

Chapter Five

Titchener's Introspective Training Manual

The conformity of replies from so many different sources... and the evident effort made to give accurate answers, have convinced me that it is a much easier matter than I had anticipated to obtain trustworthy replies to psychological questions. Many persons, especially women and intelligent children, take pleasure in introspection, and strive their very best to explain their mental processes. I think that a delight in self-dissection must be a strong ingredient in the pleasure that many are said to take in confessing themselves to priests.

– Francis Galton, *Inquiries into the Human Faculty and Its Development*

(1883/1907), p. 60

[After quoting the passage above] Such statements suggest that introspective exercises may be paralleled, as a form of polite recreation, with the word-puzzles in the magazines, and that the circular of questions is a royal road to the attainment of psychological truth.... [But] we get replies upon the introspective level of the average educated man. This level is low.... The ordinary observer, untrained in psychological method, can give an opinion as to the match of two colours upon a colour mixer, while he is wholly unable to follow the course of an after-image.

– E.B. Titchener, *Experimental Psychology: A Manual of Laboratory Practice*

(1901-1905), Vol. I, Pt. 2, p. 389.

We err, or at least are susceptible to confusion, about much of our stream of experience: our dreams (Chapter 1), our imagery (Chapter 3), our visual experience of depth (Chapter 2), our auditory experience of echolocation (Chapter 4). And more examples will come. In such matters our judgment is easily led astray. If I'm right about this, a natural question is, can we get better? In the late 19th and early 20th centuries, at the peak of the introspective psychology movement, many psychologists thought so – none more explicitly and insistently than E.B. Titchener, whose introspective training techniques are the topic of this chapter.

Experimental psychology arose as a distinct academic discipline in the latter part of the 19th century. At its center, sociologically, was Wilhelm Wundt (1832-1920), who founded the first formal psychology laboratory in 1879, who trained many of the next generation's leading psychologists, including Titchener, and who launched one of the first psychology journals (*Philosophische Studien*) in 1881.¹ Although the rise of psychology as a rigorous empirical science had been urged and anticipated for more than a century (e.g., in Hume 1740/1978), what finally enabled it to flourish appears to have been the emergence of *quantitative introspective methods*, especially regarding sense experience. These methods were first developed a half-generation before Wundt by researchers such as Gustav Fechner (1801-1887) and Hermann von Helmholtz (1821-1894). Early psychologists' emphasis on *introspection* helped cleave their new discipline apart from physiology (which was the formal background of many early and proto-psychologists, including Wundt and Helmholtz) by focusing it on the operations of the

mind; at the same time, their experimental methods distinguished their work from previous work on the mind by philosophers such as Locke, Hume, and Kant. Although early psychological research was not always introspective – for example, in studies of test performance (e.g., by Galton and Alfred Binet), memory (e.g., by Hermann Ebbinghaus), and reaction time (including by Wundt himself) – such non-introspective research was not at the programmatic center of 19th century psychology. Thus Wundt (1888, 1896/1897) describes psychology's subject matter as experiential processes and its method as “inner perception”, and William James in his summary of the methods of psychology writes that “introspective observation is what we have to rely on first and foremost and always” (1890/1981, p. 185).

The emphasis on quantification rendered scientific psychologists' claims more precise than their philosophical predecessors' had been, thus opening up a field of questions that seemed to invite the straightforward application of laboratory methods. Such questions included: By how much must two sensory stimuli differ to produce an introspectively observable difference in experience? What is the threshold or “limen” of experience in various sensory modalities (e.g., what tone and light frequencies are audible and visible, and how weak a stimulus can still be consciously heard or seen)? What is the relationship between the experienced intensity of a sensation and the physical intensity of the stimulus producing it? (The Weber-Fechner law – a topic of early dispute and still occasionally a topic of dispute today – holds that the relationship is logarithmic.) Along what dimensions, exactly, can sense experience vary? (Color experience, for example, was discovered to vary systematically in three dimensions: hue [e.g., red, green, yellow, blue], saturation [how intensely colored vs. how pastel or grayish], and lightness

or brightness [distance along the dimension from black to white], with the combination of these three dimensions creating the “color spindle” or “color cone” still accepted by psychologists today.) One can see how quantifying questions of this sort invites the systematic research, number-and-equation-filled journal articles, and progressive consensus building characteristic of scientific disciplines.

Consensus building, though, notoriously, only to a degree. By the early 20th century it was becoming clear that on many issues consensus was elusive. For example, the dimensions of variation in emotional experience was a hot and frequent topic of dispute with no resolution on the horizon (see, e.g., James 1890/1981; Wundt 1896/1897; Titchener 1908); so also was the existence or non-existence of “imageless thought” (briefly discussed in Chapter 3 and briefly returned to in Chapter 7) and the nature of attention (James 1890/1981; Pillsbury 1908; Titchener 1908). Although the most straightforward psychophysical questions about the relationship between external stimuli and sensory experience often yielded readily to the introspectivists’ methods, experiences whose connection to sensory stimuli was more complex proved trickier.

By the 1910s, behaviorism, which focused simply on the relationship between outward stimuli and behavioral response, had declared war on introspective psychology, portraying it as bogged down in irresolvable disputes between differing introspective “experts” (e.g., Watson 1913). In the 1920s and 1930s, introspective studies were increasingly marginalized.² Although strict behaviorism declined in the 1960s and 1970s, its main replacement, cognitivist functionalism (which treats functionally defined internal cognitive processes as central to psychological inquiry), generally continued to share behaviorism’s disdain for introspective reports about conscious experience. The

last decade or two has seen a resurgence of interest in consciousness (in the shadow of, and partly co-operative with, cognitivist functionalism), but few psychologists have looked seriously at early introspective methods to see what methodological insights might be recovered from that period.

Titchener's introspective training manual, *Experimental Psychology: A Manual of Laboratory Practice* (1901-1905), was perhaps the pinnacle of early introspective method, running 1600 pages, with separate parts for student and instructor. It presents and discusses a vast number of introspective exercises. It is peerless in its combination of depth, variety, concreteness, and attention to detail. Reading it, I feel – and will attempt to convey in this chapter – both the potential promise of formal introspective training and, simultaneously, the daunting obstacles for any such program.

ii.

