not invite easy discovery by the untutored.

It may seem odd, then, that we have so little scientific knowledge of what lies closest at hand, apparently ripe for easy discovery, and of greatest importance for our quality of life: our own conscious experience – our sensory experiences and pains, for example, our inner speech and imagery, our felt emotion. Scientists know quite a bit about human visual capacities and the brain processes involved in vision, much less about the subjective experience of seeing; a fair bit about the physiology of emotion, almost nothing about its phenomenology.

Philosophers began in earnest in the seventeenth and eighteenth centuries to describe and classify our patterns of conscious experience. John Locke (1690/1975), for instance, divided experienced “ideas” into those that arise from sensation and those that arise from reflection, and he began to classify them into

types. David Hume (1739/1978) distinguished what we would now call images from perceptual experiences in terms of their “force” or “liveliness.” James Mill (1829/1967) attempted a definitive classification of sensations into the traditional five senses (sight, hearing, touch, taste, and smell) plus muscular sensations and sensations in the alimentary canal. However, despite such efforts, not even the most basic taxonomy of experience was agreed upon; and it is still not agreed upon.

The study of conscious experience acquired a more scientific look with the introspective psychologists of the late 19th and early 20th centuries. Researchers such as Gustav Fechner (1860/1966), Wilhelm Wundt (e.g., 1896/1897), and E. B. Titchener (1910/1915), presented carefully measured stimuli to subjects who had been trained to “introspect” – to take careful note of their immediately occurring (or just passed) experiences. These psychologists aimed to understand how these introspected experiences covaried with changes in stimulation. However, as is well known, after a few decades, behaviorism (which stressed measuring relationships between stimulus and behavioral response rather than stimulus and introspected experience) won the day in mainstream experimental psychology, driving out or marginalizing the study of consciousness. Subsequent elaborations of behaviorism, and later “cognitivism,” allowed more room for the postulation of internal states and mechanisms mediating behavioral responses; yet these internal states and mechanisms were generally assumed to be nonconscious.

Central to the behaviorists' complaint about the introspective study of consciousness was the unreliability of the introspective method, the fact that several decades of work yielded little consensus on even the most fundamental issues. John B. Watson, the early standard-bearer for behaviorism, in his seminal 1913 article "Psychology as the behaviorist views it," criticized the lack of consensus in introspective psychology as follows:

One psychologist will state with readiness that the attributions of a visual sensation are quality, extension, duration, and intensity. Another will add clearness. Still another that of order. I doubt if any one psychologist can draw up a set of statements describing what he means by sensation which will be agreed to by three other psychologists of different training.... I firmly believe that two hundred years from now, unless the introspective method is discarded, psychology will still be divided on the question of whether auditory sensations have the quality of 'extension', whether intensity is an attribute that can be applied to color, whether there is a difference in 'texture' between image and sensation and upon many hundreds of others of like character.... The condition in regard to other mental processes is just as chaotic... (p. 164-165).

The considerable truth in this complaint partially explains the success of the behaviorist overthrow of introspective methodology. The fact that introspective

psychologists had failed to reach consensus about such issues revealed a serious weakness in their methodologies. Furthermore, much of the consensus they did manage to reach was undermined by an early 20th-century shift, among those still interested in consciousness, away from the early introspectionists' focus on the basic "elements" of experience in favor of a more holistic conception of a sensory "Gestalt," indivisible into individual elements. Thus, despite the obvious importance of conscious experience to our lives, and its apparent ready availability for research, conscious experience had largely resisted systematic attempts at scientific description, and its study fell into disrepute.

Although research on consciousness has enjoyed a considerable resurgence since the 1990s, the most basic structural and methodological questions remain unanswered. With little examination, introspection has re-entered psychology and philosophy. Even hard-nosed cognitive neuroscientists ask their subjects about their subjectively felt experience while in the fMRI magnet. However, it should be clear from the history just described that such casual and haphazard introspection cannot be trusted to yield robustly replicable results and accurate generalizations. Furthermore, it seems to us that the introspective methods employed by most current researchers in consciousness studies are less careful than the methods used by introspective psychologists a century ago. Unless better methods can be found, we fear that the scientific study of consciousness may again stall. And if there simply are no better methods, the scientific study of

consciousness may prove wholly impossible in principle: vacuous without introspective report, intractably conflictual with it. Scientists could perhaps elude this difficulty if they found a way to study consciousness without the help of introspective report. We doubt such an enterprise makes sense, but we will not argue the point here. We will assume that any science of consciousness must take, as a fundamental source of data, people's observations and descriptions of their own experience. Thus a re-examination of the adequacy of introspective reports is of central importance to consciousness studies.

That leads us to the question that stands at the heart of this book: To what extent is it possible accurately to report conscious experience? One author of this book, Russ Hurlburt, has argued that we can profit from the demise of classical introspection and create methods for reporting conscious experience that largely avoid the old pitfalls. He has developed one such method, Descriptive Experience Sampling (DES), to be described in the next chapter, that he has claimed (Hurlburt, 1990, 1993; Hurlburt & Heavey, 2006) does provide largely accurate descriptions of experience. The other author, Eric Schwitzgebel, without addressing DES in particular, has argued that introspective reports in general are greatly prone to error, even in what would seem the most favorable of cases (Schwitzgebel & Gordon, 2000; Schwitzgebel, 2002a-b, 2004, forthcoming, in preparation-b).

In this book, Russ and Eric confront each other directly and concretely on the adequacy and accuracy of introspective reports, using the particular reports of an actual subject as the starting point. Throughout the book, we will use the term “introspection” to refer only to the observation of particular instants of experience as they occur, or immediately thereafter. Sometimes, but not in this book, introspection refers to chewing over, musing, reflecting – to a certain type of self-oriented, retrospective or prospective contemplation. Our usage is quite specific: we wish to discuss whether, or to what extent, it is possible for people to report what is ongoing in their experience as it is currently happening.

1. The Origins of This Book

In April, 2002, Russ presented a paper titled “Describing inner experience: Not impossible but also not trivially easy” at an interdisciplinary conference in Tucson called Toward a Science of Consciousness. This paper, co-authored with Chris Heavey, criticized earlier attempts at introspection but argued that if one employed a proper method, it was possible to describe the features of inner experience (thoughts, images, feelings) with considerable accuracy. Russ had been working for decades developing just such a method.

At the same meeting, Eric presented a paper titled “Some reasons to distrust people’s judgments about their own conscious experiences.” In this paper, Eric argued that the introspection of emotion, sensory experience, imagery, and

thought – which together comprise much if not all of our experiential life – is unreliable, and that even in favorable circumstances of extended reflection on these aspects of our mental lives as they transpire, we often make gross mistakes regarding their basic features. Thus, he advocated a skeptical position that seemed to be considerably at odds with Russ’s cautious optimism. Eric was in the midst of publishing a series of papers defending this view (see the citations above).

Prior to the Tucson 2002 convention, we had never met, but the papers and our conversations showed that we shared a substantial intellectual history, despite Russ’s training in psychology and Eric’s in philosophy. We had both independently encountered the introspective literature on conscious experience and concluded that there was good reason for skepticism. We had both examined the methodology of the early introspectionist school and had written criticisms of those practices (Hurlburt, 1990; Schwitzgebel, 2002a). We had both written criticisms of the armchair introspections that underlie philosophical and psychological thought about consciousness (Hurlburt, 1990; Schwitzgebel & Gordon, 2000; Schwitzgebel, 2002a-b, 2003a-b).

However, despite these similarities, we had by 2002 reached opposing positions. Russ had responded to the methodological inadequacies of introspection by creating, in the late 1970s, a method of exploring inner, conscious experience that sought to avoid the pitfalls that had doomed earlier introspective attempts. This method came to be known as Descriptive Experience

Sampling (DES), and the project had culminated in two books (Hurlburt, 1990, 1993). Russ argued in those two books, as well as in the paper at Toward a Science of Consciousness, that his method solved enough of the methodological problems that DES could be taken as providing largely correct descriptions of inner experience (and perhaps other methods could as well). Russ will describe DES more completely in Chapter 2, but for now it is enough to know that DES uses a beeper to signal the subject to pay attention to the “inner experience” that was ongoing at the moment of the beep. Subsequently, the subject and investigator meet to discuss the details of such beeped moments.

Eric was not won over. Over the centuries, many people had made enthusiastic claims about the accuracy of their introspections, and most if not all of them had not proven credible. Why should he regard Russ’s claim about DES any differently? He agreed that the DES beeper did seem likely to overcome some of the difficulties involved in introspective report, but it appeared to aggravate other difficulties, and he thought it likely that, all things considered, substantial doubt would still be warranted. Yet at the same time, he had never examined the DES methodology closely.

We both recognized that it was crucial to determine whether it was possible to provide trustworthy accounts of conscious experience. The pressure was rising both in psychology and in philosophy to explore inner experience, consciousness, the phenomenology of thought and emotion. If Russ was right, then we should

redouble our efforts to explain to psychologists and philosophers how it is possible accurately to observe conscious experience. If Eric was right, even the most apparently credible reports of inner experience should not be accepted at face value without substantial independent support from non-introspective data.

We agreed that Eric would serve as a DES subject for a few days, right there at the Toward a Science of Consciousness conference. This would give Eric the opportunity to explore Russ's approach from the inside, to gain a more direct and intimate knowledge of it. Furthermore, it would provide a series of concrete occasions on which to discuss introspective methodology. We would thus move from the realm of general statements to the realm of concrete particulars. Eric's being a subject would turn Russ's method inside out, would let the fox explore the chicken coop from the inside. It would also test Eric's commitment to skepticism when his own experiential report was the one on the table.

We recognized that Eric was by no means a typical subject. He was open to participating in DES, but at the same time he had already thought extensively about the difficulties of introspection and was on the public record as a harsh critic of it. Thus, whereas most of Russ's subjects are simply trying to report the features of their experiences, Eric was trying both to report and at the same time to examine the limits of that reporting.

These interviews initiated a conversation that was continued by email over the next six months. We wrote each other at length, discussing the history of

introspection, examining Eric's experience as a subject, considering and reconsidering both of our skepticisms and Russ's explanations of how DES attempts to limit the risks inherent in earlier methods. That correspondence could be simplified as follows: We agreed that the history of introspection showed that most introspective reports were not to be trusted. But we disagreed about the extent to which the failure of earlier methods reflected general, ineliminable difficulties in introspection. Russ was optimistic. He argued that an interviewer like himself, carefully avoiding bias and focusing the interview on individual moments of experience, could often generate largely reliable reports. Eric remained relatively pessimistic, even when he himself was the subject.

2. Sampling with Melanie

To continue the conversation usefully, we felt that Eric needed more experience with interview techniques where his roles as skeptic and investigator wouldn't be complicated by his also simultaneously serving as the subject. So Russ proposed a new endeavor. We would jointly take the role of investigator and interview a naive subject, someone who had not previously been interviewed by Russ. In these interviews, Eric would be free to cross-examine the subject in whatever way he found useful, probing the subject's opinions about her sampled experiences without being confined to DES interviewing principles. For the role of subject, Russ found Melanie, a friend of a friend. Melanie had just graduated

from college with a joint degree in philosophy and psychology and was new in town, looking for a connection to the local psychology scene. Before coming to town, she had had no prior direct contact with either Russ or Eric or their views.

Until then, our conversations had been either about introspection in general (“should we trust introspective reports?”) or about Eric’s own (atypical) DES experience. The first kind of question was too broad. The second was confounded by Eric’s dual role and prior investigations. Now, however, the questions would be specific, concrete, and relatively straightforward: Should we believe Melanie’s report about her experience at 11:34:21? We could explore the question in any way we wished. To what extent would we agree, when faced with specific, individual reports? Would we disagree broadly about all the reports, or would the disagreement be concentrated on just a few reports, or a few aspects of them? We would be faced throughout with a concrete person, Melanie. It would not be adequate to say the impersonal, “I don’t believe introspective reports”. We would have to be concretely personal: “I don’t believe Melanie’s report”.

Our aims were also personal. Russ wanted candidly to expose his views to Eric, who seemed an open-minded but unsympathetic audience, to gain a skeptic’s perspective on his methodology, to refine his own skepticism, to reconsider how much skepticism about Melanie’s reports might indeed be warranted. Eric was exploring the limits of his skepticism, wavering between the radical pessimism about introspection with which he was flirting in his papers and

a more nuanced caution that admitted the possibility of progress and discovery. Our collaboration was intended to be a private conversation between the two of us, facilitated by Melanie's willingness to be questioned. We did not begin with the intention of making our conversations public.

