

Preface

I have two aims in this book. First, I aim to persuade you that people in general know very little about what might seem to be very obvious features of their stream of conscious experience – where by “conscious experience” I mean things like sensory experience, visual imagery, inner speech, emotional experience, and the whole variety of subjective phenomena that constitute what we sometimes think of as our “inner” lives. Second, I aim persuade you that you yourself know very little about such matters, that the thoughts and judgments you arrive at when you consider such subjective phenomena likely merit doubt and suspicion. Obviously, these two aims intertwine and support each other. My scholarly emphasis is on the general claim, as it probably should be, but I’ve found in the course of writing and revising these chapters that I care at least as much about the effect these reflections might have on the reader’s conception of her own self-knowledge. I want to undermine your self-confidence, which, however modest, on this topic is likely to be more than my own. I want to make it seem not obvious to you where the truth lies, on various issues you might have thought easily settled.

Currently ongoing conscious experience, subjective phenomenology, might seem a singularly unpromising topic for doubt. Even Descartes in his first two *Meditations* and Hume in the first book of his *Treatise*, despite all their great talent for skepticism, couldn’t bring themselves to doubt such matters. Both thought, or appear to have thought, that however great our errors may be about the outside world, we can’t in the same way err about the current contents of our own consciousness: I may not know whether there’s a red tomato in front of me – maybe I’m being systematically deceived in my sensory

inputs by some powerful being – but I know for sure what the character of my *visual experience* of that tomato is, that I *seem* to be seeing a red thing, in this hue or hues, in this shape, over this apparent distance. Likewise, I may be wrong about having dropped a barbell on my toe, but surely I cannot be wrong about the fact that I am now feeling severe pain. The whole Western philosophical tradition is nearly univocal on the special privilege, or at least excellent accuracy, of our knowledge of our currently ongoing stream of experience. Though Freud and Richard Nisbett and many others have embarrassed us with our errors about *some* features of our minds – our motives, traits, and the like – most philosophers have thought that nonetheless current conscious experience is one special corner of the mind where our knowledge is, well, maybe infallible, or if not so then at least remarkably secure.

The chapters each aim to support my pessimism somewhat differently, concerning different aspects of consciousness. They are not cumulative; with the exception of Chapter 7 (and the partial exception of Chapter 5), each serves as a kind of case study of our ignorance in some particular domain. As case studies, then, these chapters can be read in any order. I recommend starting with the chapter whose topic you find most intrinsically interesting. Chapter 7 is the most general statement and defense of my pessimism, drawing force from three brief cases studies embedded within it and, more powerfully I hope, from the more detailed case studies that constitute the bulk of the book. Since my goal is skeptical, the chapters sometimes take the form of descents into confusion, with no clear final thesis but rather a tossing up of the hands; you'll either share my uncertainty or think I'm being dense.

I will not have much to say about the *metaphysics* of consciousness, the question of whether we are purely material beings, and if so what aspect of materiality is key to the stream of conscious experience. I am, however, skeptical about metaphysical accounts of consciousness too. Partly this is because I think that one thing that became evident in the late 20th century (if it wasn't evident earlier) is that *all* metaphysical accounts of consciousness will have some highly counterintuitive consequences, if their consequences are confronted frankly. (If functionalism is true, some weird assemblages with the right functional properties will be conscious; if consciousness depends on the stuff we're made of, then aliens behaviorally indistinguishable from us may nonetheless be totally unconscious; etc.) *Something* apparently preposterous must be true of the consciousness. Thus, our ordinary untutored intuitions cannot be a reliable guide to what kinds of systems that are conscious, nor is there any reason to think that they would be. What's more, there seems to be no solid basis for choosing among the various alternatives: Armchair philosophical reflection leaves us only idiosyncratic hunches about equally unsupported half-intuitive theories, while empirical observation of physical structure and behavior is uninterpretable – cannot be accepted as showing the presence or absence of conscious experience – without a prior theory of consciousness, thus creating a tight vicious circle. (More on this last point near the end of Chapter 6.)

Nothing important in this book turns, I hope, on complex, abstract philosophical argument. History has not been kind to such arguments; evidently the cognitive powers of even the best philosophers are generally too frail for such arguments to help them gain much purchase on the truth. The problem does not generalize across fields: Complex mathematical and scientific arguments do often survive scrutiny. The difficulty, I suspect,

is principally in the background assumptions, more easily agreed upon by mathematicians and scientists than philosophers, and too shaky in the case of philosophy to support grand edifices.