Psychologists of Titchener's era often accepted as a condition of sound scientific method that introspective reports come from subjects – or as Titchener preferred to say, “observers” – with substantial introspective training. (Galton – see the epigraph and Chapter 3 – was a notable exception.) In published research, it was standard to rely exclusively on observers with graduate training in psychology and thus presumably at least several months, often several or many years, of intensive experience with introspective methods. Wundt was reputed not to have admitted data from observers who had performed fewer than 10,000 laboratory introspections (Boring 1953).

In his 1898 *Primer of Psychology*, Titchener compares the development of skill in introspection to the development of skill in physical measurement and chemical analysis (1898/1900, p. 25). Just as a chemist would never rely on an untrained assistant for any but the simplest measurements, so also the laboratory psychologist cannot rely on untrained introspectors for any but the crudest observations. In fact, Titchener argues soon afterward in his laboratory manual, quantitative introspection is considerably more difficult than similarly precise work in chemistry (1901-1905, II.2.cliii-clvii³). Consequently, “the average student, on entering the laboratory, is simply not competent” to participate as an introspective observer in demanding experiments (II.2.cliv; see also I.2.389). Difficulties include maintaining consistent attention, avoiding bias, knowing what to look for, and parsing the complexity of experience as it flows rapidly past (1898/1900, p. 24-25; see also 1915, p. 20-22). For example, without introspective training, Titchener asserts, it’s difficult to compare the relative brightness of two different colors (I.1.13; I.2.31); to differentiate a very low tone sensation from a sensation of atonal noise (II.1.1; II.1.3); to make the quantitative assessment that two sensations are each an equal distance, in different directions, from a third (e.g., that one tone sounds as high in pitch above a reference tone as another tone sounds below it) (II.2.201-204; II.1.xxxii-xxxiv); or even just to follow the course of one’s afterimages (see this chapter’s epigraph). Experienced introspectors, Titchener says, are also more likely than untrained introspectors to maintain a consistent standard of judgment and to accurately report lapses of attention and interfering influences.⁴ Reason enough already, I think, to hope that rigorously trained introspectors could bring something to the scientific table that naive introspectors would mostly lack.

Titchener turns on its head the standard argument against introspective training, that it introduces bias. Especially regarding our own minds, Titchener believes, everyone is subject to bias and preconceptions. People do not generally approach psychology neutral between theses, even when those theses are dry, psychophysical ones – and when people do start out relatively open-minded, after a few introspections they're apt to speculate and form hypotheses. Titchener consequently rejects the ideal of an introspective account “furnished by a naive, commonsense, non-scientific observer, who has not yet adopted the special attitude of the psychologist” and thus supposedly takes a “neutral standpoint” (1912b, p. 489). Such a neutral standpoint is unattainable. “We can hardly, with the pressure of tradition and linguistic forms upon us, consider mental phenomena in a really naive way, with a truly blank prescientific impartiality” (ibid.).⁵ In Titchener's view, avoiding bias requires not naiveté but expertise. Introspective practice and an “objective” frame of mind aid the observer in setting aside expectations to report mental phenomena accurately (I.2.xxv-xxvii; I.2.151; II.2.133-134; II.2.202). “The trained observer, psychologist or physicist or what not, can take the suggestion [i.e., the hypothesis toward which he might be biased] for what it's worth; he does not allow it to affect his observation. But the beginner is exceedingly liable to be led by interest into partiality” (1896/1906, p. 45; see also Müller 1904).

Although the intractable laboratory-vs.-laboratory disputes of his era suggest that Titchener may have overestimated the ability of trained introspectors to overcome bias, it doesn't follow that naive introspection is any better. I take the evidence of Chapters 1-4 to suggest that ordinary introspectors are often misled by preconceptions, theories, and culturally available metaphors. To see some of the other limitations of untrained

introspection – and also simultaneously (contra Titchener) to see some of the limitations of introspective training in Titchener’s style – let’s consider three exercises from his laboratory manual.

iii.

If two tones of frequency U (for the upper tone) and L (for the lower tone) are sounded together, it’s sometimes possible simultaneously to hear a third, lower (and generally quieter) tone, called a *difference tone*. The pitch of this difference tone will resemble that of a tone of frequency U minus L .⁶ For example, when two flutes simultaneously play the notes F_6 (fundamental frequency 1396.9 hertz) and C_6 (1046.5 hertz), listeners may also report hearing a note at about the pitch of F_4 (349.23 hertz) (Stickney and Englert 1975). Combining sine waves in a sound editor program can produce similar effects. The standard view, and Titchener’s, is that difference tones so generated do not exist in the environment but rather are a consequence of “non-linearities” in the human ear – that is, they result from the ear’s failure to respond proportionately to all frequencies and energies of auditory input, distorting the signal somewhat as an overdriven amplifier does (Plomp 1976; Hall 1980/2002; Rossing et al. 1982/2002).⁷ In addition to the (first) difference tone at $U-L$, a *second difference tone* (also called a *cubic difference tone*) may sometimes be heard at $2L-U$, and more rarely other tones, including a *third difference tone* at $3L-U$ and disputably a *summation tone* at $L+U$. As a class, these are known as *combination tones*.

Titchener introduces his introspectors in training to combination tones in the seventh experiment series in the first volume of his laboratory training manual (I.1.39-46). He begins by directing their attention to a particularly salient difference tone produced by two Quincke's tubes with fundamental frequencies of approximately 1584 hertz and 1980 hertz. (Quicke's tubes consist of a glass whistle connected to a resonator, producing a relatively pure tone; these two frequencies are approximately equivalent to the third G and B above middle C on the standard scale.) Titchener remarks that the difference tone's "moderate loudness" combined with its depth (two octaves below the lower primary tone or "generator") should make it "easily recognizable" to the student (I.1.41). He advises the student to produce the tone repeatedly until he "is entirely satisfied with his introspections". Titchener next recommends the student listen for the difference tone of two Quincke's tubes of 1584 and 2376 hertz, which he describes as particularly loud and one octave below the lower generator. After these two hopefully easy introspections are each rehearsed several times, the student is instructed to proceed up and down an octave's worth of musical intervals, then to practice hearing difference tones when one or both of the generating tones is quiet and when the duration of the tones is short. Titchener suggests that several tones be produced and the students be required to say which tone is closest in pitch to the difference tone they purport to hear (I.2.70). Finally, the student is instructed in similar procedures for the second and third difference tones and the summation tone. Titchener expects students to have only limited success in hearing the more difficult of these tones. Still, by the end of the experiment series – presumably conducted within one or a few sessions over the course of a week or less – the student should be able to discern combination tones that previously would have

eluded her. She has, apparently, become something of an “introspective expert” in this limited domain.