After half a dozen sampling interviews with Melanie, spread over a month or so, we felt we had sufficient material to drive our discussion to the next phase, so we thanked Melanie for her participation and had the interviews transcribed by Sharon Jones-Forrester, one of Russ' students. The transcription was intended to serve as the basis for our continuing personal conversation about the trustworthiness of Melanie's reports in particular, and about DES reports and introspective reports in general. We independently read the transcripts and emailed comments about specific details to each other. We then replied to each other's comments and replied to those replies and so on, back and forth until we judged we had reached a point of diminishing returns. Over the course of the interviews and subsequent discussions, we gradually came to think that our concretely based considerations of the limits of skepticism, designed originally to be a private and candid conversation, might have value to others facing some of the same issues. Thus this book was born.

3. The Format of This Book

The sampling interviews that form the heart of this book were thus intended to be a personal confrontation between Russ and Eric. Because these interviews were real-time exchanges, we occasionally meandered, repeated ourselves, misunderstood each other, assumed shared knowledge unavailable to an outsider, phrased things poorly. In making these interviews available to the reader, therefore, we cut such portions of the transcripts; these cuts were never made unless we both agreed the remaining interview material stayed faithful to the original whole. We also slightly eased the remainder, removing some of the vocalized pauses and false starts, for example, again only where we jointly agreed to the fidelity of the alterations. Our aim in editing was to remove unnecessary distractions, thus focusing the remainder more sharply on what we took to be the issues of greatest general interest. We will make the complete, unaltered interview sound files and their transcripts available on the World Wide Web (see www.mit.edu/hurlburt-schwitzgebel.html) for those who wish to compare.

The heart of this book is therefore the transcripts of our interviews with Melanie along with 88 boxed discussions of issues raised in those interviews. To a large extent, those boxes are streamlined versions of the personal e-mail exchanges between Russ and Eric as we tried to hammer out our similar or differing takes on the adequacy of some particular aspect of our interviews with Melanie. We could have presented our views in the more traditional format for a co-authored pro-and-con book, each writing a discursive essay and reply.

However, we felt that the presentation of a verbatim transcript, with inserted comments and replies, would have substantial advantages over the more standard format. The transcript format forces the reader to begin with, and constantly confront, the particular. By contrast, most other discussions of introspective method begin with abstractions and general considerations, invoking particular instances, if at all, only selectively for the advancement of the author's more general thesis. While there is nothing inherently wrong with such an approach, we feel that there is something salutary in presenting the reader with randomly obtained particular reports, one at a time, prior to reaching general conclusions, with each report confronted on its own terms before proceeding to the next. Russ's and Eric's reactions and comments, both in the course of the original dialogue and in their later amplifications, may help the reader get some bearing on the kinds of doubts that may reasonably be raised and the resources available for responding to them.

Although this book looks wholly at the reports of one subject, Melanie, the reader will swiftly discover that the issues it raises are quite general. If the reader finds some of Melanie's claims about her experience to be believable and others to warrant doubt – as we think most readers will – this book invites consideration of what might drive these evaluations, and it offers different and sometimes conflicting suggestions on that topic. Temporarily replacing the factious and general debate about the trustworthiness of introspective reports with a personal

and particular look at the details of Melanie's reports will, we think, take us a long way toward honing or refining, trimming or amplifying, shifting or otherwise altering the skepticism that is desirable when encountering reports about conscious experience.

Thus this book is not a debate between opposing partisans, each trying to convince the other. Instead, it is a forthright collaboration between opposing partisans, each genuinely seeking to refine his own level of skepticism and to replace, as much as possible, partisanship with balanced critical judgment. The result, we hope, is an illumination of some of the major issues from two sides at once.

Our confrontation and dispute has also produced one potentially very useful byproduct: an examination, in unprecedented detail, of random moments of one person's experience. To the extent readers accept Melanie's reports, they will find a wealth of information about imagery, emotion, self-awareness, inner speech, and so forth, as experienced by a particular individual at particular moments in time. In the upcoming chapters we comment frequently on general issues pertaining to such experiences, such as the bearing of Melanie's reports on various psychological or philosophical theories, and the apparent similarities and differences between Melanie and other subjects we have read about or studied, including ourselves.

A Note to the Reader

Chapter Two presents the general rationale behind Russ' belief that satisfactory introspective methods may exist; Chapter Three presents Eric's general rationale for doubting such claims. We're ambivalent about including these chapters here. On the one hand, this background seems worth presenting, and this is the natural place. On the other hand, we've just argued for the value of starting with concrete instances instead of theoretical generalities, and on that logic it would be better for you to dive right into our interviews with Melanie beginning with Chapter Four. The interview transcripts don't assume knowledge of Chapters Two and Three, though you may have a fuller sense of what we're up to if you read these chapters first. We encourage you to follow your inclinations in this matter.

Chapter Two

Can There Be a Satisfactory Introspective Method?

Russ Hurlburt

Eric's and my interest in introspection stem from the same source: we agree that most attempts at the observation of inner experience have not been successful. But we have diverged in our response to that source. I have tried to capitalize on earlier introspective failures and build a better method than was used in the

previous attempts; so far, the best method I have discovered is Descriptive Experience Sampling (DES; Hurlburt, 1990, 1993; Hurlburt & Heavey, 2006). Eric has publicized the skeptic's position, criticizing all attempts at introspection without excluding new and perhaps better ones. In a nutshell, I want Eric to examine DES (or any other method that avoids earlier pitfalls) on its own merits without damning it by association with other not-so-sophisticated attempts; he wants me to recognize that history includes many enthusiastic supporters of introspective methods that have ultimately proven to be problematic. What makes this conversation engaging is the fact that we both recognize the legitimacy of the other's point of view, and are both pretty darn honestly trying to figure out the appropriate balance of these necessarily confrontational positions. Neither of us is trying to win the argument; both of us are as happy to hone the other's position as our own in the service of more adequately answering the Can we believe people's reports about their inner experience? question.

This chapter makes the case that there might well be introspective methods that deserve the scrutiny of even the most skeptical observer of introspection. I use DES as an example of such a method, not because it is the best method, but because it is the best method that I know of. I will show why it is reasonable to suppose that it is enough different from previous attempts to escape from the broad criticisms that have been leveled against introspection repeatedly over the last century. My attempt in this chapter is not to argue that DES actually does

provide accurate descriptions. Here I simply wish to demonstrate why I think it possible that introspective methods can be devised that avoid the earlier pitfalls.

This chapter is in many ways a reconstruction for the reader of the extended conversations that Eric and I had prior to deciding to sample with Melanie. The reader will recall from Chapter 1 that the outcome of those conversations was that Eric came to see that introspective methods might be able to be improved upon and to see DES as potentially interesting, sufficiently worthy of his skeptical attention to devote a substantial chunk of his professional time. In this chapter I have the same aim for the reader.

The chapter has the following organization: First, I survey a century of psychological science to discover what the characteristics of a good method might be. Then I describe DES, a method that embodies those characteristics. Then I describe ten reasons that DES reports might be considered plausible, and then describe a few compelling idiographic cases. [See Box 2.1 for a note about the terminology “inner experience.”]

Box 2.1. A note about terminology: “inner experience” or “conscious experience”?

Eric: Russ, you’ve called the subject matter of your work “inner experience.” I don’t like that term, because I think it favors experiences like thoughts and feelings (which are generally thought of as inner) over things like sensations (which are more outwardly directed). I prefer to call it “conscious experience” or

even just “experience.” I’m also concerned about how the phrase seems to build in the idea of the mind as interior and the world as external. I’m sympathetic with recent trends in cognitive science that reject a strict inner/outer division (sometimes called “embodied” views of the mind, or “externalism” or “contextualism”).

Russ: I agree that the “inner” in “inner experience” has the disadvantage that you point out – it does seem to favor thoughts over sensations. But DES subjects don’t seem to be affected by that; and it avoids the psychological and philosophical traditions in ways that I find highly desirable.

“Experience” (unmodified) can refer not only to inner experience, but also to “external” or “environmental” or “surrounding” experience, as in “I was affected by the space-shuttle-disaster experience” or “I took the job to get management experience.” Thus I think we need some kind of an adjective to indicate that “experience” refers to thoughts, feelings, sensations, and the like.

“Conscious experience” seems to awaken either (a) the contrast to the “unconscious” in Freud and many others’ sense awaken the existence of “states of consciousness”; or (b) the contrast to sleeping, dreaming, drug-altered, and so on experience.

“Attention” and “awareness” have an implication of a meta-awareness that I do not intend.

There is thus no nonproblematic terminology to refer to what might variously be called inner experience, conscious experience, experience, awareness, attention, or whatever. I have preferred “inner experience” as the being the least misleading, but it is far from perfect.

The good news is that in DES it simply doesn’t seem to matter what you call it, and therefore, I alternate quasi-randomly between all those terminologies in the attempt to distance myself from any one particular connotation. For example, in the set of interviews that we will display in Chapters 4-9, we use the term “inner experience” a total of 5 times, “experience” about 250 times, “awareness” about 100 times and “attention” about 70 times.

Eric: I’m not entirely convinced that it doesn’t matter what you call it, but I do agree that every terminology has shortcomings. “Conscious experience” also suggests a possible contrast to “unconscious experience” – a phrase that sounds incoherent to me. And does the phrase “conscious experience” invite the idea that we’re normally conscious of our experiences, in some self-observational way? Though some philosophers appear to endorse such a view (e.g., Rosenthal, 1986; Lycan, 1996), I’d prefer not to be committed to simply by the terminology. So maybe the phrase “inner experience” isn’t worse than any other. The reader will notice that I’ve reconciled myself to having it in the title of this book.

Thread: Loose language. Next: Box 4.1.

1. Toward a Better Introspective Method: 15 Guidelines from a Century of Science

The question this book is exploring is whether it is possible (or the extent to which it is possible) to obtain accurate descriptions of inner experience. Chris Heavey, Todd Seibert, and I (Hurlburt, Heavey, & Seibert, 2006) surveyed the last century or so of psychological science research to ascertain what that literature (most of it not introspective) has to say about the characteristics of a good introspective method. That paper extracted 15 guidelines for any good introspective method; this section paraphrases those guidelines; the reader is referred to the original article for amplification.

Guideline 1: The Stakes Are High. Bluntly stated, introspective methods failed and non-introspective methods came to dominate psychology largely due to introspection's failure. Should psychological science reawaken an interest in introspection without adequate discussion and improvement of introspective method, there may be an even more severe reaction (if that is possible) to a reawakened introspective era.

Guideline 2: Skepticism is Appropriate. Except perhaps for think-aloud procedures, all introspective procedures require memory to greater or lesser extent. [For a brief description of think-aloud procedures, see Box 2.2.] Psychological science robustly shows that human memory is prone to a variety of errors.

Box 2.2. Summary of sampling methods

Russ: For comparison purposes, here is a brief description of some current methods that attempt to explore inner experience.

Descriptive Experience Sampling (DES; Hurlburt, 1990) uses random beepers to trigger the qualitative description of experience. DES differs from all other sampling methods in that it is descriptive, not quantitative.

Thought sampling (or cognition sampling; Hurlburt, 1979) uses beepers to trigger subjects to fill out questionnaires. These questionnaires examine a variety of features of thought and mood.

The Experience Sampling Method (ESM; Larson & Csikszentmihalyi, 1983) is predominantly a quantitative methodology that collects standardized data about internal and external aspects of experience and situational/contextual variables. ESM differs from thought sampling primarily in its interest in situational variables and in the standardization of the questionnaires.

Ecological Momentary Assessment (EMA; Stone & Shiffman, 1994) is also a quantitative time-sampling method that differs from ESM in that it is not exclusively a random time sampling method; instead EMA sampling may occur at regular intervals (every hour, for example) or triggered by specific events (while jogging, for example).

Think-aloud procedures (Ericsson & Simon, 1980) ask subjects to verbalize their ongoing inner processes while performing some particular tasks (solving an

anagram, for example). Sometimes these methods are called “verbal protocol analyses.”

Articulated Thoughts in Simulated Situations Paradigm (ATSS; Davison, Robbins, & Johnson, 1983) is a verbal protocol analysis approach where subjects listen to audiotapes describing “stimulus scenarios” designed to elicit particular responses (social anxiety, for example). Subjects are to imagine actually being involved in the scenarios; immediately after hearing each scenario, they verbalize what they were thinking and feeling during the simulated situation.

Guideline 3: Introspect with Little Delay. It is well known that (a) if something is not encoded, it will likely not be recalled (Klatzky, 1975); that (b) meaningful chunks, not random details, are likely to be encoded (Bower, 1970); and that (c) this encoding must take place within a few seconds of the event. Because the features of inner experience that might be requested by introspection are not necessarily the meaningful portions of an event, those features are not likely to be encoded and therefore not likely to be reported accurately unless the introspection takes place very soon (within a few seconds) after the event.