Jakob Hohwy has kindly pointed out to me three recurring structures in my skeptical reflections. The core arguments are simple enough, I hope, and employed often enough with different examples, they can support my overall perspective even if several of my examples fail. First is the *argument from variation*: People often differ enormously in their judgments about their stream of experience, either across cultures, or between individuals within the same culture, or within the same individual over time. Sometimes, in such cases, it seems unlikely that people's actual underlying experiences vary correspondingly. Consequently, some of the relevant judgments – we don't necessarily know which – are probably pretty badly wrong. Second is the *argument from error*: Ordinary people often think that their experience has some feature that more careful introspection, perhaps combined with other evidence, suggests it does not have. Ordinary reflection, in such cases, is therefore prone to err. Third is the *argument from uncertainty*: When instructed to carefully inquire and asked probing questions, people often find they feel uncertainty about even the most basic aspects of their stream of experience. Such doubt suggests a substantial possibility of error in judgments of that general type, not only when people are reflecting carefully and asked probing questions but also when they reflect more casually. I ask you not to evaluate these arguments now. I wouldn't expect you to find them compelling, independent of detailed examples. I am only noting their general form, which may be of help in understanding the argumentative structure of the chapters.

This book has been improved by the discussion and criticism of many people over a long period, with most of the chapters based on essays conceived between 1999 and 2004 which have undergone continuous revision and updating (Schwitzgebel and Gordon 2000; Schwitzgebel 2002a&b, 2004, 2006, 2007, 2008). Timothy Gann, Linus Huang, Alan Moore, Russell Pierce, Cati Porter, and Daniel Price all provided detailed comments on the entire book as it was first being shaped into a whole. The following people helped me in thinking about one or more individual issues or chapters: Donald Ainslie, David Barlia, José Bermúdez, Richard Bett, Ned Block, Jorg Büchholz, Curt Burgess, Mason Cash, Dave Chalmers, Jonathan Cohen, Nelson Cowan, Dan Dennett, Josh Dever, John Dilworth, Bill Domhoff, Denise Durance, Dillon Emerick, Kevin Falvey, Bill Faw, Carrie Figdor, Maggie Friend, Kirk Gable, Jim Garson, Brian Glenney, Alvin Goldman, Jakob Hohwy, Changbing Huang, David Hunter, Russ Hurlburt, Jonathan Ichikawa, Manyul Im, Tony Jack, Michael Jacovides, Brian Keeley, Sean Kelly, David Kirchner, Bryan Lee, Steven Lehar, Felipe Leon, Dom Lopes, Jessica Ludescher, Pete Mandik, Justin McDaniel, Tori McGeer, Stuart McKelvie, Pat Missin, Christopher Mole, Jennifer Nagel, Eddy Nahmias, Shaun Nichols, Alva Noë, Gualtiero Piccinini, Bill Prinzmetal, Erich Reck, Richard Reis, Teed Rockwell, Larry Rosenblum, David Rosenthal, Frank Russo, Josh Rust, Colleen Ryan, Sandy Ryan, Wade Savage, Brian Scholl, John Schwenkler, Susanna Siegel, Charles Siewert, Carol Slater, Declan Smithies, Maja Spener, Nigel Thomas, Arnold Trehub, Penny Vinden, Glenn Vogel, Gary Watson, Gideon Yaffe, Jeff Yim, Chuck Young, Yifeng Zhou, and Aaron Zimmerman. I also thank the many readers of my blog who commented on these issues, mostly by first-name

or pseudonymously, and the many formal and informal “introspective observers” (to use Titchener’s phrase; see Chapter 5) who have given me reports on their experience. Let me also apologize to those of you in that last group for not believing most of your reports – not that I necessarily *disbelieved* them either – as well as to the many people whose help I have unjustly forgotten.

I thank my son Davy for discussing philosophy with me on evening walks. I thank my daughter Kate for showing me the meditative powers of play sand. My greatest debt is to my wife, Pauline, who not only commented in immense detail on the entire manuscript, but more importantly suffers through and supports my long hours and academic obsessions, though I cannot adequately justify them to her or even to myself.