If you’d like to try a bit of the training yourself, I’ve produced a web adaptation of Titchener’s difference tone training procedure. It’s available from the “Difference Tone Training” link on my home page (<http://faculty.ucr.edu/~eschwitz>) and also from the free online journal *Psyche* (***) insert URL when updated (***) ; vol. 11 [2005], no. 6). For understanding the issues, nothing beats actually hearing the stimuli.

One thing to note is this: Discerning combination tones is *difficult*. Titchener begins, therefore, with comparatively easy cases, proceeding to the more difficult ones only after the easier are mastered. Also, since there’s good theoretical reason to expect each difference tone to be heard at a particular pitch – reasons having to do with acoustics and the ear and confirmed by accomplished introspectors – the students’ introspective reports can be verified, making it a good occasion for training. Many, but not all, of Titchener’s exercises share these features of *scaled difficulty* and *corrective feedback*. Indeed, so do many ordinary non-introspective training procedures.

Let’s back up a bit, though, and ask: Are students in this experiment really *introspecting*? When I try it, it seems to me that attempting to discern a combination tone differs not at all from attempting to discern a faint tone of the ordinary sort. It seems just like listening for sounds in the external environment. As far as I can tell, it’s not in any way distinctively introspective.⁸

One might hope to defend the view that the training is nonetheless introspective on the grounds that combination tones, being (in general opinion) an artifact of distortions in the ear, do not exist in the world in the same way that ordinary tones do,

and thus that in attending to them one cannot be attending to the outside world. Since it sounds odd to say that one is attending to one's ear, it is easy to suppose that one must be attending to some part of one's experience, or introspecting. However, this argument would prove too much. If every sensory or perceptual judgment about something that doesn't exist outside the observer is an introspective judgment, then many of the perceptual mistakes we make due to illusion must likewise involve introspective judgments. Perhaps, indeed, we should regard combination tones as similar to double images, color adaptation effects, or the floating black spots experienced by people with a certain sort of eye damage – that is, as a kind of illusion, a product of our sensory apparatus not straightforwardly reflecting how things stand in the world beyond. If you hold your finger six inches before your nose and focus on something in the distance while continuing to attend to the finger and you consequently notice a doubling, are you necessarily introspecting – or could you just be looking, in a particular way, at your finger? I could see the argument going either way. (By the way, did the double image exist before you attended to it? See Chapters 2 and 6.) But surely you needn't be introspecting if, with yellow-adapted eyes (perhaps unbeknownst to you), you mistakenly judge a white wall to be baby blue. The blueness is, in some sense, only in your own mind – but you do not introspect it. You're making an ordinary (mistaken) perceptual judgment about the color of the wall, not an introspective judgment about anything going on in your mind. Similarly, I'd suggest, with difference tones.

To see how Titchener's procedure qualifies as introspective training we must take a different tack. Consider the naive introspector asked to describe his auditory experience of an interval sounded by a musical instrument. If he has a minimum of

musical knowledge, he might be able to describe the interval as, for example, a major third, considerably above the middle of the scale, and indicate the instrument played if it's a familiar one. But his experience is vastly richer than those words suggest – or at least it's plausible to suppose that it is – influenced by harmonics, resonances, echoes, deficiencies in his ear, and sundry other acoustic and aural phenomena, including (presumably) combination tones. Some of these facts are indicated indirectly by his statement that it was a major third played upon (say) a piano; others are not. Auditory experience is far too complex for ordinary people to parse. Thus a new student entering Titchener's laboratory, asked to describe her auditory experience with care and in detail, would be baffled. To provide introspective reports of any value, she needs concepts and a vocabulary, a sense of what to look for, and practice in discerning these aspects of her experience as it occurs. Training in discerning combination tones is thus introspective training not because reporting such tones is necessarily an introspective act (involving a judgment about one's own mind) but because for the person antecedently interested in reaching introspective judgments about her auditory experience the training provides a way of identifying and labeling one aspect of it. "Introspective training" need not always involve acts of introspection – just as athletic training need not always involve acts of athletic prowess – as long as the training improves the quality of introspective acts when they do occur.

Trained musicians and psychophysicists, therefore, although they don't generally conceive of themselves as "trained introspectors" in the Titchenerian sense, and although they have ordinarily not undergone any general course of training and reflection on the methodology of introspection, possess some tools for apprehending their conscious

experience that others lack and that it's part of the Titchenerian introspective training procedure to provide. Indeed, to the extent that their aim *is* to apprehend their own experience – as opposed to, say, just improving (or mapping) their capacity to discriminate and label aspects of the publicly shared audible world – their project does resemble Titchener's and might profit from broader reflection on introspective methods and epistemology.

When an untrained observer at first can't discern a combination tone and then later, after training, in an acoustically identical situation, can do so, what exactly has happened? At one extreme, we might suppose that while on the second occasion she genuinely experiences the difference tone, on the first occasion the difference tone was in all respects so thoroughly absent from her experience that we couldn't even say that it contributed in some ineffable way to its richness. At the other extreme, we might hold that the auditory experience remains in all respects completely identical from one occasion to the other, the only difference consisting in a separable introspective process and judgment. Neither extreme seems to me especially inviting. Most philosophers and psychologists now take for granted that general knowledge can influence sensory experience, so that two people with the same sensory stimulation may nonetheless have different sensory experiences (e.g., Bruner and Postman 1949; Polanyi 1966; Hanson 1969; Lamme and Roelfsema 2000). If so, it seems likely that knowledge of combination tones and practice in discerning them will affect one's auditory experience, at least when one is deliberately listening for them. On the other hand, if we grant that auditory experience is rich, beyond the capacity of most observers fully to parse and articulate, and if we grant that combination tone sensations are not wholly *created* by the training

procedure but can in some sense be discovered in experience, then despite the “top down” effect of general knowledge on sensory experience, a gap of ignorance still divides the naive introspector’s auditory experience from his judgments about it; and if Titchener is right, introspective training can help reduce this gap.