Guideline 4: Target Specific, Concrete Episodes. People often engage in theory-guided recall when retrospectively characterizing their experiences (Pearson, Ross, & Dawes, 1992). Characterizations of experience over time are also likely to be distorted by features of the experiences themselves. For example,

Kahneman and colleagues (e.g., Kahneman, 1999; Redelmeier & Kahneman, 1996) have found that people asked to characterize pain over time do not perform some sort of average across actual events, but rather are unduly influenced by the peak level of pain and the current level of pain. Targeting specific moments of experience will minimize these biases.

Guideline 5: Keep the Target Experience Brief. There are “severe limitations on the amount of information that we are able to receive, process, and remember,” as Miller summarized in his highly influential “Seven, Plus or Minus Two” paper (Miller, 1956, p. 56). The introspectionists recognized such limitations a century ago. For example, Watt (1905), in his introspective analysis of problem solving, “fractionated” the problem-solving event into four parts, the preparation, the period prior to the presentation of the problem, the presentation of the problem, and the search for the solution, so that each part was no longer than a second or so. The implication is that the shorter the experience to be introspected, the better.

Guideline 6: Disturb the Experience as Little as Possible. James (1890/1981) famously suggested that it would be impossible to capture ongoing inner experience because the attempt to capture it would destroy the experience:

As a snow-flake crystal caught in the warm hand is no longer a crystal but a drop, so, instead of catching the feeling of relation moving to its term, we find we have caught some substantive thing, usually the last

word we were producing, statically taken, and with its function, tendency, and particular meaning in the sentence quite evaporated. The attempt at introspective analysis in these cases is in fact like ... trying to turn up the gas quickly enough to see how the darkness looks (p. 158).

John S. Mill suggested that it might be possible to capture ongoing experience through the medium of memory just after the experience has passed: “A fact may be studied through the medium of memory, not at the very moment of our perceiving it, but the moment after: and this is really the mode in which our best knowledge of our intellectual acts is generally acquired. We reflect on what we have been doing when the act is past, but when its impression in the memory is still fresh” (Mill, 1882/1961, p. 64). James and Mill were correct in pointing out we should try to disturb the targeted experience as little as possible

Guideline 7: Explore Natural Situations. External validity (Campbell & Stanley, 1963), “mundane realism” (Aronson & Carlsmith, 1968), and “ecological validity” (Brunswik, 1949) concerns about generalizability indicate that explorations should take place in the subject’s own natural environments.

Guideline 8: Minimize Demands. Explorations of private phenomena should seek to minimize “demand characteristics” (Orne, 1962) or the “Pygmalion Effect” (R. Rosenthal & Jacobson, 1968), employing double-blind testing when possible (R. Rosenthal, 1976) and scrupulously bracketing presuppositions when double-blind testing is not possible (as is often the case in DES).

Guideline 9: Terminology Is Problematic. B. F. Skinner observed that verbal behavior about private events may be impoverished because it is difficult for the verbal community to shape a person's speech about inner experience:

The verbal response "red" is established as a discriminative operant by a community which reinforces the response when it is made in the presence of red stimuli and not otherwise. This can easily be done if the community and the individual both have access to red stimuli. It cannot be done if either the individual or the community is color-blind. The latter case resembles that in which a verbal response is based upon a private event, where, by definition, common access by both parties is impossible. How does the community present or withhold reinforcement appropriately in order to bring such a response as "My tooth aches" under the control of appropriate stimulation? (1953, pp. 258-259, italics in original)

Thus Skinner established that talk about inner experience, such as "I was thinking...", "I am feeling...", "I am depressed," and so on, are not likely to have the same precision as talk about external events.

My DES colleagues and I have made this observation frequently in our sampling studies. For example, people often use the term thinking to mean something entirely non cognitive; others use the word feeling to refer to cognitive events (see Box 4.1). However, we have also observed that these people can

substantially improve or clarify their meanings if given repeated DES opportunities to try to speak accurately about their experience. Thus, we should recognize that some speakings cannot be adequately differentially reinforced, and we should therefore be very cautious in those arenas. However, where we can improve the differential reinforcement of speakings, we should do so. The implication is that methods must clarify to the extent possible precisely what is being described.

Guideline 10: Don't Ask Participants to Infer Causation. Nisbett and T. Wilson (1977), in a highly influential paper, reviewed research examining the attribution of causality and concluded that people often cannot describe “why” they behave/think the way they do. The moral seems clear: Avoid asking “why” questions.

Guideline 11: Abandon Armchair Observation. It follows from all that has gone before that casual observation about inner experience is not likely to yield scientifically valid results. Merely asking someone about their inner experience is simply not good enough. Furthermore, asking someone to perform armchair observations about their own experiences is problematic, even if that observation is done with careful instruction or by sophisticated observers:

I have conducted this brief examination of our introspective knowledge of visual imagery to promote the more general thesis that we can be, and often are, grossly mistaken about our own current conscious experiences

even in favourable circumstances of quiet attention.... We must abandon not only research paradigms in psychology and consciousness studies that depend too trustingly on introspection ... but also some of our ordinary assumptions about our knowledge of our own mental lives and what it's like to be ourselves. Human judgment about anything as fluid, changeable, skittish and chaotic as conscious experience is bound to error and confusion (Schwitzgebel, 2002, p. 50).

Guideline 12: Separate Report from Interpretation. Neuroscience has effectively used introspective reports throughout the past century. Reports of experience by those suffering from brain damage and disease have led neuroscience to an ever greater understanding of brain processing. Neuroscience has been successful because they have appropriately separated the introspective report from the interpretation of that report. It is the patient's job to provide the introspective reports, and the neurologist's job to provide the interpretation.

Guideline 13: Don't Require Too Much. Classical introspection observed many or most of the above guidelines and still Titchener's group disagreed vehemently with the Würzburg school about the existence of imageless thought: The Würzburgers thought they had discovered a new "imageless" element of thinking, whereas Titchener thought that images were present but very faint. Many observers see this lack of agreement as a primary cause of the fall of introspection a century ago (Misiak & Sexton, 1966; but see Danziger, 1980).

However, Monson and Hurlburt (1993; see also Hurlburt & Heavey, 2001) reviewed the introspectionist reports and found that Titchener and the Würzburgers substantially agreed about the phenomena in question, even though they disagreed about the interpretation of those observations. Had the introspections limited themselves to the careful description of phenomena, rather than trying to resolve an issue in their theory of mind, they would not have disagreed and introspection might not have been discredited.

Guideline 14: Value Prospective Research. Prospective designs offer the possibility of tapping a wide range of information relatively irrespective of theoretical perspective, collecting evidence that may or may not be related to some later question. Particularly at this early stage of the science of inner experience, this ability to allow the emergence of perhaps unexpected relationships or characteristics is especially important.

Guideline 15: Situate introspective observations in a nomological net. Those who would use introspective observations should explore the relationships of those observations to other kinds of research results.

These 15 guidelines highlight desirable features of any introspective method. There are doubtless other ways of slicing the century-of-psychological-research pie, which would yield a somewhat different set of guidelines. That is, I'm not claiming that this is the only nor the best set. Yet, it does seem to me that

this set is a reasonable summary of the desirable characteristics of introspective methods.

2. Descriptive Experience Sampling

Beginning in 1974, I began developing a method shaped by the thinking that is embodied in the guidelines we have just reviewed. That method is called Descriptive Experience Sampling (DES), which is my best shot at an accurate method for describing inner experience.

I emphasize that I do not think that DES is the ultimate method, only that it is the best method that I know of at this time. Should a method come along that I judge to be better than DES, I'd be happy to abandon DES in its favor. That is, I am personally, and this book is specifically, much more committed to the high quality study of inner experience than to the DES method in particular.

I have described DES in a variety of places (Hurlburt, 1990, 1993, 1997; Hurlburt & Heavey, 2001, 2002, 2004, 2006; Hurlburt & Akhter, in press) and will discuss its basics and rationale only fairly briefly here. Readers interested in more detail are referred to the works cited above. DES uses a random beeper in the subject's natural environments to signal the subject to pay attention to the experience that was ongoing at the moment of the beep. The subject then jots down notes about that now-immediately-past experience. The subject collects a half-dozen such beeped experiences and then meets with the investigator within

24 hours for an expositional interview, whose aim is to describe the experiences that were ongoing at each of the six beeped moments.

The beep/interview procedure is repeated over a number (usually between three and ten) sampling days. The “iterative” nature of the procedure interviews allows the subject’s observational and reporting skills to improve over the course of the several sampling days: Each day’s interview informs/refines/differentiates the next day’s observations, and in turn those newly refined observations inform/refine/differentiate the subsequent interviews (Hurlburt & Akhter, in press).

Occasionally critics of DES have disparagingly referred to the “magic beeper,” but whereas there is nothing magic about it, its characteristics are important (Hurlburt & Heavey, 2004, 2006):

- The beep is random. This makes it clear (a) that I and my subject are on equal footing with respect to the beep (that there is no manipulation involved); and (b) that I have no presuppositional expectations about what are important or unimportant occasions or events.
- The beep has a rapid onset or “rise time.” This makes it clear that I am interested in a precise moment, measured to the fraction of a second, perhaps. A vibrator of the type used in pagers is not adequate, for example.

- The beep should be easily detectable. A beep that is too loud will startle the subject, and the startle response will destroy the contents of experience. A beep that is too soft will trigger the subject's asking, "Is that the beep? Is that the beep? Yes! That's the beep!" but by now the experience that was occurring at the moment of onset of the beep may be lost.
- The beep is unambiguous. It means "Sample now!" and nothing else. Some critics have attempted to simulate the DES procedure by using, for example, a telephone ring as the signaling device. That doesn't work, because the subject's response must be, "That's a telephone ring, but I'm not supposed to answer the telephone, I'm supposed to pay attention to my experience." However, that response is likely to destroy the experience that was ongoing at the moment of the beep.
- The beep should be private. DES subjects generally use an earphone. If the beep is delivered through an external speaker, the subject must think about what she will say to anyone who might also have heard the beep, or must hasten to stop the beep so as not to annoy others. Either way, the ongoing experience has been lost.
- The beeper must be easily portable, so it can be easily used in the subject's natural environments.

The expositional interview asks essentially one and only one question: “What were you experiencing at the moment of the beep?” The object is to get as complete and detailed an answer to that question as possible, while at the same time avoiding confabulation. We want “the whole truth and nothing but the truth,” and the interview (in fact, the entire DES project) is aimed at that result. The interview is not structured, but instead asks that question over and over, in as many different forms as necessary, to focus the subject on the precise moment of the beep and nothing else. [See Box 2.3 for Russ’s comment about “the truth, the whole truth, and nothing but the truth.”]

Box 2.3. The truth, the whole truth, and nothing but the truth

Russ: Society often takes the statement “the truth, the whole truth, and nothing but the truth” to mean substantially less than its literal meaning. In the courtroom, “nothing but the truth” sometimes cynically means “anything that is not technically a lie.” Witnesses are routinely admonished not to provide the “whole truth” in the sense that they are instructed to answer only the question being asked and not to volunteer additional information, even if that additional information seems necessary to the understanding of the whole truth.

However, in DES, we mean “the truth, the whole truth, and nothing but the truth” to be taken as completely literally as possible. We give subjects the explicit choice: It’s okay not to tell us anything. But if you decide to tell us something

about a beeped experience, we would like you to tell the truth, the whole truth, and nothing but the truth as straightforwardly is possible. Our intention is the opposite of an attorney's. We want to discover the complete truth, not to hide behind a technical truth or show only one side of the truth. We want you to help us get to the heart of your experience, not to lead us away from it. We want you to help us discover as accurately as possible the details of your experience, not to blur them in the service of hiddenness. If we overlook something in what you've said, bring that to our attention. If we distort some feature of your experience, bring that to our attention. If our questions don't help you describe accurately your phenomena, help us to ask better questions. If you are unwilling to expose as accurately as possible the details of a beeped experience, then we would prefer not to talk about that experience at all.

Thread: Interview techniques. Next: Box 2.4.

By "the moment of the beep" we mean the last undisturbed moment before the beep begins – a millisecond before the beep. That is, we are not interested in the subject's reaction to the beep; we are not interested in what led up to the beep; we are not interested in what caused the experience; we are not interested in whether the experience is typical or rare. We are interested in the experience that was naturally ongoing at the millisecond before the beep began. We often use the

metaphor of a flash snapshot: we are interested in whatever the flash (beep onset) happens to catch.

