Chapter Three

Galton's Other Folly

There was a time when I could visualize the obverse, and then the reverse. Now I see them simultaneously. This is not as though the Zahir were crystal, because it is not a matter of one face being superimposed upon another; rather, it is as though my eyesight were spherical, with the Zahir in the center.

– Jorge Luis Borges, "The Zahir" (1949/1962)

i.

Francis Galton was one of the great intellectual pioneers of the Victorian era. He explored Africa; he introduced the pressure-system weather map; he invented the statistical concepts of correlation and regression; he published influentially on heredity, fingerprints, and genius. His interest in quantification was sometimes whimsical: He invented a measure for boredom at meetings of the Royal Geographic Society (by counting fidgets), and he worked on a "beauty map" of Britain based on counting the number of "attractive, indifferent, or repellent" girls seen in the streets (he ranked London highest, Aberdeen lowest: 1908, p. 278, 316). Like many pioneers, he also sometimes misjudged the new territory, most notoriously and problematically in his advocacy of eugenics (a term he coined), an idea with (as I'm sure you know) a horrible history in the 20th century. The prefigurings of this history can be seen in Galton's cavalier assumptions about how much human variation, including between the races, is

due to heredity, and in his readiness to forbid the reproduction of "degenerate" and "unfit" people like criminals and paupers (he himself was independently wealthy). "Folly" is perhaps too gentle a word.

Galton's other folly was, in my view, one of the contributions for which he is most praised: his cataloguing of human differences in visual imagery, based on self-report and anecdote (1880, 1883/1907). I will argue in this chapter that the differences in people's reports about their imagery fail to reflect real differences in their underlying imagery experience. For 130 years, the psychological study of imagery has been burdened with excessive optimism about subjective report, inherited partly from Galton.

ii.

Close your eyes and form a visual image. (Are your eyes closed? No, I can tell, you're peeking!) Imagine – as Galton (1880) suggests in his first classic study of imagery, which we'll discuss in more detail shortly – your breakfast table as you sat down to it this morning. Or imagine the front of your house as viewed from the street. Assuming that you can in fact form such imagery (some people say they can't[1]), consider this: How well do you know, right now, that imagery experience? You know, I assume, *that* you have an image, and you know some aspects of its content – that it's your house, say, from a particular point of view. But that's not really to say very much yet about your imagery experience. Consider these further questions:

How much of the scene can you vividly visualize at once? Can you keep the image of the chimney vividly in mind at the same time that you vividly imagine your

front door? Or does the image of the chimney fade as you start to think about the door? How much detail does your image have? How stable is it? Supposing you can't visually imagine the entire front of your house in rich detail all at once, what happens to the aspects of the image that are relatively less detailed? If the chimney is still experienced as part of your imagery when your image-making energies are focused on the front door, how exactly is it experienced? Does it have determinate shape, determinate color? In general, do the objects in your image have color before you think to assign color to them, or do some of the colors remain indeterminate, at least for a while (as, in Chapter 1, I suggested may be the case for many dream objects)? If there is indeterminacy of color, how is that indeterminacy experienced? As gray? Does your visual image have depth in the same way your sensory visual experience does (*if* it does – see Chapter 2), or is your imagery somehow flatter, more sketch-like or picture-like? How much is your visual imagery like the experience of seeing a picture, or having phosphenes (the little spots of color many people report when they press on their eyes[2]), or afterimages, or dreams, or daydreams? Do you experience the image as located somewhere in egocentric space – inside your head, say, or before your eyes, or in front of your forehead? – or does it make no sense to attempt to assign it a position in this way?