Here’s an issue, though, that Titchener did not, I think, sufficiently consider: When should we regard an introspector as sufficiently attentive and well trained that we may take at face value her claim not to hear a combination tone?⁹ Besides the combination tones so far described, combination tones of 2U-L, 3L-2U, 4L-3U, 2U-2L, and others are sometimes reported for various stimulus intensities and frequency ranges, as well as combination tones arising from the interaction of the harmonics of the fundamental tones. Let’s say you deny hearing 3L-2U in some context in which others do say they hear it, and let’s suppose also that you’re a pretty well trained expert by now and firm in your opinion. Is there any way to determine whether you’re missing a tone that is a (perhaps subtle) part of your experience? To say you couldn’t possibly be missing it, even if it’s ever so subtle, seems to attribute to you an implausible infallibility. On the other hand, it seems wrong simply to assume that you must be missing it. Could we simply measure activity in auditory regions of the brain to settle the question (assuming we can sufficiently improve the resolution of neuroimaging equipment)? To do that, we’d have to know in considerable detail the relationship between neural activity and auditory experience – at that might require knowing, first, the answer to the very issue I’ve just raised. It’s not hopeless perhaps, but I see no *simple* resolution.

The issue of course isn’t limited to combination tones. Difficulties of this sort will emerge anywhere we admit the possibility of erroneously reporting the absence of

experience (see also Chapter 6) – potentially creating a major stumbling block both for evaluating individual reports, whether in the course of training or as scientific evidence, and in evaluating the final success of training methods. Perhaps tellingly, Titchener himself seems to slide, either deliberately or in confusion, between speaking of unreported difference tones as absent and speaking of them as merely introspectively undetected – most often choosing to say, ambiguously, that the observer does not “hear” them (I.1.39-46; I.2.66-72, *passim*).

iv.

The fourth series of laboratory training experiments in Titchener’s manual concerns afterimages. The fifteenth and last of these afterimage experiments begins with an observer sitting for five minutes in a dark room before a curtained window. When his partner gives a signal, the observer looks toward the window, the curtain is removed to reveal the upper two panes, and he stares fixedly for 20 seconds at the vertical bar separating the window panes. He then closes his eyes and with the room again darkened he reports the next few minutes’ visual experience. The experiment is to be repeated, Titchener says, until the observer reports similar visual experiences on every trial (I.1.29-30).

What the observer sees is a sequence of shifting afterimages, known as the “flight of colors”. You can easily induce a flight of colors in yourself by staring for a few moments at an incandescent light bulb or by glancing at the sun then closing your eyes; or you might attempt a closer replication of Titchener’s experiment.

I'll now quote at length from Titchener's discussion of this experiment in the instructor's part of the first volume. Notice the specificity of the report Titchener says his observers converge on, after their initial "mere chaos":

This experiment shows, in a striking way, the effects of practice. The report of a wholly unpracticed observer is a mere chaos. With attention, the uniformity of the phenomena soon becomes apparent; and presently the observers who at first gave radically different accounts of the after-image will reach agreement upon all essential points.

With an unclouded sky, or a sky thinly covered with clouds and presenting an even white surface, the flight of colours is as follows:

(a) A momentary positive and same-coloured image.

(b) Interval of 5 or 6 sec.

(c) Positive image, fluctuating in colour; sometimes with patches of red and green. After 1 or 2 sec., the image settles down to a sky *blue*, the vertical bar remaining dark.


(d) The blue passes, with or without interruption, into a *green*. The green is at first very vivid; it disappears and reappears five or six times, growing gradually paler; at last it is almost whitish. – These initial changes show a good deal of individual variation. Some *O*'s [observers] now see

(e) A *yellow* image. This (or the whitish green preceding) is regularly followed by

(f) A deep *red* image. The black bar becomes luminous and slightly greenish, the light appearing at first as a crack in its length. This is the stage of transition from the positive to the negative image. The red undergoes several fluctuations. Then follows

(g) A deep *blue* image, with yellowish bright bar, more lasting than any of the preceding phases. The blue darkens, and the image gradually disappears, with or without passing into

(h) A dark *green* image.... Note the periodicity* of stages *c* to *h*:



It's not clear whether Titchener is right that practiced observers eventually settle on similar descriptions of the flight of colors. Titchener cites Helmholtz (1856/1909/1962) and Washburn (1899), who report roughly similar sequences of colors. However, Helmholtz's description is a rather bare statement that intense white light produces afterimages that proceed white-blue-green-red-blue. And Washburn was not really an independent source, having recently earned her Ph.D. under Titchener. (Washburn was, incidentally, both Titchener's first doctoral student and the first woman ever to receive a Ph.D. in psychology.) One might wonder whether Titchener's explicit instruction that observers are to settle on a single sequence influenced his findings. It's also unclear what influence, including theories discussed in the laboratory, might incline observers to report one sequence rather than another.

* Blue and yellow are generally treated as opposing colors in visual perception, as are red and green.

By far the most detailed treatment of the flight of colors, complete with color plates depicting the afterimages, is a 1913 article by Paul Homuth. Homuth is even more insistent about training than is Titchener, claiming that expertise requires several months of intensive practice in observing afterimages. Homuth divides his images (which do not include a vertical bar) into four parts, the center, border, outer frame, and extreme periphery, which undergo different color shifts. In his primary condition with bright white light, Homuth reports the center of the afterimage to be mainly blue alternating with reddish-violet, magenta, or pinkish-violet, with the sequence concluding in brownish-yellow. The resemblance to Titchener's description is minimal at best.