Of course it would be naïve to think that we actually get to a perfectly undisturbed moment; the beep has to have been processed by the subject to identify the “last undisturbed moment.” One of the aims of this book is to get a sense of how undisturbed that moment is likely to be. Most subjects report that it seems that something like a “sensory store” for experience seems to exist, giving them time to “freeze” the experience and then to report it. But the believability of those reports is part of what is at issue in this book.

The DES interviewer tries to grasp the subject’s experience, as experienced by the subject. That requires suspending preconceptions about what the characteristics of the subject’s experience are, listening carefully to what the subject says, and trying to help the subject describe her own experience accurately. [See Box 2.4 for a discussion of the DES questioning technique.]

Box 2.4. Open-beginninged questions

Russ: DES questions are sometimes called “open-ended,” but I think it makes as much sense to call my questions “open-beginninged” as open-ended. An open-beginninged question is one that does not presume the content about which it asks.

“Tell me about your image” is an open-ended question, because it allows the respondent to elaborate as much or as little as possible about images. But its beginning is fixed: the question is about images, nothing else.

By contrast, “Tell me about your experience, if any, whatever it happened to be” is an open-beginninged (as well as an open-ended) question, because it allows the respondent to discuss images, speech, emotions, sensations – whatever was occurring at the moment of the beep, including none of the above or no experience at all.

The failure to appreciate the importance of open-beginninged questions has been, in my opinion, one of the major problems in the development of the science of inner experience, including most of the approaches described in Box 2.2. One researcher assumes that visual experience always exists, and asks about the characteristics of visual experience. Another researcher assumes that emotional experience always exists, and asks about the characteristics of emotional experience. Another researcher assumes that verbal experience always exists, and asks about the characteristics of verbal experience. Our DES research shows that there is no form of inner experience that comes anywhere close to always existing; if that’s true, the assumptions of all those researchers are incorrect, and their results therefore problematic.

It is possible to have a particular interest (say, in images) and still ask open-beginninged questions. You ask, in an open-beginninged way, what was going on

at this moment. If the experience happens to include an image, then you include that in your study. If the experience happens not to include an image, then you discard it. Such a study is, it seems to me, the only way to gain an accurate view of the way images are actually experienced. The argument (as I have heard it) that such a way is too inefficient actually proves my point. If it is inefficient, that must be because many moments do not include images, and to ask about images at those times must have been misguided.

Thread: Interview techniques. Previous: Box 2.3. Next: Box 4.3.

Thread: Richness. Next: Box 3.4.

We accept that Skinner was correct in his observation that people, including our DES subjects, are not differentiated observers or reporters of their inner experience (see Guideline 9 above). That is, subjects say many things about their experiences that are false or misleading, not because they wish to deceive but because in their life encounters until now they have not learned an adequate vocabulary to describe their experiences accurately; they have not learned to discriminate adequately between their actually occurring experience and their self-theories about their experiences; they have not learned to focus on one moment. The series of DES expositional interviews must therefore provide training on those important observational and reporting skills at the same time as it is acquiring reports of inner experience. Therefore, the first sampling day or two

(or more in some cases) is generally considered entirely training, not data gathering, and training continues past that time when necessary.

Thus the expositional interview consists of the subject's saying some things that are faithful and some that are misleading about her inner experience. The interviewer's task is to help the subject, over the course of sampling, say more and more of the faithful and less and less of the misleading. A metaphor that appeals to me is that I'm standing under the chute of a thresher with wheat and chaff pouring down. I try to grab the wheat and just ignore the chaff. (I actually don't know whether threshers work like that.) As the subject finds out that I'm very interested in the characteristics of particular moments and I'm not interested in the extra-sampling general statements, almost always there eventually becomes more wheat and less chaff, more talk about moments and less about general characterizations.

All this assumes that the subject is truly motivated to provide faithful descriptions of her inner experience. There may be some subjects who are motivated to lie, and probably nothing can be done about that. But DES does take seriously the attempt to enlist the subject's interest in faithful descriptions. First, we present ourselves as co-investigators: the subject has something (her experience) and we have something (the DES method), and together we can discover something that probably neither of us separately can do.

Second, we actually are, and present ourselves as being, genuinely interested in the faithful apprehension of her experience, as it occurred, with as few embellishments as possible. We demonstrate that genuine interest in a variety of ways: We question carefully to make sure we understand precisely what is being said; we encourage the careful focus on the precise moment of the beep by discouraging wandering away from that moment; we encourage the careful focus on the precise moment of the beep by discouraging speculation about what might have caused the currently experienced phenomenon; we consistently try to keep our own presuppositions out of the picture, maintaining a focus on the subject's experience as the subject experiences it; we let a random beeper choose the moment, rather than presume to know what moments are important.

Third, we protect the subject's privacy, telling her that we will not divulge her experiences until she explicitly agrees that we may do so; that she should feel free to discontinue sampling at any time without prejudice; that she should feel free to decline to discuss any experience for any reason (we have things that are none of her business and presume that she has things that are none of ours). We do ask that that if she wishes to decline to discuss an experience, that she tell us up front, and we will simply omit discussion of that beep entirely. Then if we do discuss a sample, we can delve as thoroughly as we desire (certainly the subject knows that she can change her mind and discontinue reporting or sampling at any time).

The result of all this is that the subject typically comes to realize that our aim is actually to apprehend the reality of the subject's experience, one moment at a time. Most subjects, I think, find that a very powerful and quite rare event: Someone really cares about my experience! Most subjects, I think, find it an unusual opportunity to be as honest as possible about personal experiences. Most subjects, I think, find it an opportunity to discover something about themselves, and the more accurate the better. [See Box 2.5 for a comment on Nisbett and Wilson's criticism of introspection.]

Box 2.5: Nisbett and Wilson's critique exempted DES, and indeed (contrary to myth) consciousness generally.

Russ: Nisbett and Wilson's 1977 criticism of introspection is so widely quoted as to require comment: "The accuracy of subjective reports is so poor as to suggest that any introspective access that may exist is not sufficient to produce generally correct or reliable reports" (Nisbett & T. Wilson, 1977, p. 233).

Critics of introspective-like methods have often taken the Nisbett & Wilson article to be an unconditional refutation of introspection in general. However, it is not widely known that Nisbett and Wilson, later in that same 1977 paper, recognized the possibility of accurate reports about inner experience:

We also wish to acknowledge that the studies do not suffice to show that people could never be accurate about the processes involved. To do so

would require ecologically meaningless but theoretically interesting procedures such as interrupting a process at the very moment it was occurring, alerting subjects to pay careful attention to their cognitive processes, coaching them in introspective procedures, and so on (p. 246, italics in original).

DES, as we have just seen, involves precisely “interrupting a process at the very moment it was occurring, alerting subjects to pay careful attention to their cognitive processes, coaching them in introspective procedures, and so on.” It is thus fair to say that Nisbett and Wilson, among the staunchest critics of introspection, agreed that methods like DES were at least “theoretically interesting” and might “be accurate about the processes involved.” (I think Nisbett and Wilson were mistaken about their further claim of ecological meaningfulness, but readers may judge for themselves by the end of this book.)

Eric: Let me go further, Russ, and point out that – despite the mythology that Nisbett and Wilson repudiated introspection generally (and the many citations of them to that effect) – they very explicitly emphasize that they mean only to challenge our introspective access to our own “cognitive processes” and not our “mental content.” In fact, they devote an entire section of their famous paper to making this point (“Confusion Between Content and Process,” p. 255-256). They grant, with what they take to be “almost all psychologists and philosophers,” that individuals have “a great storehouse of private knowledge... that can be known

with near certainty” (p. 255), including knowledge of our current sensations and emotions. They aim only to show that we have poor introspective knowledge of the processes leading up to – the causes of and influences on – such things as our judgments, decisions, feelings, and other conscious events. They do not claim that we are can be mistaken about what those judgments, decisions, feelings, and other conscious events themselves are. They challenge, for example, self-reports about why we prefer a particular pair of socks, not self-reports that we prefer them or self-reports of one’s current sensory experience (if any) in seeing the socks. Wilson continues to be quite explicit about this distinction in his more recent work, where he stresses our ignorance of “the adaptive unconscious,” as distinguished from consciousness (e.g., T. Wilson 2002, p. 17-18).

In general, psychologists have done a poor job separating skepticism about the self-reports of nonconscious processes, traits, behavioral dispositions, etc., from skepticism about self-reports of inner experience or consciousness; and when they do distinguish the two, it often turns out – as with Nisbett and Wilson – that they are only skeptical about the first.

Russ: I agree with all that.

3. Does DES-Apprehended Inner Experience Faithfully Mirror Inner Experience?

At the outset, I acknowledge that DES reports about inner experience mirror inner experience absolutely accurately only in rare cases if ever. So the issue is not whether the mirror is perfect, only whether it is scientifically adequate.

There are, it seems to me, two kinds of evidence for believing that DES reports might faithfully reflect inner experience. First, there are what I will call plausibility arguments – characteristics of the world and the method that lead me to think that accurate characterizations is the most plausible state of affairs. Second, and by far more important to me personally, are what I will call compelling idiographic observations – one-case-at-a-time observations of single individuals.

3.1. Ten Plausibility Arguments

Here are ten plausible reasons to believe that DES reports accurately reflect inner experience. None of them, by themselves, carry the day – one can argue against any of them. But all of them together are, to me, pretty persuasive. However, I emphasize that I do not think that arguments based on plausibility are ever an adequate foundation for science. They are important in that they clarify issues, but one person's plausibility is another's doubt. Science must be built on direct observation, not plausibility; that is why I believe that the compelling idiographic observations that I discuss in the next section are far more important than the plausibility arguments discussed here. I see these plausibility arguments

only as setting the stage for what I find to be the convincing idiographic observations.

. 1. The DES method is sophisticated. There is, historically, much good reason to doubt introspective reports. However, those introspective reports have been gathered in ways that I find seriously methodologically flawed. By following the guidelines and employing the characteristics described earlier in this chapter, DES in a sophisticated way may avoid those flaws. That of course doesn't imply that its introspections are necessarily successful, but it does open the possibility for more accuracy than earlier introspections.

2. Prospective DES subjects are skeptical too. Nearly all prospective DES subjects think DES will be difficult or impossible, but they find it easy once they actually engage in the DES procedure. It seems reasonable to suppose that the subjects' initial skepticism is somewhat similar to others' (perhaps the reader's) skepticism: it is based on armchair attempts at observing inner experience. But as I observed in 1997,

critics [should] not dismiss the descriptive experience sampling method on the basis of informal attempts at replicating the procedure. Informal sampling attempts such as asking oneself on occasion, "What am I thinking right now?" are nearly always discouraging, leading the typical critic to believe that he or she would be unable to perform the sampling task. However, I reported (Hurlburt, 1990, p. 269) that most subjects

find the actual task of responding to the random beep to be quite easy and unambiguous, stating that “unsuccessful [informal] attempts at thought sampling should not lead you to conclude that [descriptive experience] sampling ... is impossible; but rather should lead you to an appreciation of the relative delicacy of the method” (Hurlburt, 1997, p. 947).

The fact is that most subjects at the outset believe they will have a hard or impossible time capturing their inner experience, but over the first day or so of DES they become convinced that they can in fact capture their inner experience. This often-repeated trajectory from skepticism to acceptance based on their own directly-observed experience seems an argument against unrelenting skepticism.

3. DES subjects say they give accurate and complete reports. Despite the fact that I, in a skillfully repetitive way, give DES subjects the opportunity to say that there is more in their experience that they can't quite describe, they say the opposite – that they are giving pretty complete reports. They are convinced of that, and I am confident that that is not the result of my asking of leading questions.

4. Variability in within-subject reports implies their openness to a variety of experience. People often give quite different reports at different beeps, for example, inner speech at one beep, an image at another, unsymbolized thinking (the experience of thinking without words, images, or any other symbols;

Hurlburt 1990, 1993, 1997; Hurlburt & Heavey, 2006) at another, a combination of inner speech and feelings at another, and so on. This seems to indicate that people have a willingness and ability to report a variety of kinds of inner experience. It is therefore not the case that these subjects have a “favorite” kind of inner experience, or are “blind” to all other kinds of inner experience. (Certainly they might be blind to things they never report.)

Said another way, if one believes that reports of putative inner experience are purely artifactual, you’d expect the reports to be always the same. They are not.

There are other possible explanations for variability within subject’s reports; for example, that a subject has a self-theory of himself as highly variable, and therefore gives variable reports. However, in my experience most people think of variability in content, not variability in form. A person would have to be quite sophisticated about inner experience (to recognize the existence of unsymbolized thinking, for example), for self-theory to influence the form in this way.