Most of the people I've interviewed about their imagery, when faced with such a series of questions, at some point stumble or feel uncertainty. The questions seem *hard* – or at least some of them do (different ones, I've found, to different people). They seem like questions one might get wrong, even when reflecting calmly and patiently. And if the questions are hard, so that people can in fact easily err – well, on the face of it that would seem to conflict with many philosophers' optimistic views about the infallibility or

practical accuracy of our knowledge of our ongoing stream of experience. It would seem to conflict, for example, with Descartes' (1637/1985, 1641/1984) claim, and Price's (1932, as quoted in Chapter 2) and Locke's (1690/1975) and Sydney Shoemaker's (1963) and David Chalmers's (2003), that it's impossible to doubt or be mistaken about your own experiences, at least to the extent your judgments pertain entirely to your ongoing phenomenology (that is, your subjective experience or consciousness). It calls into doubt, I think, the wisdom of psychologists' (including Galton's) widespread trust in the general accuracy of people's reports about features of their stream of experience. Now I don't want to reject such optimistic assessments of our self-knowledge too hastily. This whole book is aimed at undermining such optimism, and we'll return repeatedly to optimists' possible responses, qualifications, and countermoves. But for now I just want to note the surface plausibility of the following inference: *If* people can easily err about their own ongoing conscious experience of imagery, then there is at least one major type of conscious experience about which people can easily go wrong; and that in turn would seem to spell trouble for broad claims about the impossibility, or rarity, or necessarily pathological origins, of mistakes about one's own currently ongoing stream of experience.

Of course, not every reader faced with these questions will feel the same uncertainty that I do and that most of my interviewees seem to. If you are one of those confident readers, then I suspect you'll lack sympathy with my critique of Galton's and others' optimism. The sense of doubt I hope to evoke in this section is what gives force to the remainder. It's what lends plausibility to the thought that there may be widespread error in the cases I'll soon be describing. If you don't feel such doubt yourself, then

you'll probably find more appealing some alternative explanation of the phenomena to be discussed. So if while reading the opening two paragraphs of this section you didn't actually attempt to form an image and answer the questions, I urge you to do so now and frankly assess, as well as a person can in such circumstances, the difficulty of the questions and potential for error.

Let's suppose you do feel some uncertainty or room for error. You may still perfectly well know your experience. Possible error doesn't entail actual error or even, always, a failure of knowledge. Or your feeling of uncertainty may be a poor index of the likelihood of error. Furthermore, it could be that any uncertainty you feel flows not from some shortcoming in your epistemic relationship to your imagery experience but instead only from the tangle of concepts I've invoked in my questions. Your confusion might, that is, be like the confusion I would feel if someone asked me whether a particular shade of red I'm looking at is scarlet, maroon, vermillion, or magenta. I know perfectly well what shade I'm looking at, but my command of the sub-vocabulary of redness is too weak for me to confidently apply such labels. Or maybe it's like the confusion I might feel if my accountant (as a practical joke) were to ask me if I've stopped cheating on my taxes yet. I've never cheated on my taxes that I recall, and ordinarily I'd say that I know this fact about myself; but I'm not *certain* that I've never cheated (maybe I cheated in a small way that I've since forgotten), so the question leaves me flustered. It builds in a false presupposition that I can't decisively reject. Such sources of doubt are artifactual, incidental to how the questions are posed, and don't reflect the sort of potentiality for error that I have in mind.

Granted, then: Some of the vocabulary I've used may strike you as strange ("image-making energies"?), and some of the questions may have problematic presuppositions (such as that an image is the kind of thing to which it makes sense to attribute stability or instability). That's probably unavoidable. (Why is it unavoidable? I suspect it's because our limited understanding of our imagery hobbles our thought and talk about it; but to put any weight on that suspicion here would be to assume what I hope to conclude.) So feel free to set aside what seem to you the problematic questions. Or better, recast them in your own terms. Return to them at leisure, when they're not cheek to jowl with a dozen others. Don't worry about keeping to a shared, stable vocabulary, but try to consider just your knowledge of the experience itself undressed, in its broad characteristics. For me at least, even trying all this, the feeling of difficulty and uncertainty remains.

There may be a level of detail beyond which it's inappropriate to ask questions. One can imagine a striped tiger without imagining it as having a determinate number of stripes (Price 1941; Dennett 1969; Block, ed., 1981). Or at least maybe one can: George Berkeley would disagree (see the next section below). It would then be perverse to insist on a precise answer about the number of stripes and take any resulting uncertainty as a sign of introspective ineptitude. But I hope my questions aren't like that. They concern, for example, *whether* the imagined tiger (or chimney) has a determinate number of stripes (or bricks) or not – which is a substantial, or at least middle-sized, feature of the imagery experience, namely how detailed or sketchy it is (in some sense of "sketchy"). (Of course, it may be indeterminate whether an image is indeterminate; but advocating

higher-order indeterminacy is quite different from accepting the more ordinary view that images often *do* have some indeterminacy.[3])

iii.