In 1922, William Berry reviews the literature on the flight of colors all the way back to Aristotle and finds great variability of this sort among researchers. He concludes that there's no consistent sequence in the flight of colors, a point he later (in 1927) supports with a study with his own graduate student observers. On the other hand, V.M. Robertson and G.A. Fry (1937) point out that earlier observations were conducted in a wide variety of conditions and might be expected to produce variable results even if there is consistency in the flight of colors in any one condition. They report consistency among their own observers, with results fairly similar to Titchener's (as do Weve 1925 and Barry and Bousfield 1934). The very sparse more recent research that I've been able to find on the flight of colors does nothing to resolve the issue. The matter was less settled than dropped.¹⁰

But – you might wonder, I at least wonder – how difficult could it be, really, to sketch, in broad outline, the flight of colors? What exactly is the challenge here? It may be difficult to articulate precisely what's going on in Homuth's four regions, or to

describe accurately an exact shade, or to characterize (or maybe even just inarticulately appreciate) the complexities in an afterimage with shifting, ragged borders. But to label the general approximate color range of an afterimage if that same general color persists over several seconds – that seems pretty easy, doesn't it? When I've tried to replicate Titchener's windowpane exercise, the color of the resulting afterimages seems to me usually to be very obvious. This morning, for example – May 6, 2009 – I tried the exercise again. It seemed to me that my afterimage evolved, over the course of about a minute, from pretty much white (with a dark bar in the middle), to a kind of lavender and then brilliant red, finally fading into blackish blue (with the bar lightening up near the end). Could I be wrong about this? Might the image actually have been, for some of this period, vivid green (as required in Titchener's sequence), while I thought it was lavender or brilliant red? I find this rather difficult to imagine. I don't feel the same kind of uncertainty here that I felt about my imagery experience at the beginning of Chapter 3 (when I considered how stable the image was, how flat, etc.) or about my echoic experience at the end of Chapter 4. Nor does it seem likely to me that others would err in a similar situation.

This feeling of confidence, this feeling of the ease of the introspective task, is of course just an intuitive assessment. I can imagine Titchener rejecting that assessment. There are features of the retina we all share, he says, that explain the evolution of afterimages (I.2.49), so the flight of colors should be the same in all normal perceivers (assuming the retina to be properly whitewashed with light). Consequently, my flight *must* have had a green phase. Or if not my flight – since I may somehow have done the experiment wrong – then at least the flights of Titchener's own students, despite the

novices' divergent reports. Furthermore, Titchener might say, if it's hard even for novices to go grossly wrong about the flight of colors, and if his students do in fact move from diverse to convergent reports, then it would seem to follow that his training procedure somehow *regularizes* the flight of colors; but that strains against the natural idea that the flight of colors is driven by relatively early (e.g., retinal) features of the visual system. So if I or some other insufficiently trained observer detect no period of "vivid [green]... growing gradually paler" in some properly induced flight, then maybe the best conclusion to draw, the conclusion that Titchener presumably would draw, is that we're just wrong, any intuitive sense of certainty notwithstanding. Why should we think, anyway, that our intuitive sense of certainty or uncertainty about such matters is particularly trustworthy?

This argument, I confess, leaves me cold. I just can't bring myself to doubt that experience of lavender followed by brilliant red. There was no green. But now I wonder: Am I guilty here of the same unjustified blasé confidence that I so readily attribute to others? Or does my judgment instead reflect instead some reasonable attunement to what is and is not relatively dubious?

v.

The compelling visual illusions that one generally sees in books and demonstrations tend to mask the introspective difficulties that arise for weak or non-obvious illusions. Confronted with Poggendorff's illusion (fig. 1, from I.1.165 in Titchener's laboratory

manual), most people feel unambivalently comfortable in reporting that, in some sense, the partly occluded line which we may know to be straight nonetheless “looks” crooked.




fig. 1

If we accustom ourselves only to such easy cases – the “best” illusions – it might seem inconceivable that one would have to look hard to find an illusion, that one might be talented or inept at the introspection of illusions, and that the criteria of illusoriness and visual appearance are evasive.

Carefully examine the figure below, from page I.1.154 of Titchener’s manual:




fig. 2

Titchener invites his students to consider the following questions, which I ask you also to consider:

How does the figure *A* strike you at first sight? Fixate on some point on [the line segment] *be*. What is the appearance of the figure? Move the eye slowly from *b* to *e*, and back again. Does the figure change its perspective? Move the eye from *b* to *c*, and back again. Is there any change? Is there any uniformity of perspective, according as you move in the directions *bc*, *ba*, *ef*, *ed*, or in the opposite directions?

How does the figure *B* strike you at first sight? Fixate, first, a point upon *bd*, and then a point upon *ac*, *ad* or *cd*. Is there any difference of perspective? Move the eye slowly in the direction *ba* or *bc*; and then in the direction *ab* or *cb*. What happens in the two cases? What secondary modifications of the appearance of the figure are conditioned upon the shift of perspective? How many perspective illusions, in all, is the figure capable of producing? (I.1.154)

If you're like me, following these directions is rather difficult – perhaps surprisingly so. The difficulty lies partly in controlling one's attention and the movement of one's eyes, resisting the temptation, for example, to glance at point *c* as one is supposed to be moving one's fixation slowly along *ab*. Since controlling attention and eye movement is crucial to many introspective tasks, such practice is part of Titchener's introspective training (see note 4 for references).

In the instructor's part, Titchener says that in both figures, the central line is usually seen as closer to the observer but that fixation on any point on a line tends to

bring that line forward, with consequent modifications of perceived transparency and orientation (I.2.310-311).¹¹ Also, he says, moving one's eyes along a line tends to bring forward the initial vertex rather than the terminal one. Although my own introspections of the experiences produced by figure 2 were initially rather disorganized, I find them now mostly to conform to the pattern Titchener describes. But I'm unsure whether I'm now judging my experience of the figures more accurately or whether accepting Titchener's generalization has altered my experience.

Let's consider another figure in which perspectival Gestalt shifts are less salient, figure 3 below (from I.1.160; with apologies for the imperfect reproduction, especially of B).




fig. 3

Titchener asks his students to view the objects in figure 3 serially, first with both eyes, then with one eye at a time. He asks: "Is there any illusion of extent [i.e., in the apparent length of the lines]? Is there any other illusion? Look very carefully, in both cases, and do not be satisfied with your first discovery". (All the figures are perfectly regular on the printed page.)

Faced with this task, I find that I feel considerable uncertainty about how the shapes in figure 3 actually look to me. Maybe you'll feel the same way. In figure 3A, does the vertical line look longer, shorter, or the same length as the horizontal? With the right eye closed, does the right horizontal limb look longer, shorter, or the same length as the left? Of course, you can quickly toss out a response, confident that no one will prove you wrong (if such a proof is even possible); but approaching these questions conscientiously, I at least am unsure of myself.