5. Variability in between-subject reports implies my openness to a variety of experience. Different people have quite different patterns of responding. For example, one person reports nearly all inner speech; another reports nearly all images; another reports a mix of forms of inner experience. This seems to indicate that I, as one particular DES investigator, am open to a variety of experience. It is

therefore not the case that I have a “favorite” kind of inner experience, or am “blind” to all other kinds of inner experience.

6. The analogy from visual perception. The phenomenology of figure-ground perception has been well known at least since the Gestalt psychologists. Their work was largely in the visual realm; they showed that people spontaneously, seemingly immediately, create strongly felt patterns out of the visual arrays, and they proposed laws that govern such perception: proximity, similarity, closure, good continuation, and so on. Their main point was that people do not see everything that is available to be seen; they create, as part of the active perceptual process, a well-defined object to “see.”

It seems likely (and this is the way it is reported by DES subjects) that a similar process occurs across modalities. Thus, in much the same way that the faces disappear when I pay attention to the vase way of seeing the face/vase ambiguous figure, it seems reasonable to conclude that the sounds around me disappear when I pay attention to the visual, and that the visual disappears when I pay attention to the tactile, and so on. Certainly there are cases where I can pay attention to two or several aspects of the environment, but for most people most of the time, the number of such things is apparently small.

There are exceptions to that, but it is the exceptions that prove the rule. Some subjects do not “filter out” alternative modalities of alternative perceptions in the same modality. That indicates, it seems to me, that I am prepared to hear

such reports if they are given (that is, that I am not biased against them); because most people don't make such reports, even when thoroughly discussed in the expositional interviews, that I would be willing to hear more complex accounts if experience was more complex. (See Box 5.5.)

7. Compare the alternatives. An alternative that is sometimes advanced is that there is always visual experience ongoing. If the DES subject doesn't report it, it must therefore be neglected. I am not persuaded by that as a possibility, because the same argument can be made for other sensory modalities. Auditory experience must also always be ongoing because if someone says my name, I'm likely to hear it even if I'm paying attention to something else. Therefore, the argument goes, a piece of my awareness must have been auditory. Kinesthetic experience must also always be ongoing, because if I'm walking down the street and the pavement suddenly becomes spongy, I spontaneously adjust my gait. Therefore, the argument goes, a piece of my awareness must have involved the pavement feel and my body's reaction to it. And I see no reason to stop there: taste, smell, and so on are equally arguably always ongoing. So, on this model, I am always simultaneously experiencing many simultaneous multimodal things. I just don't think that's true. We certainly process input from multiple modalities at once, but most of that input does not become a recognizable part of our stream of experience, as the response of our immune system to invading bacteria or the expansion and contraction of our pupils as lighting conditions change do not

become recognizable parts of the stream of experience. See also Box 4.8 and the discussion of Eric's rich / thin study in Chapter Ten, Section 3, and Chapter Eleven, Section 2.1.

8. Subjects are not reluctant to report everything. As we saw in the previous section, part of the DES method is to impress on subjects that if a feature of their experience is none of my business, we shouldn't discuss that sample. I tell them that it is far easier if they just say up front, "This sample is none of your business," rather than try to disguise or hedge. I say that I will try to get a complete account, and if they're hiding something, we'll just go 'round and 'round; I won't feel a sense of completion.

Subjects occasionally do say "None of your business," which indicates that the message is heard. But they don't say it often, primarily (I think) because the beeped moments are usually pretty mundane.

I conclude that subjects are usually not reluctant to report as completely as possible; if they were, they'd say "None of your business" more often. In fact, subjects often report things that are embarrassing or run counter to their self concept [as indicated by verbal (You're sure this is confidential...") and non-verbal (blushing, stammering, etc.) evidence].

9. I myself am pretty good at bracketing presuppositions. I don't mean to be arrogant, or to single myself out, but the ability to bracket presuppositions probably has to be evaluated one person at a time. The evidence for the adequacy

of my own bracketing efficacy: (1) My reports vary dramatically from subject to subject, indicating that I am not “out looking for” my favorite characteristic. (2) I have reported many phenomena that were surprising to me myself (unsymbolized thinking, the absence of figure-ground phenomena, the absence of inner experience altogether). (3) I have worked at it and written about it. (4) I have been observed by at least one skeptic (Eric) who acknowledges that I seem to be pretty good at it (see Chapter Ten, introduction).

10. Leading the witness is less a problem with reports about actually occurring events than with general statements. Descriptive psychology is plagued by the problem of demand characteristics of the communications. I believe that the ability for demands to be effective in altering a subject’s perceptions diminishes as the situation becomes more concretely immediate. “See that stop sign there? It’s blue with white polka-dots” is not likely to be effective in the face of a red stop sign because your own immediate perception can refute it. DES tries to limit reports to immediately occurring events, and thus avoids much of the problem with demand.

3.2 Compelling Idiographic Observations

The plausibility arguments that I have just discussed suggest to me, in a weight-of-the-evidence kind of way, that the general answer to the question “Does DES-apprehended inner experience mirror inner experience?” is Yes. But I

recognize that someone else might advance ten plausible reasons to the contrary, and conclude that the answer is therefore No. There is, I believe, no clear-cut way out of this dueling plausible generalities scenario.

However, I believe that the general attempt to answer the “Does DES-apprehended inner experience faithfully mirror inner experience?” question is somewhat misguided. The fact is that, while I believe the ten factors I just listed do support the plausibility of the Yes answer, I myself am not persuaded by those arguments. The arguments I gave in the previous section tried to give an analytic answer to a question that may at heart require an inductive answer.

So let me recast the question to read, “Does the DES-apprehended inner experience of Allen faithfully mirror his inner experience? Does the DES-apprehended inner experience of Beatrice faithfully mirror her inner experience? Does the DES-apprehended inner experience of Chuck faithfully mirror his inner experience? Does the DES-apprehended inner experience of Dolores faithfully mirror her inner experience? and so on. And if the answer to many of those sub-questions is Yes, then we can perform the true inductive generalization and conclude that the DES-apprehended inner experience of (many) subjects mirror their inner experience.”

I have performed many DES investigations, and my answer to most of those inductive questions about them are “Yes, yes, yes, yes...and therefore Yes.” And, furthermore, that inductive series is capable of compelling me to believe the

final Yes in a way that the analytically plausible generality series of arguments simply cannot. I, as an individual, am pretty darn sure that DES-elicited reports of inner experience often or usually mirror actual inner experience; and I believe that I have been compelled to that belief by observing a series of single individuals for whom a contrary position seems bizarre. I'll cite two such cases here, both of which I've written about elsewhere. You'll get to judge for yourself the adequacy of the account of Melanie later on in this book.

3.2.1. The Case of Fran. In 1993 I reported the case of Fran, a woman who had been diagnosed as having a borderline personality disorder. In 1997 I discussed the “idiographic validity” of that case, arguing that my DES characterization of her inner experience reflected her actual inner experience:

“Fran” [was] a woman diagnosed as having a borderline personality (Hurlburt, 1993). Hurlburt described many salient characteristics of Fran’s inner experiences, of which I discuss three. First, Hurlburt reported that Fran’s inner experience was frequently populated by multiple (as many as five or ten) visual images, all occurring simultaneously and in the same “visual space” (that is, these images were not a side-by-side collage, but were instead all viewed straight ahead in a physically impossible overlaying that somehow did not provide any confusion for Fran herself). Fran’s case is thus an example

of the extreme complexity that inner experience can attain as reported by the descriptive experience sampling method. Such complexity cannot possibly be reported by any method other than sampling. For example, had Fran used a think-aloud technique, the most detailed non-sampling method, she simply could not have had time to report adequately one image, to say nothing of five or ten simultaneous images.

Second, Hurlburt (1993) reported that some of Fran's visual images (usually those with extremely negative content) often lasted for hours or days, nonstop, uninterrupted. (By contrast, the descriptive experience sampling method finds that images in healthy participants last for only a moment.) For example, Fran reported a visual and auditory image of her father "telling her off." In this image, Fran was seated at the dining room table. Her father was standing over her, pointing his finger at her, telling her she was "no good – a failure." Her mother was seen at the kitchen sink in the background looking over her shoulder at Fran. This image appeared in several successive samples, with the description being the same at each sample, and apparently continued uninterrupted during the time in between, for a total of at least several hours (pp. 202-205). This long-duration-image phenomenon might be considered impossible without sampling evidence.

Third, Hurlburt (1993) reported that Fran had no figure–ground phenomenon in either her inner image perception or her external perception—she took in an entire visual scene without focusing on any of its aspects. This conclusion was based on the fact that in repeated descriptive experience sampling interviews, Fran consistently denied the occurrence of phenomena associated with figure and ground: no part of an image appeared to be “closer” or “in better focus,” and when she shifted her gaze from one image (or external object) to another, she had no experience of “zeroing in” or of the previous center of attention “losing focus.”

A major question is of course whether Hurlburt’s (1993) descriptive experience sampling reports about Fran accurately reflect Fran’s inner experience: Fran was clearly the only person in a position to know that experience. Direct reliability studies are therefore impossible, so reliability must be indirectly inferred from validity considerations. Furthermore, one cannot apply standard validity-checking procedures (which intrinsically use across- group measures) to the idiographic observations of a single person; instead, one must infer validity idiographically, considering the unique characteristics of the particular description. I can identify five such idiographic validity considerations regarding the case of Fran:

First, the question of idiographic validity applies not to the descriptive sampling method per se but to the particular individuals who apply the method. In Fran's case, I was the investigator (Hurlburt, 1993). I might be expected to be a valid applier of the method because my previous descriptions of different people differ dramatically from each other, are sometimes surprising even to me myself, and are in agreement with other observers in those cases where more than one observer have sampled jointly (Hurlburt, 1993).

Second, the lack of figure-ground phenomenon in inner experience leads to an obvious but risky prediction that if Fran viewed the classical ambiguous figures such as the faces-vase or Jastrow's duck-rabbit, they would not "alternate" in her experience. I (Hurlburt, 1993) performed this informal validity experiment and found that Fran did in fact see both aspects of each drawing simultaneously with no alternation. A correct risky prediction can be taken as support for an underlying proposition (Popper, 1963) and therefore here as evidence of validity.

Third, I (Hurlburt, 1993) ruled out miscommunication, misunderstanding, or language deficit as alternative explanations of her failure to report figure-ground experience as follows. Fran asked to borrow the ambiguous figures to show to her coworkers, believing that I was mistaken about the existence of the alternation phenomenon. She

telephoned me a few hours later to report that to her surprise, her coworkers did in fact report the experience of alternation. In this conversation she gave an accurate description of her coworkers' alternating experiences but still denied that such alternation occurred for her. Thus it seemed clear that Fran understood what figure-ground phenomena are and was capable of describing them if they had existed for her.

Fourth, the descriptive experience sampling descriptions of Fran's inner experience provided plausible explanations of two characteristics of her external behavior. First, during Fran's discovery of her coworkers' figure-ground phenomenon, the coworkers came to realize, much to their surprise, that Fran could pay attention to many aspects of one thing or many different things simultaneously (such as her frequent multiple images), as had been discovered by descriptive experience sampling. The coworkers observed that this multiple-attention ability explained a trait that angered them all: They worked in a bank, and a frequent task was counting money. Each person would stand at a counter and count their own individual stacks of bills. Fran irritated her coworkers by repeatedly initiating conversations while counting, causing them to lose count. The simultaneous tasks of counting and conversing were impossible for her coworkers but simple for Fran. Thus, it seemed

clear to me that the multiple-experience characteristic of Fran's inner world had real ramifications in Fran's exterior everyday world.

The second sampling-based plausible explanation of external behavior came from Fran's psychotherapist. Before Fran had become involved in the sampling study, her psychotherapist had responded to her complaints of being preoccupied with negative thoughts by training her in thought substitution—a cognitive-therapeutic technique aimed at teaching her to think about something positive, based on the rationale that increasing her frequency of positive thoughts would lower the frequency of negative thoughts. However, that therapeutic intervention had been unsuccessful; sampling provided the plausible explanation that Fran was quite capable of thinking about something positive without ceasing to think about something negative.

Fifth, changes in external behavior were reflected in changes in inner experience. Near the end of sampling Fran experienced a remarkable improvement in her borderline symptoms: her exterior disorganization and chaotic psychological fragility vanished. Samples obtained after this improvement were now much less complex and now included the experience of figure–ground phenomena.

Taken together, these observations led me (Hurlburt, 1993) to conclude that the idiographic descriptions of Fran were indeed valid. If

their validity is at least tentatively accepted, they are extremely provocative; for example, to my knowledge, no reports of visual perception without figure–ground phenomenon appear in the perception literature, and no mention is made of the possible connection of the lack of figure–ground to psychopathology (Hurlburt, 1997, 946-947).