The history of psychology and philosophy has seen three major debates about imagery, which I want to briefly address before returning to Galton. The most recent – what in the 1980s was often called "the" imagery debate – concerns underlying cognitive structure, whether imagery is underwritten by language-like symbolic structures (as per Zenon Pylyshyn 1973, 2002) or whether, instead, it requires some more irreducibly pictorial (or "quasi-pictorial") cognitive structure (as per Stephen Kosslyn 1980, 1994). It's unclear whether this debate turns on any disagreement about the actual conscious experience of imagery. Maybe theoreticians in both camps could entirely agree about imagery experience while continuing to disagree about cognitive architecture.[4]

However, the two other major historical debates about imagery do seem to involve – or maybe they just *are* – disagreements about conscious experience or phenomenology. The "imageless thought" controversy of the early 20th century concerned (as one might guess from its label) the possibility of conscious thought without imagery – where "imagery" was construed to include not just visual imagery but also auditory imagery (such as sentences in "inner speech" or silent tunes in one's head), kinesthetic imagery (such as in imagining the feeling of waving one's arms), and imagery in any other modality. Oswald Külpe and his collaborators (as reported in, e.g., Ach 1905 and Bühler 1907) asserted that one could have conscious thoughts that involved

nothing imagistic whatsoever. Titchener (1909) rejected this view, as did (probably) Wilhelm Wundt (1907, 1908b). (For contemporary reviews, see Angell 1911; Odgen 1911; for a recent review see Kusch 1999.[5]) Titchener claims that he regularly entertains visual images of all sorts of "unpicturable notions" such as of the notion of *meaning* – which he normally sees as "the blue-grey tip of a kind of scoop, which has a bit of yellow about it (probably a part of the handle), and which is just digging into a dark mass of what appears to be plastic material" (1909, p. 18-19). Of course, Titchener allows that others may more often think via auditory images such as the sound of spoken words. Külpe, in contrast, records thoughts like the thought that "with Darwin, chance is considered as an explanation of purpose" with no trace of any image, including of the word "Darwin" (Bühler 1907, p. 318-319; trans. from Humphrey 1951). In Chapter 7, I'll briefly discuss the contemporary revival of a similar controversy. Even Aristotle appears to take a stand, saying "the soul never thinks without an image" (if we may translate "φαντάσματος" here as "image" in the relevant sense; 4th c.BCE/1936, 431a).

Still earlier was the debate between Locke and Berkeley about abstract ideas. Locke seems to have felt that he could form an image of a triangle that is "neither oblique, nor rectangle, neither equilateral, equicrural, nor scalenon; but all and none of these at once" (1690/1975, p. 596).[6] Berkeley writes in response:

> If any man has the faculty of framing in his mind such an idea of a triangle as is here described, it is in vain to pretend to dispute him out of it, nor would I go about it. All I desire is that the reader would fully and certainly inform himself whether he has such an idea or no (1710/1965, p. 12).

It's only tongue in cheek that Berkeley grants the possibility of such a faculty; the whole point of his discussion is to demonstrate the impossibility of such "abstract ideas". Locke, he thinks, is wrong about his own experience.

Is it clear who's right in this matter? My own inclination favors Locke, though I have to admit that abstraction and indeterminacy seem to me supportable only up to a point: Would Locke allow that we can visually imagine a circle and a triangle side by side without imagining which one is on which side? That is, could we imagine its remaining indeterminate or unspecified whether the circle is on the right and the triangle is on the left or vice versa, while nonetheless picturing (if I may use that word) the two as next to each other? That gives me pause. On the other hand, I doubt that images are always as perfectly determinate as Berkeley seems to insist. Russ Hurlburt offers what strikes me as a plausible example of an indeterminate image:

> Susan, a college student, was critical of her roommate Helen's relationships with boys. Susan [reported having (when a random beeper went off to cue her to think about her stream of inner experience)] an image of Helen, seen from the waist up sitting on their couch with a boy. Helen in the image was wearing only a bra. Helen and the couch and the bra were seen clearly in the image, but the boy's face was unelaborated or indistinct (Hurlburt and Schwitzgebel 2007, p. 106).