If you share this feeling, perhaps you'll also share the sense that to find yourself in this difficulty is, in a way, peculiar. How could it be hard to reach a judgment about how things appear to you? Although judgments about how things *are* understandably carry some risk, judgments about how things *look to you right now* seem insulated in a particular way. Could you really go wrong in such a judgment? And if you couldn't go wrong, where does the difficulty lie?

Some readers will not feel any difficulty or have any sense that they could be mistaken. If that feeling arises from general temperament or philosophical conviction, this chapter, indeed this whole book, will probably seem to you misconceived. Others of you, however, may have approached the task too casually, since after all no Titchener is standing over your shoulder forcing you to write a detailed lab report. Consider then, in more detail, figure 3A. Examine it both binocularly and monocularly. On first glance I've found that most viewers report no illusion: The two lines look to be equal length and to bisect each other perfectly at right angles. Nevertheless, figures of this sort are standardly presented as examples of the "horizontal-vertical" illusion (e.g., Robinson 1972, p. 97; Coren and Girgus 1978, p. 29).¹² Experts in visual illusion appear to agree

that, in some sense, the vertical line in 3A *does* look longer to normal perceivers. Now maybe something about the arrangement of this particular figure, with other shapes and a frame nearby, compromises this illusion, but Titchener at least appears not to have thought so (I.2.309; I.2.315). For comparison, maybe 3E or 3D show the illusion more obviously. Return, then, to figure 3A. Are you still confident that the lines look the same length? With one eye closed, the horizontal-vertical illusion purportedly is reduced or vanishes (I.2.315; also Prinzmetal and Gettleman 1993). Titchener also claims that in monocular vision the outer horizontal limb (on the side away from the nose) looks longer than the inner, and that in binocular vision the upper vertical limb looks longer than the lower (I.2.315).¹³ Is it obvious to you which of these illusions is present or absent in your own experience?

Part of the difficulty here – if I’ve managed to persuade you that there is a difficulty – may be that it’s not clear what it is for two lines to “look” the same length. Clearly, it can’t be a matter of one’s *overall* judgment about the length of the lines, since one can judge that two lines look different lengths even when one knows them actually to be the same length. Is it a judgment about what your assessment of the lines would be, if you were to depend only on immediate visual cues? I doubt visual cues operate separately from general knowledge in the way presupposed by such an approach. And in any case, the necessary judgment would seem to be a difficult hypothetical one, requiring us to ascertain the bases of and influences on our assessments – which we seem to be rather poor at, generally speaking, for reasons famously reviewed in Richard Nisbett’s and Timothy Wilson’s work (Nisbett and Wilson 1977; Nisbett and Ross 1980; Wilson 2002).

Do two lines look the same length if they extend equal lengths across what we might think of as “the television screen of visual experience”? Many psychologists and philosophers now think that there is no one locus of visual experience, where everything comes together as on a screen, but rather a sequence of processes, some in parallel, that may yield differing results. (Such a perspective is engagingly explored in Dennett 1991.) There’s empirical evidence that different parts of the visual system are differently subject to illusion. One influential series of experiments studied the Ebbinghaus (or Titchener Circles) Illusion, in which a circle surrounded by larger circles is judged verbally to be smaller than a circle of the same size surrounded by smaller circles. The experimenters found that the part of the visual system that guides reaching, as measured by the distance between the fingers during grasp, is largely unaffected by the illusion.¹⁴ One question, though, is whether grasp aperture is a good measure of illusion, conceived of as a feature of visual experience. Maybe it’s a better measure of one’s assessment of how things stand in the world. (The two may of course come apart, for example in cases of known illusion.) Or maybe part of the illusion is that one’s grasp aperture is different in the two cases, so that two grasps of the same aperture may nonetheless reflect different experienced sizes. Furthermore, even if we decide there is something like a single television screen of visual experience (flat or otherwise, see Chapter 2), it’s unclear whether how things look should be judged by their appearance on it. Does an oar half in water “look” straight or bent (Ayer 1940; Austin 1962)? Does a cat behind a picket fence “look” like so many cat slices (Noë 2004)? Presumably there’s an illusion in figure 3A just in case the lines look different lengths. But now I’m unsure exactly what this means or how we’re to come to a dependable judgment about it.

To add a different sort of difficulty: Suppose that when you focus on the horizontal line your sense is that the vertical line, as peripherally experienced at that moment, is the same length as the horizontal, but you recall the vertical line to have looked longer when it was in focus. Or suppose that you don't feel sure whether the vertical line looks longer while you are focused (attentionally and/or by visual fixation) on the horizontal, but feel a kind of compulsion to focus on the vertical to make the judgment. Or suppose that to avoid these sorts of problems you resolve to fixate steadily on the center of the cross and attend equally to the whole thing, but then you find that it shifts in some way either from or to looking perfect. If visual experience is a complicated flux, there may be no stable experience of the lengths of the lines to underwrite a stable judgment about which looks longer, in close cases of this sort.

One way to approach the question of whether there are illusions in figure 3, even for people who claim to see none, would be to present a variety of similar figures in which the lines differ in length. The subject might then be required to state which figures are skewed which direction, and the researcher could check for a tendency to err one way or another. (Künnapas 1955 did this for figures like 3D and 3E.) Alternatively, the subject might be given the opportunity to adjust the lines until they seem equal length (as Gardner and Long 1960a-b did for figures like 3D and 3E). Problems also dog this approach, however. Such judgments either replace judgments about how long the lines *look* with judgments about how long the lines *are* or blur the two sorts of judgment together. Maybe this is acceptable if participants are sufficiently naive, but someone aware of the possibility of illusion might treat the two questions differently. Furthermore, the presentation of multiple figures in sequence, or the ability to control the

length of the lines, significantly alters the cognitive situation. Gardner and Long found that as small a variation as whether the horizontal line is fixed and the vertical adjustable or vice versa had a pronounced effect on the magnitude of error. It's therefore conceivable that people may consistently err on such tests and yet experience no corresponding illusions in figure 3.

So if someone reports no horizontal-vertical illusion in figure 3A, should we conclude that she genuinely does not experience such an illusion? Or might one line look longer than the other despite the observer's being an insufficiently capable introspector to discover that fact about her visual experience? I can't see how we might easily resolve the issue. To insist on the former seems unrealistically to deny the possibility of inaccuracy in assessing the complex stream of visual experience. To insist on the latter risks opening the door to a world of illusions that no one reports and that never deceive us.