For reasons of focus in my 1993 and 1997 accounts, I did not include the following additional anecdote. Recall the conversation between Fran and her coworkers, when they discovered that Fran’s multiple-attention ability was the reason that she could count money and hold a simultaneous conversation. During that same conversation, Fran discovered that her coworkers had only one TV in their living rooms. Fran herself had three (didn’t everyone!?!), and watched them all simultaneously—not one after the other but all at the same time without switching her attention back and forth. She was surprised when her coworkers reported that they could not do the same thing! Furthermore, after the improvement in her borderline symptoms, she reported that she had lost this simultaneous-TV-watching ability, a substantial disappointment.

This case compels me to believe that my DES characterizations of Fran’s inner experience correspond in some important way to her actual inner experience. Sampling had putatively “discovered” a highly unusual phenomenon of Fran’s inner experience (no figure–ground phenomenon in image and external perception). This was “corroborated” by three highly unusual external

characteristics: (a) no alternation of ambiguous figures; (b) the ability to count and hold a conversation (the ability was actually stronger than that – she could count, participate in one conversation, and simultaneously listen without difficulty to one or more other simultaneous conversations); and (c) the existence of three TV's in her living room and the ability to watch all of them at the same time. And as if that weren't enough, when Fran's remarkable recovery occurred, both her inner experience and her external skills dramatically (literally overnight) lost their unusual characteristics.

It is therefore difficult for me to believe that the DES multiple-image characterization of Fran's experience was not substantially correct. How else can one explain these remarkable characteristics? It is of course possible that Fran was lying and inventing reports to seem "special," as those diagnosed with borderline personality sometimes do, was feeding off my interest and trying to confirm what she supposed to be my hypotheses. That doesn't seem likely to me, because she would have had to have been very psychologically sophisticated, because I myself had no hypotheses to confirm, and because I was quite skeptical about her reports. The most reasonable conclusion seems to me that sampling discovered and accurately reported important characteristics of her experience; substantially more sampling case studies and corroborating objective investigations will be required to be fully confident.

3.2.2. The Case of Robert. In 1994 Asperger Syndrome expert Uta Frith, her student Francesca Happé, and I reported the case of Robert, a 25-year-old man diagnosed as having Asperger's syndrome, a form of Autism where the level of intellectual functioning can be quite normal (Hurlburt, Happé, & Frith, 1994). Robert's IQ was 90, and he was quite able to perform the sampling task. Here are excerpts from our account:

The characteristics of all Robert's 16 samples were strikingly uniform. All 16 involved visual images, with no other aspects of experience reliably available to be reported—no feelings, no inner speech, no bodily sensations, etc. All Robert's images were seen clearly and in accurate colour, with the centre of the image being most clear and losing focus at the periphery, apparently exactly the same as his real-world perception...

Robert's samples were marked by the absence of any characteristics of inner experience except images. Except for the imagined sensation of a cat scratch on the back of his hand in one sample of an image of a cat, no samples included inner speech, feelings, bodily sensations, or other features of inner experience that have been reported by other subjects. Robert clearly had adequate ability to describe such features, and on occasion we specifically enquired whether such features were present, so as to rule out the possibility that they were simply being overlooked. Our

conclusion was that they simply did not occur to Robert as aspects of experience at any of the sampled moments.

Because the lack of non-image forms of inner experience was so striking, we structured [informal, non-DES]exercises during the interviews to explore the ways in which Robert experienced unambiguous strong bodily sensations. For example, with Robert's consent one of the authors (R.H.) leaned him forward and sideways to very tilted body positions; his inner experience (seeing a recalled image) remained constant, and a bodily awareness did not occur to him. In another such experiment, R.H. twisted the skin of Robert's wrist in opposite directions, creating what in most people would be a moderately painful experience. The wrist sensation did not create its own image or disturb the image that was present in his real-time inner experience: the image that he had been describing to us remained constant. Robert said he could feel the skin twisting but insisted it was not painful (Hurlburt, Happé & Frith, 1994, pp. 388-389).

On the basis of DES, Hurlburt, Happé, and Frith characterized Robert's inner experience as being almost always exclusively visual. The question we are dealing with here is whether that characterization is true. The informal experiments we performed were specific attempts to induce non-visual experience (leaning him to the side, applying painful twists to his wrist). Those manipulations

were only slightly effective: what in most individuals would immediately dominate experience became, apparently, only slightly or not at all a part of Robert's experience. That fact that images persisted despite explicit attempts to elicit non-image experience seems to corroborate the characterization of Robert's experience as being largely visual. Nonetheless, it is still possible that this visual focus is simply a characteristic of Robert's report, not of Robert's actual experience – maybe Robert felt pain but simply didn't have the vocabulary to report it.

However, there are other facts not reported by Hurlburt, Happé, and Frith (1994) which compel me to the view that his reports accurately mirrored his actual experience. During the session in which we had explored the painful wrist-twisting, Robert told the following anecdotes. When he was a child and lost his first baby tooth, his parents instructed him to put it under his pillow; the next morning, the tooth was gone, replaced by “a quid left by the tooth fairy.” Later that day, Robert took a pair of pliers and pulled out four more teeth! A more recent story occurred a few months before I had met Robert. He was in his apartment kitchen, and he smelled something burning. Looking around, he discovered it was his hand, which was accidentally resting on a hotplate!

Those remarkable and objectively corroboratable stories compel me to believe that pain does not figure Robert's his inner experience, just as sampling had shown. It is highly implausible that Robert's pain experience was similar to

that of most other people, and that only his pain reporting differed from the norm. Certainly such accounts do not verify that Robert's experience is visual, but they do lend credence to the accuracy of his no-pain description of our arm-twisting experiment. That, in turn, lends credence to the accuracy of his ongoing-undisturbed-visual-image portion of that description, and that in turn does support, in my view, his credibility as a reporter of ongoing imagery. I simply cannot accept the notion that we should treat Robert's DES accounts as "mere reports." By contrast, they are, it seems to me, substantially related to what Robert actually experienced.

By the way, neither absence of pain nor ubiquitous presence of images are known to be frequent characteristics of Asperger individuals, although there are similar reports by others (e.g. Grandin, 1995). It is therefore difficult to argue that we set out, knowingly or unwittingly, to look for those characteristics in Robert. Thus it seems difficult to avoid the conclusion that these were actual characteristics of Robert's experience. [See Box 2.6 for Eric's comment on the case of Robert.]

Box 2.6: How Compelling is the Case of Robert?

Eric: I find the case of Fran more compelling (assuming she's being honest) than the case of Robert. Your phraseology here confuses me; and if I'm confused, I worry that Robert may have been confused too. What do you mean when you say,

for example, that Robert had no “bodily awareness” as he was tilted forward and sideways? Was he completely ignorant of the fact that he was being tilted? You report no general disorder in his sense of balance. Or was Robert in some sense aware of being tilted (and able to report it?), though aware without “inner experience”? Do you mean to suggest that Robert navigated the world for the most part entirely non-consciously, no more having tactile or auditory experience than the rest of us have experience of our immune system or the growth of our fingernails? Or do you mean only something weaker?

How confident are you that Robert was alive to such distinctions as I’m asking about here and that he interpreted your questions as you intended them? Did you ask him about sensory visual experience, which you seem to assume he had, though it sounds like it’s not reported in any of his samples? Especially without verbatim transcripts to look at, I don’t feel I can give much credit to this strange material, confusingly presented, and therefore possibly born of confusion in the original interview. Pardon my frankness!

I do concede that you have some anecdotal evidence that comports nicely with Robert’s denial of pain. However, it should be noted that total incapacity to feel pain is a rare and serious disorder, typically accompanied by serious injury and deformity, due to insufficient self-care (Rosemberg, Marie & Kliemann 1994; Nagasako, Oaklander & Dworkin 2003). You don’t report this in the case of Robert.

Russ: I offer the case of Robert only to open the reader to possibilities, not as proof. You, Eric, discredit the report because it seems “strange” to you, and doesn’t match your presuppositions about our experiences of balance and touch, the relation between bodily injury and the experience of pain. I, too, found Robert’s reports initially rather strange. But as I argue repeatedly throughout the book, we must set aside (or “bracket”) such presuppositions when faced with DES reports (see Chapter Eleven, Section 1.7). Furthermore, I think the general rarity of pain insensitivity strongly supports my point. We discovered Robert’s pain insensitivity as a result of an exploration of inner experience that had nothing to do with pain. The ubiquity of Robert’s images led to the risky prediction of little or no bodily or pain experience, and as in the case of Fran, a correct risky prediction is supportive evidence.

I agree that our discussion of Robert is incomplete. One advantage of the present project, Eric, is that you can explore any similar presupposition-based doubts you have about Melanie’s reports as deeply as you like.

Thread: Bracketing presuppositions. Next: Box 3.3.

Thread: Human similarity and difference. Next: Box 3.3.

3.2.3. Discussion. Fran and Robert’s cases provide corroboration in instances where you might expect such corroboration the least. These were individuals with serious disorders, and yet their characterizations of their inner

experience seemed compellingly accurate. If seriously disturbed individuals can be faithfully accurate reporters of experience, healthy individuals should be able to be at least as accurate.

When I consider the many subjects I have examined with DES, some as dramatically compelling as Fran and Robert, I see little choice in believing that DES is about the exploration of inner experience, not merely about the reports of inner experience. To say that we are just examining reports of inner experience is, of course, true in a fundamental way that I can fully accept – everything has to be filtered through and understood in the context of reporting. But to say that we are just examining reports of inner experience seems substantially far-fetched, at the same level of far-fetchedness as to say that we are just examining perceptions of reality with nothing substantial implied. Just as I stop at red traffic lights because I believe in the substantial existence of the oncoming cars, I believe in the substantial existence of the inner experience that DES intends to describe. Just as I do not understand the nature of the reality of the oncoming cars, I also do not understand the nature of the reality of inner experience. But just as I do in fact get out of the way of oncoming traffic, I do treat inner experience as a fact.

It is possible to argue that the cases of Fran and Robert were exceptional – that's why I discussed them – or that perhaps my characterizations of Fran and Robert were somehow biased by my personal characteristics. The present book seeks to examine such reservations. We chose as a subject Melanie, who was not

thought to be particularly exceptional in the Fran or Robert sense – in fact, we knew little about her other than that she had been a successful college student. In the coming chapters we will expose the entire process, so that you can decide for yourself the extent to which the account of Melanie’s experience that our interviews provide should be taken to be accurate.

Chapter Three

Descartes Inverted

Eric Schwitzgebel

1. Some History

René Descartes argues, in his famous Meditations on First Philosophy (1641/1984), that the mind – including especially conscious thought and experience – is better known than the body. He supports this view with his dream doubt and demon doubt thought experiments, which are now standard fare in introduction to philosophy. You may think it certain that there's a book in front of you. But can you, really, be absolutely sure? Haven't you had the experience of dreaming that you were reading, falsely confident that you were awake? Or perhaps a demon is bent on deceiving you, and is thus feeding you false sensory impressions – or (in the more contemporary version) perhaps a genius neuroscientist from Alpha Centauri removed your brain last night while you slept, relocating it to a vat where it's being stimulated so as to mimic exactly the inputs it would receive from a normal waking day, including the reading of a hallucinatory book.

Of what can you be certain, according to this argument? Only that you exist, that you're thinking, that you have certain conscious experiences – a visual experience of blackish figures against a whitish background, a tactile experience

as if there were a book in your hands. Such facts about currently ongoing events in your mind, Descartes argues in his first two Meditations, are known indubitably and infallibly as no external fact could be. Indeed, later thinkers, such as Locke (1690/1975), and perhaps Descartes himself (on standard interpretations, like Russell's [1945]), argue that our knowledge of our own minds serves as the basis of our knowledge of the world outside. We apprehend our sensory experiences first; our judgments about the external world flow indirectly, derivatively, from a primary and more secure knowledge of our own consciousness. You know there's a book in your hands only because you know, antecedently and with greater certainty, that you're having visual and tactile experiences of a certain sort.

Philosophers of the mid-twentieth century, despite the skepticism about introspection that was commonplace in research psychology, commonly accepted something like Cartesian introspective infallibility or incorrigibility (e.g., Lewis, 1946; Ewing, 1951; Ayer, 1963; Shoemaker, 1963; Rorty, 1970), often supporting their claims by appeal to the example of pain. How could one be mistaken, or justifiably correctible by an outside observer, in the judgment that one is, or is not, in severe pain? A Saturday Night Live spoof (pointed out to me by Ned Block and Lex Newman) highlights the intuitive appeal of this idea through a mock commercial advertisement for a "home headache test" that requires users to draw and centrifuge blood to determine if they have a headache. The commercial ends as follows:

She: Oh, God, I'm in agony!

He: [looking at her test results] Honey, you don't have a headache.

She: Oh, thank God!

He: [looking at his test results] Neither do I! [They hug.]

Spokesman: The Home Headache Test. From Leland-Myers. Because what's worse – having a bad headache or not knowing if you have a bad headache?

The humor in this derives, of course, from the preposterous idea that one would need to look to outside sources to determine if one has a headache. Similarly, it seems difficult to imagine, when one is looking attentively, in good conditions, at a nearby bright red shirt – and consequently having a visual experience of red across a large swath of one's visual field – that one could possibly be mistaken in the judgment that one is experiencing the visual phenomenology of “redness” (though one might be incorrect in using the label “red” for that experience, due to a purely linguistic mistake). Seeking outside confirmation in such matters seems absurd.

In the last few decades infallibilism about current conscious experience has fallen out of favor among philosophers, or has at least been sharply curtailed. Philosophers such as Armstrong (1963) have argued, and others such as Shoemaker (1994) have conceded, that it's at least possible in principle to be mistaken about one's own current conscious experience, and that external sources

of evidence might sometimes justifiably override one's own introspective judgments. Churchland (1988) offers a typical thought experiment of the sort many have found convincing: You have been touched on the back nineteen times with a hot poker. On the twentieth trial, an ice-cube is surreptitiously substituted, and for a fleeting moment as it touches you, you mistakenly think you feel pain. Or: You're blindfolded and told you'll be tasting orange sherbet, but really it's lime (which tastes very similar in blind tests). For a moment – perhaps until you taste the actual orange – you erroneously think you experience an orangey taste. (See Box 3.1, however, for some concerns about these thought experiments.)

Box 3.1: Putative examples of introspective error

Eric: Though I'm sympathetic with Churchland's aims, his examples seem to me to invite easy objection: Perhaps you really do feel pain for a moment, or perhaps your judgment in that case isn't genuinely introspective? Perhaps you're not wrong about the taste experience itself but only about how to relate it to previous taste experiences and the world, or perhaps the suggestion of orange in combination with the lime-sherbet flavoring is enough to generate an actual orangey experience?

In fact, those who would defend infallibilism seem always to have an array of options when faced with a putative case of introspective error: Maybe the subject is simply using words differently than the rest of us, or maybe she is mistaken

only in her classification of an experience (an experience she perfectly well knows) as of a kind with other experiences she remembers, or maybe she really does experience what she says she experiences despite the behavioral or physiological evidence. To undermine all such potential responses generally requires more work (and more empirical work) than is possible in a typical philosopher's example; and definitive refutation may be out of reach – an infallibilist can always stomp his foot and insist that the behavioral and physiological evidence is misleading. I have devoted whole articles to developing plausible cases of error, in which I think typical infallibilist responses are at least a bit strained (the auditory experience of echolocation in Schwitzgebel & Gordon, 2000; visual imagery experience in Schwitzgebel, 2002a).

Although such examples are intended to undermine Cartesian infallibilism, embracing the possibility of error in some cases is quite compatible with holding – as the great majority of contemporary philosophical fallibilists do (including Churchland, 1985) – that introspection is generally a reliable process for coming to know one's own current experiences (e.g., Armstrong, 1963; Hill, 1991; Audi, 1993; Shoemaker, 1994; Dretske, 1995; Lycan, 1996; Goldman, 1997; Chalmers, 2003). Fallibilists almost always confine their examples of error to marginal cases, like those discussed above – matters of fine discrimination, or mistakes made only for a moment, or in circumstances of stress and distraction, or in

pathological or science-fiction cases, or only about the causes of our experiences rather than about the experiences themselves. Thus, the debate between fallibilists and infallibilists within philosophy has almost always been conducted under the umbrella assumption that introspection – that is, whatever process(es) drive our ordinary judgments about our currently ongoing or immediately past conscious experience – is a broadly trustworthy method, at least in favorable circumstances. No prominent philosopher has clearly and unequivocally put forward a case for thinking that we often go grossly wrong about our current conscious experiences, even in calm and ordinary circumstances of extended reflection (though see Box 3.2). Even psychologists suspicious of introspection have tended not to be entirely clear about whether they mean only to impugn self-reports of such things as motives, skills, traits, and so forth or whether they mean also to reject introspective reporting of current conscious experience; and some have explicitly cordoned off the latter sort of reports from their doubts – including the much-cited Nisbett and Wilson (1977; see Box 2.5).

Box 3.2: Dennett on introspection of current conscious experience

Eric: Daniel Dennett's work, especially his 1991 book, Consciousness Explained, is often read as arguing for the possibility of pervasive and radical mistakes about conscious experience. However, I find Dennett far from clear on the point. Sometimes he seemingly takes himself to be providing examples of gross

introspective error about ongoing conscious experience (e.g., regarding the level of detail in visual experience, in Ch. 11); but elsewhere, he compares our authority in reporting our experiences to the authority an author has over his fictional creations – which seems to imply that we could no more go wrong in our reports than Arthur Conan Doyle could go wrong in reporting the color of Holmes’s easy chair (e.g., 1991, p. 81 & 365-366). Dennett also asserts that we can come “close” to infallibility when charitably interpreted (Dennett, 2002, p. 13, 16), and he allies himself explicitly with Rorty (1970) and other “incorrigibilists” who argue that we can never be justified in overturning a sincere introspective report on the basis of outside evidence (Dennett, 2000, 2002). I doubt a coherent view can be made of all this: See Schwitzgebel (forthcoming-a).

2. My Point of View

This seems to me an odd state of affairs. Why should philosophers – an ornery lot who rarely reach general consensus about anything – almost universally regard the introspection of one’s ongoing phenomenology, or stream of experience, as trustworthy and reliable? People aren’t especially trustworthy and reliable, most think, in reporting the real grounds of their judgments and decisions (e.g., why they chose a particular pair of socks: Nisbett & T. Wilson, 1977) or in reporting their implicit attitudes (e.g., about the characteristics of other races:

Gawronski & Bodenhausen, 2006). Is it obvious that the introspection of current conscious experience deserves better epistemic credentials?

Autobiographically, my interest in this issue was first aroused through my work in Alison Gopnik's developmental psychology laboratory in the 1990s (while I was simultaneously a philosophy Ph.D. student under Elisabeth Lloyd and John Searle). Gopnik, at Berkeley, and the Flavells at Stanford, were exploring the extent to which children could be mistaken about their own attitudes and experiences (e.g., in Gopnik, 1993a-b; Flavell, Green & Flavell, 1993, 1995, 1998, 2000; Flavell, Green, Flavell & Grossman, 1997; Flavell, Flavell & Green, 2001; Flavell & Flavell, 2004). It seemed from their research that young children – 4- and 5-year-olds – could be vastly mistaken, denying the presence of beliefs just expressed, holding that people could go for days without a single thought, grossly misreporting their own ongoing or immediately past reflections and expectations. One typical result from Flavell is this:

The experimenter showed the child a library-type bell and then hid it under the testing table. She said, "In just a few seconds, I'll ring it." After a 4-sec delay she rang the bell and said, "OK, in just a few seconds I'll ring it again." After another 4-sec delay she rang the bell and reiterated, "OK, one more time. In just a few seconds I'll ring it again." This time she failed to ring the bell and after a 10-12 sec delay instead asked... "Are you wondering or thinking about anything right now or not?" If the answer was "yes," the experimenter asked, "What are you wondering or thinking about?" (1995, p. 72-73).

Only 38% of five-year-olds responded that they were wondering or thinking about anything having to do with the bell or the experimenter's striking it, and 44% said

they weren't wondering or thinking about anything at all. Flavell et al. assume that most of the children must have been thinking about the bell (and that virtually all must have been thinking about something), so they interpret the majority as mistaken in their replies. Flavell et al. perform a variety of similar experiments, varying their methodology and sometimes going to great lengths to explain to the children what "thinking" means. I don't view the results as decisive – surely, children may fail to understand the tasks or the language, despite the Flavells' near-heroic efforts in various versions of these experiments. Still, I'm drawn to the Flavells' overall conclusion that young children know very little about the processes of thinking and are remarkably poor introspectors –and, as Gopnik suggests, if children struggle enormously with understanding the task or the language, that itself already tells you something. Unfortunately, I know of no attempts to replicate Flavell's work on this outside of his laboratory. (For a concern about Flavell's interpretation of his results, see Box 3.3.)

Box 3.3: How should we interpret Flavell's children?

Russ: An alternative explanation of Flavell's results is that many five-year-olds simply do not have thoughts, and that therefore they are not mistaken at all, let alone "vastly mistaken." I fear that the Flavells and Eric reject this possibility because it is too different from their own experience. One of my recurrent themes in this book is the desirability of bracketing presuppositions about others'

experience, in particular resisting the temptation to believe that everyone's experience is like one's own. We shall see that I think that bracketing presuppositions is of vital importance to the advancement of a science of inner experience, while Eric thinks my dismissal of prior research is too cavalier. See my extended treatment of the Flavell studies in Chapter Eleven, Section 1.7.8.

Thread: Bracketing presuppositions. Previous: Box 2.6. Next: Box 4.6.

Thread: Human similarity and difference. Previous: Box 2.6. Next: Box 4.1.

If Gopnik and Flavell are right, young children are far worse at introspecting their experiences than they are at perceiving outward objects. Now, are we to suppose that this situation reverses itself by adulthood? Are ordinary adults, unlike young children, at least as accurate in their judgments about their stream of experience as they are about the physical and social world around them? Do they largely avoid gross errors? I was surprised to find that little empirical work has been done on the question.

It's a difficult question. We can't directly observe someone else's experience. We may be tempted to infer error or accuracy in particular cases, but with no dependable, general theory to hand about the relationship between conscious experiences and publicly observable brain states or behavior – no theory more dependable than the subjective reports themselves, anyway – we're necessarily on shaky ground in the many cases where a person's judgment about

her experience diverges from what one might think would be her experience based on other evidence. Witness my criticisms of Churchland's examples (in Box 3.1) and Russ's criticism of Flavell (in Box 3.3). We have too much liberty to posit whatever conscious experience best matches our theory, whether we wish to shield the subject's report or refute it. How can this be good science?

Still, I think the issue admits of exploration through a careful blend of introspection (by researchers and independent subjects), theoretical reflection, and attention to outwardly observable empirical facts. In the face of ambiguous evidence, we can at least weigh the plausibilities. The remainder of this chapter will provide a taste of the kinds of consideration I find persuasive.

Consider, for example, not the infallibilists' favorite cases of severe pain and foveal red, but rather the experience of visual imagery. Create a visual image right now, if you can: for example, an image of your breakfast table as you sat down to it this morning (following Galton, 1880), or an image of your house as seen from the street. Now consider the following questions. Can you, indeed, form and retain such an image? If so, how stable is that image? Does it fluctuate as you think about different aspects of the scene, as your attention waxes and wanes, or does it stay relatively constant? Does it have a focal center and a periphery, or is everything equally present at once? How detailed is it? Are objects to the side well articulated in your image before you specifically think about them (and "fill them in," as it were)? If you concentrate on one object in the

image (assuming you can do so), how are the other objects experienced? Do they “fade” in some way? What is that fading like? Is everything present in color all at once, or does some of the image have indeterminate color? How is that indeterminate color, if there is any, experienced – as black and white, or gray, or in some other way?

Most people of whom I ask such questions – although these are questions about major features of presumably (but see Box 5.1) ongoing conscious experience – eventually feel confusion or uncertainty. Certainly I do. I don’t think this confusion is merely about words and theories, about how best to describe a patently obvious visual imagery experience. Rather, it extends, to some degree, to the experience itself. It’s not absolutely certain what our visual imagery experience is like. And consequently, it shouldn’t be surprising if some people, at least, occasionally go wrong about it. Furthermore, subjective reports about visual imagery experience vary widely – in apparently normal people, all the way from the complete denial of visual imagery experiences to descriptions of visual imagery as detailed as ordinary visual experience of outward objects or even more detailed (Galton, 1880, 1907, still provides the most useful compendium). Yet psychologists have generally failed, over a long history of hundreds of experiments, to find consistent correlations between subjective reports of visual imagery experience and performance on the types of tasks commonly thought to involve visual imagery, such as mental rotation tasks, visual memory tasks, and

tests of visual creativity (see McKelvie, 1995, and critics, especially Slee, 1995; Schwitzgebel, 2002a).