Perhaps we can imagine an observer who, when presented with a variety of figures like those in figure 3, reports experiencing several small illusions in one direction or another for each of the figures, though most observers report no such illusions; and perhaps further it turns out that both this observer and those who deny the illusions, when given tests like those described two paragraphs above, err slightly but similarly in the directions predicted from the reports of the first observer. It might then be plausible to suggest that all the observers experienced illusion in the original figures – that the lines actually looked to them, in *some* relevant sense of “looked”, to be different lengths despite their contrary report – and thus that the first introspector has a talent for discovering non-obvious illusions that others misreport. Maybe such a talent could be

nurtured with proper training. This is, I think, what Titchener hoped and believed.

Whether things would actually turn out so neatly is anyone's guess. Neither Titchener's trainees nor anyone else I'm aware of has been put to systematic test.

vi.

Probably no part of early introspective methodology was more thoroughly and durably overthrown than the emphasis on introspective training. No contemporary research avails itself of observers that Wundt or Titchener would regard as well trained. Yet if accurate introspection is difficult, one might expect training to produce substantial benefits.

Plausibly, a trained introspector would at least, as Titchener says, employ more stable standards of judgment, better maintain consistent attention, know better what to look for, and deploy more sophisticated concepts for describing complex experiences. Maybe, too, she'll better follow the course of an afterimage, more dependably spot subtle visual illusions, more accurately detect combination tones. Yet the effects of classical introspective training were never adequately tested. Instead, when introspective psychology seemed just on the cusp of becoming rigorously self-conscious about evaluating its training methods – Müller (1904) and Wundt (1907) come to mind as well as Titchener – it perished amid squabbles and at the eager hands of behaviorism.

It would never have been entirely straightforward to prove the value of training, as I hope the examples of this chapter illustrate. It's challenging to evaluate the accuracy of introspective reports not only when the relationship between stimulus and experience is likely to be complex and highly variable – as with emotional experience and abstract

thought – but also even sometimes when the relationship between stimulus and experience appears to be comparatively simple and dependable, as with combination tones, afterimages, and visual illusions. To what extent does introspective training alter the target experience? When is someone expert enough that we should accept her denials of subtle experience? With how much skepticism should we view reports (of novices, of experts) that conflict with the most straightforward interpretations of non-introspective evidence? Might attempts to verify introspective reports problematically change the target experience? How much weight should simple assertions of confidence receive, and might practiced introspectors' confidence be better calibrated? Such questions are not, perhaps, unanswerable. The right conjunction of evidence might confirm the value of trained reports. Given the dubious quality of introspective judgments by untrained subjects and by philosophers and psychologists across history, as argued throughout this book, some more formal training regimens may be worth at least a try. I'd wager on mixed results.

¹ The first psychology journal, founded by Alexander Bain in 1876, was the British journal *Mind*. Wundt's journal was second. *Mind* is now one of the "big three" leading *philosophy* journals, with rarely more than a whiff of empirical psychology. Notably, *The Journal of Philosophy*, another one of the current big three philosophy journals, underwent a similar transition. It was founded in 1904 as *The Journal of Philosophy, Psychology, and Scientific Methods* by philosopher Frederick Woodbridge and psychologist James McKeen Cattell, featuring work by Titchener, George Trumbull Ladd, and other leading psychologists in its inaugural issues. The final four words disappeared from its title with volume 18 in 1921.

² The consequent amnesia for early introspective procedures was compounded by the simultaneous rise, Gestalt psychology (e.g., Köhler 1929/1947) as the chief competitor to behaviorism in academic research psychology. (Freud, his competitors, and his followers played a large role, of course, in *clinical* psychology – a field that during the 20th century overlapped surprisingly little, scientifically or sociologically, with laboratory research psychology.) Gestalt psychology, though it gave an important role to introspection, regarded classical introspective techniques as objectionably reductionist in their analysis of experience into discrete sensory elements. Gestalt theorists held the holistic sensory "Gestalt" (literally, figure or form or shape) to be primary in ordinary experience.

³ References to Titchener's laboratory manual will list the volume in capital Roman numerals, followed by the part in Arabic and the page in either Arabic if it is from the body or lower-case Roman if it from an introduction. The second part of each

volume was intended only for the use of the laboratory instructor. The 1971 reprint of *Experimental Psychology* omits the second part of the first volume; and each part of the first volume is itself misleadingly divided into two “parts” (perhaps explaining the reprinter’s confusion).

⁴ For more specific discussions and examples of the benefit of practice in focusing on stimuli, steadying one’s sense organs, and controlling one’s attention, see I.2.30-31; I.2.121; II.2.cliv-clvi; II.2.307. For the benefit of practice in attaining a consistent standard of judgment, see I.2.87; II.1.xxxiii; II.1.1-2; II.1.25-26; II.2.307. On knowing what to abstract, attend to, or look for in a complex sensation, see I.1.41-42; I.2.48; I.2.52; I.2.75; I.2.87; I.2.217; I.2.300. On reporting lapses of attention and interfering influences see I.2.167, I.2.220-222; I.2.341-345; II.2.402; and in conjunction with each other II.1.104-106; II.2.2-3; II.2.260. See also Nahmias 2002.

⁵ However, E.G. Boring claims that later in his career Titchener put “considerable faith” in the method of naive phenomenological report that he here criticizes, though he never published on the subject (1929/1950, p. 416; see also Boring 1927, p. 502; Evans 1972).

⁶ For readers unfamiliar with these terms: *Frequency* is a physical measurement of rate of vibration, in this case of a sound wave, in hertz or cycles per second. *Pitch* is a subjective phenomenon pertaining to how high or low a tone sounds on the musical scale. Generally speaking, higher frequency tones sound higher in pitch (doubling the frequency increases the pitch by one octave), though as with most psychophysical phenomena the relationship between stimulus and experience is complex when examined in detail.

⁷ I'm not entirely convinced there isn't a sense in which difference tones exist in the environment (see Hall 1981), but the philosophical and acoustic issues are complex, turning in part on what counts as a proper component of a soundwave structure.

Ultimately, I think no major points in this chapter hang on the issue, as should become evident later in the section.