Or consider emotional experience. What's your emotional experience right now? Do you even have any? Try to conjure some if you think you don't. Is it completely obvious to you what it is? Even if you're fairly confident in giving it a general label, how much do you know about its particular experiential character? Does introspection reveal its details, its somatic manifestations in experience (if any), its full phenomenological structure, as clearly and certainly as visual inspection reveals the contents of your desktop? Does this seem to you like a topic on which you could not possibly go wrong? Speaking for myself, I'm tempted to suggest that my wife reads my emotional phenomenology better in my face and posture than I do in my own introspection.

These reflections are meant only to be suggestive. Not every reader will find the same uncertainty and doubt that I do. Think of these reflections not so much as arguments but rather invitations to a point of view too little defended by philosophers – as, similarly, infallibilists' reflections on the apparent impossibility of being mistaken about one's headaches or about vivid color experiences are really more invitations to embrace infallibilism than demonstrations of its truth. Or, if you like, for the purposes of this chapter, think of these reflections as merely an expression of my point of view as it stood prior to our interviews with Melanie.

3. Sources of Introspective Error

The introspection of conscious experience is difficult for several reasons. First, experience is fleeting and changeable – or so it seems to me right now as I reflect, introspectively, upon it. The page before me, as I reread these paragraphs, is relatively steady, but my visual experience as I look at the page is in constant flux. As my eyes move, the portion that's clear, the portion that's hazy, constantly changes. I blink, I glance away, I change my focus, and my experience shifts. My eyes slowly adapt to the black and white of the page, to the contrast with the surrounding desk, to the changing light as the sun goes behind a cloud. I parse some bit of the page into familiar words as my eye scans down it; I form a visual image, reflecting the content of the discussion; my attention wanders. All this, it seems, affects my visual experience. Consider your own experience as you read this paragraph. The text in your hands changes not a whit, but your visual phenomenology won't stay still a second, will it? (Or will it?) The same is true, I'm inclined to think, for our auditory experience, emotional experience, somatic experience, conscious thought and imagery, taste, and so on: Even when the outside environment is relatively steady, the stream of experience flies swiftly. It won't hold still to be examined. (For some of Russ's concerns about my claims here, see Box 3.4.)

Box 3.4. On experience while reading

Russ: Your description of the experience of reading and vision, Eric, is the state of the art view held by philosophers and psychologists, but I think that view is mistaken. The experience of reading is never or at most rarely like that. Certainly the eyes move and the retinal representation of words actually becomes more or less distinct as it is closer or farther from the fovea. But I think that level of processing takes place always (or almost always) outside experience. The actual experience while reading, I think, has little or nothing to do with such things.

The reason that we can disagree so dramatically about such a fundamentally important issue is that science has done a horrible job of exploring experience. That's why I think this book is so vitally important. If it's possible, as I think it is, to obtain accurate observations of the experience while reading, then science should do so, and then replace the state-of-the-art (and your) view with a more accurate one based on carefully obtained experiential data. On the other hand, if it's not possible to obtain accurate observations of experience, then science (and I) should refrain from embarking on a second round of introspection.

Eric: I agree that we know very little about the experience of reading and about naturally occurring visual experience generally – an amazing lacuna, really! Perhaps because I incline toward a relatively “rich” view of experience (see Box 4.8 and Chapter Ten, section 3), I've assumed that the reader has a visual experience as she reads the text, and that this experience changes with the position

of her eyes and the overall state of her visual system. Maybe I could be talked out of this, though, by clever enough arguments and experiments. My main point, of course, is that we shouldn't simply trust our own introspective judgments about our experience – and I don't exempt my own from doubt.

For more on the experience of reading, see Box 5.3.

Thread: Richness. Previous: Box 2.4. Next: Box 4.8.

Thread: Sensory experience. Next: Box 4.8.

Second, we're not in the habit of attending introspectively to experience. Generally, we care more about physical objects in the world around us, and about our and others' situation and prospects, than about our conscious experience, except when that experience is acutely negative, as with the onset of severe pain. This may seem strange, given the importance we sometimes claim for "happiness," which we generally construe as bound up with, or even reducible to, emotional experience – but despite the lip service, few people make a real study of their phenomenology. We spend much more time thinking about, and have much subtler an appreciation of, our outward occupations and hobbies. And when we do "introspect," we tend to think about such things as our motives for past actions, our personality traits and character, our desires for the future. This is not, in my view (or Russ's), introspection strictly speaking; but call it what you like, it's not the sort of introspective attention to currently ongoing (or immediately

past) conscious experience that lies at the heart of consciousness studies. Introspective attention to experience is hardly a habitual practice for most, perhaps any, of us (except maybe a few dedicated meditators of a certain sort). If accurate introspection requires a degree of skill, as I suspect it does, in most people the skill is uncultivated. Furthermore, relatedly, but perhaps to some extent independently as well, experience is extremely difficult to remember: Generally what we remember are outward objects and events – or, rather, outward objects and events as interpreted, and possibly misperceived, by us – not our stream of experience as we witness those objects and events. We remember, usually – usually – that the boss said the work wasn't up to snuff, not that our visual experience as he said it was such-and-such or that we felt some particular sinking feeling in the stomach afterward. These conscious experiences fade like dreams in the morning unless, as with dreams, we fix them in mind with deliberate attention within a very short space.

Third, in part due to our disinterest in conscious experience, the concepts and categories available to characterize it are limited and derivative. Most language for sensory experience is adapted from the language we use to describe outward objects of sensation. Objects are red or square or salty or rough, and usually when we use the words “red” and “square” and “salty” and “rough,” we are referring to the properties of outward objects; but derivatively we also use those words to describe the sensory experiences normally produced by such

objects. That's fine as far as it goes, but it's prone to invite confusion between the properties of objects and the properties of experiences of those objects. The practitioners of certain specialties – for example, wine tasting and sound engineering – have refined language to discuss sensory experience, but even here our conceptual categories are only rough tools for describing the overall experience. And, anyway, isn't the gustatory experience of eating a burrito as complex as that of tasting a mature wine, and the auditory experience of sitting in a restaurant as complex as that of hearing a well-played violin? We almost completely lack the concepts and competencies that would allow us to parse and think about, talk about and remember, this complexity. (For more on this point, see Schwitzgebel, 2004.)

Fourth, the introspection of current experience requires attention to (or thought about) that experience, at least in the methodologically central case of deliberately introspecting with the aim of producing an accurate report. Problematic interference between the conscious experience and the introspective activity thus threatens. Philosophers and psychologists going back at least to Comte (1830) have complained that the act of introspection either alters or destroys the target experience, making accurate report impossible (see also Russ's Chapter Two, section 1, guideline 6). Much of experience is skittish – as soon as we think about it, it flits away. Suppose you reflect on the emotional experience of simple, reactive anger, or the auditory experience of hearing someone speak. Mightn't the self-reflective versions of those experiences – those experiences as they present themselves to concurrent introspection – be quite different from those experiences as they normally occur in the unselfconscious flow of daily life? A number of psychologists have attempted to remedy this difficulty by recommending immediate retrospection, or recall, of past experience rather than concurrent introspection as the primary method (e.g., James, 1890/1981; Farthing, 1992). However, deliberately poisoning oneself in advance to report

something retrospectively may also interfere with the process to be reported; and if one only reports experiences sufficiently salient and interesting to produce immediate spontaneous retrospection, one will get a very biased sample. Furthermore, retrospection is likely to aggravate the final problem I'll discuss here, namely:

Fifth, reports of experience are apt to be considerably influenced, and distorted, by pre-existing theories, opinions, and biases, both cultural and personal, as well as situational demands. The gravity of this problem is difficult to estimate, but in my opinion it is extreme (and considerably larger than the influence of bias and preconception now generally recognized to permeate science as a whole). Given the changeability and skittishness of experience, and our poor tools and limited practice in conceptualizing and remembering it, we lean especially heavily on implicit assumptions and indirect evidence in reaching our introspective and immediately retrospective judgments. One major source of such error is what the introspective psychologist E. B. Titchener called "stimulus error" (Titchener, 1901-1905, 1912; Boring 1921): We know what the world, or a particular stimulus, is like (we know for example that we are seeing a uniformly colored red object), and we are apt to infer that our experience has the properties one might naively expect such a stimulus to produce (e.g., a visual experience of uniform "redness"). We're much better accustomed to attend to the world than to our experience, and the difference between sensory attention to outside objects and introspective attention to the sensory experience of those objects is a subtle one; so the former is apt to substitute for the latter (for a related point, see

Dretske, 1995; Tye 2003; Siewert, 2004; Schwitzgebel, 2005; Stoljar, forthcoming). Even when experience isn't so easily traceable to an outside object, I'm inclined to think our theories can profoundly affect our reports. If we think images must be like pictures, we're more apt to instill reports of imagery with picture-like qualities than if we don't hold that view (see Box 5.2; Schwitzgebel, forthcoming-b). If we think cognition takes place in the brain, we're more apt to locate our cognitive phenomenology there than if we think it takes place in the heart (see Box 7.12). If we think that memories must be imagistic, we're more apt than those who don't think so to report memory images.

Thus, in my view, Descartes got things almost exactly backwards. The outside world of stable objects, people, and events – the world we spend most of our time thinking about – is what we know most directly and certainly. The “inner world” of conscious experience is reflected on only rarely, and known for the most part only poorly. I'm practically certain there's a tissue box here before me, and I know quite a bit about its physical details; however, I'm much less certain of my visual experience as I look at that tissue box, except at the crudest level. Furthermore, what I do know, or suspect, about my visual experience is grounded to a considerable extent in my knowledge of the properties of the box itself. My judgments about the box's shape, color, and other visible features in large part (though not exclusively) drive my judgments about my visual experience of

shape, color, and so forth – not, as many philosophers inspired by Descartes have suggested – the other way around.

4. Our Difficult Situation

Going into this collaboration with Russ, I didn't feel that introspection, or ordinary naive reflection on ongoing or immediately past "inner experience," was completely hopeless as a method for learning about consciousness, but I did – and still do – feel it must be treated with enormous caution. We cannot blithely assume that even the most credible-seeming introspective reports are likely to be true.

Yet, despite its untrustworthiness, introspection must be given a central role in the study of consciousness. Without introspection, we might not even know that we are conscious in the relevant sense – that a stream of phenomenology accompanies our outwardly visible behavior. Behavioral and physiological measures alone tell us nothing about consciousness unless it's established that those measures correlate with conscious experience; and introspection is the most straightforward way to establish such correlations. All tools for understanding consciousness are problematic in their application. Perhaps this is why consciousness studies has been so slow to find firm scientific footing. It's not as though in the face of the unreliability of introspection we can substitute some

simple set of behavioral or physiological measures that will consistently generate accurate and detailed answers to questions about our phenomenology.

Our situation, I think, is in some ways analogous to that of a foreign intelligence agency that must depend for its information on a network of unreliable, double-crossing spies. Just as the reports of spies can to some extent be corroborated or cast into doubt by such independent means as satellite photos and bank records (which by themselves may say little), so also can introspective reports be to some extent checked against behavioral, physiological, and cognitive measures; and just as consistency or inconsistency between the reports of independent spies provides at least prima facie reason to accept or doubt the reports, so also consistency or contrast between independent introspective reports, when there is no reason to suspect corresponding differences in conscious experience, may justify tentative acceptance or rejection of the reports. Given the unreliability of naive introspection, we need such methods of confirmation, shaky as they are – and perhaps introspective training as well (Schwitzgebel 2004) – before we can be truly justified in accepting introspective data.

It seemed to me at the beginning of this project that Russ's DES methodology did not change this basic situation. Although the methodology seemed likely to alleviate my fourth concern above, regarding the division of attention and selectiveness of retrospection, it seemed simultaneously to aggravate the fifth concern: distortion by pre-existing opinions and situational demands. By two seconds after the beep, it seemed to me from my own sampling, and thus still at the beginning of the articulation and categorization of the experience, all but the grossest and most salient features of the experience were gone from memory. This left a large opening, it seemed to me, for biased or theory-guided reconstruction, an opening that only expanded as time progressed. (This remains my opinion after having completed this project: See Chapter Ten, sections 4-5.)

How large a problem is this? Russ and I agree in general about the difficulties and potential sources of error in reporting experience, but we disagree about how to weigh them and the extent to which they can be alleviated with adequate care. All reasonable people must, I think, stand somewhere between thinking such reports, from experience sampling or other sources, are absolutely infallible in every detail and thinking that only sheer accident could ever allow a person accurately to report basic and apparently obvious features of her own experience. Russ and I hope the reader will find the following dialogues useful in assessing how much skepticism is warranted and the extent to which it may be overcome by careful questioning.