⁸ A possible objection is this: I know about the outside world – about things such as tones in my environment – by knowing, in the first instance, about the *stream of conscious experience* that that world produces in me and then inferring from there to what the world itself must be like (views of this sort have been associated with Descartes [1641/1984], Locke [1690/1975], and Russell [1912]), and it is by introspection that I know the stream of experience. Therefore, *all* perceptual knowledge is grounded in introspection, including the detection of combination tones. My view on this issue is rather complicated: I think we do have a sort of attunement to our experience that serves as part of the epistemic ground for our judgments about the outside world, but I don't think that such attunement rises to the level of classificational judgment, and I don't think it is properly called "introspective" (see Schwitzgebel forthcoming-a). As I'll suggest in Chapter 7, our introspective judgments are so unreliable that they could not plausibly serve as a basis for our knowledge of objects and events in our environment.

"Transparency" theorists in philosophy argue that introspective attention to experience is impossible – attention always "passes through" experience to the outside world (Harman 1990; Shoemaker 1994/1996; Dretske 1995; and Tye 1995, 2000, 2003; for critical discussions, see Kind 2003 and Siewert 2004). Harman writes:

When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experiences. And that is true of you too. There is nothing special about Eloise's visual experience. When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree (1990, p. 667).

If this is correct, it would certainly explain the apparent similarity between the introspective and the perceptual tasks in discerning combination tones.

However, on a plausible psychological notion of attention as involving the dedication of central cognitive resources or working memory, it seems clear, contra Harman (and Tye etc.), that we *can* attend to the visual experience tree. At a minimum, we can *reflect on* or *think about* what our visual experience is, in an attention-hogging way. Harman might say that he meant only we don't *visually* attend to the experience of the tree; but if that is the correct interpretation of his claim, it doesn't seem to do the work he and Tye want it to do in support of representationalism. Also it's hard to see how anyone would be tempted in the first place to think we visually attend to our visual experience: We don't *see* our visual experience, after all; no light reflects from it into our eyes. Siewert advances a more moderate claim that seems more sustainable: We don't *withdraw* attention from the tree in attending to our experience of the tree.

My own take is that introspective attention and perceptual attention differ mainly in the *goal* that's implicit or explicit in the attentive process. In introspection, the goal is to discern features of your own conscious experience or phenomenology; in perception it's to discern features of the outside world. These goals needn't compete; often they nest. Discerning features of your perceptual phenomenology generally involves, as an integral part, reaching perceptual conclusions about the outside world. Introspection differs from ordinary perception principally in what further things you do with those perceptual conclusions, in what other processes are integrated with the perceptual process, and in what the standards of success are (for example, reporting a tone generated by tinnitus [ringing in the ears] would be a mistake if your purposes were perceptual but not if your purposes were introspective). We attend to the features of our experience in part *by* attending to the features of the outside world. The two kinds of knowledge intermix. Consider looking at a tree through a slightly warped window: You might want simply to learn about the tree, or you might want to learn about the window. If the latter, you do so in part by noting how the tree looks through it. Analogously for introspection: We learn about properties of our sensory experience in part by noting the apparent properties of the perceived world. I develop this view of introspection further in Schwitzgebel (forthcoming-a). The entanglement of perceptual knowledge with introspective knowledge of perceptual experience, combined with the fact, mentioned in the main body of the chapter, that introspective training need not always involve introspective activity, explains, I think, the subtle but real difference between musical training, conceived non-introspectively, and Titchener's training here.

⁹ One might try to avoid the puzzle by adopting the position that all mathematically simple combination tones contribute to any auditory experience of a musical interval. However, this position faces two obstacles: First it's unclear where to draw the lines. Could we experience 7L-4U? 10U-14L? 16L-11U of the third harmonic? Second, even the best trained experts (e.g. Plomp) don't report *that* many combination tones in every experience of an interval, so it seems to imply a wealth of combination tones that no one can discern. This second concern, of course, interacts with the first.

¹⁰ Wallace (1979) takes individual variability in the flight of colors for granted, while Young (1948) and Feldman et al. (1974) assume the contrary. In his influential general review of the literature on afterimages, Brown (1965) seems at one point to agree roughly with Titchener's description of the flight of colors (p. 480) but at another point, apparently inconsistently, to endorse Berry's claim that the flight of colors varies greatly between people (p. 490).

A related issue is whether people experience a similar or variable evolution of colored afterimages after exposure to colored light. The evidence on this question is also divided. See Homuth 1913; Weve 1925; Judd 1927; Brown 1965; Stamper et al. 2000; Taya and Ohinata 2002. Other papers of interest include Fröhlich 1921 and Shuey 1924.

¹¹ Many researchers have discussed ambiguous figures that appear to reverse perspective, the best known example being the Necker cube (that famous drawing that looks like a transparent cube with black edges, consisting of two partly overlapping squares with lines connecting the corresponding corners). The idea that attention to a particular vertex tends to bring it forward goes back to Necker himself (1832). More

recent research suggests this tendency is not perfect and a number of other factors may be involved (Köhler and Wallach 1944; Hochberg 1950; Pritchard 1958; Gregory 1970; Girgus et al. 1977; Peterson and Hochberg 1983; Long and Olszweski 1999).

¹² Earlier sources for the horizontal-vertical illusion include Oppel (1854-1855) and Künnapas (1955). Titchener claims that every object in figure 3 shows this illusion, except the last, which he says shows no illusions. (Do you think you'd have written this in your lab book if you'd been his student?)

¹³ Titchener attributes to Kundt (1863) the view that in monocular vision the outer limb appears longer than the inner, an illusion also supposedly present in 3E. I'm not sure this is unambiguously implied by Kundt; but Kundt does clearly claim that in monocularly bisecting a horizontal line, subjects will show bias toward one side. I informally tested this claim by having acquaintances monocularly bisect, with a pen stroke, horizontal lines of varying length, but I found no consistent trends. Titchener attributes the view that the upper limb looks longer than the lower to Delboeuf (1865; see also Nicolas 1995) and claims the illusion is also present in figure 3B. Contemporary research appears to run contrary to Delbouef, with subjects bisecting vertical lines too high (Post et al. 2006).

¹⁴ Agliotti et al. 1995; Haffenden and Goodale 1998; Clark 2001; Glover 2002; Smeets and Brenner 2006; Goodale et al. 2008.