(Contrast Berkeley: "Likewise, the idea of a man that I frame to myself must be either of a white, or a black, or a tawny, a straight, or a crooked, a tall, or a low, or a middle-sized man": 1710/1965, p. 8.) Might images be capable of indeterminacy, but only to a limited degree – as perhaps ink-and-paper sketches are (or quavery, shifting, animated sketches)?

I find that view appealing, but I worry that I'm being captured by media metaphors, in the manner I criticize in Chapters 1 and 2.

Or maybe Locke and Berkeley just had very different sorts of imagery, so that each is right about his own experience and wrong only in generalizing to other people. (Locke does say it requires "some pains and skill" to form a general idea like that of his triangle; maybe Berkeley lacked the skill or wasn't inspired to take the pains.) Maybe Titchener, too, never thought without imagery while Külpe did so regularly. Or maybe there's some confusion of words, so that though all the parties have basically the same experience, they somehow end up talking past each other. I must say that in my own case, however, the uncertainty I feel in the face of questions like those with which I began this chapter warms me to the idea that Locke or Berkeley or Külpe or Titchener may indeed be simply mistaken about his experience – perhaps blinded in part by theory, preconception, analogy.

iv.

Since people differ substantially in their perceptual and cognitive abilities, they probably also differ in their visual imagery. However, the imagery reports of apparently normal people differ so enormously that one might reasonably question the veracity of those reports. For most traits (barring defect, injury, or prodigy), human variation keeps within certain limits of normality. As the ancient Chinese philosopher Mengzi says, "When someone makes a shoe for a foot he has not seen, I am sure he will not produce a basket" (4th c. BCE/1970, 6A7).

In the 1870s, Galton, as I've mentioned, asked his subjects to visualize their breakfast tables. He instructed them to describe various features of the resulting imagery, as follows:

> 1. *Illumination.*– Is the image dim or fairly clear? Is its brightness comparable to that of the actual scene?
>
> 2. *Definition.*– Are all the objects pretty well defined at the same time, or is the place of sharpest definition at any one moment more contracted than it is in the real scene?
>
> 3. *Colouring.*– Are the colours of the china, of the toast, breadcrust, mustard, meat, parsley, or whatever may have been on the table, quite distinct and natural? (1880, p. 302)

This may indeed have been the very first psychological questionnaire; I'm aware of none earlier. Any resemblance to the questions with which I began section ii is of course not coincidental.

Galton had several hundred men and boys complete his questionnaire; and he supplemented this research with anecdotal reports from a variety of sources. This classic collection of narrative reports about imagery has to my knowledge remained unduplicated to this day (no doubt partly due to scientists' preference, which Galton himself largely shared, for quantifiable and easily replicable measures), so I'll treat it as representing the scope of opinion, even if it's dated. (If the scope of opinion has substantially changed since then, that may also support my suspicions, unless there has been a corresponding change in the actual distribution of imagery.) What, then, do Galton's respondents say?

Well, they say very different things. Some claim to have no imagery whatsoever. Others claim to have imagery as vivid and detailed as ordinary vision or even more so. Though the bulk of respondents express more intermediate views, both extremes seem to be well represented among (apparently) normal respondents. Here are some quotes from respondents at the high end: "The image that arises in my mind is perfectly clear.... I can see in my mind's eye just as well as if I was beholding the scene with my real eye" (1880, p. 310); "All clear and bright; all the objects seem to me well defined at the same time" (1880, p. 305); "The mental image appears to correspond in all respects with reality. I think it is as clear as the actual scene" (ibid.). Some respondents say they can visualize an object from more than one angle at once, like Borges's fictional character in the epigraph. One respondent says, "My mental field of vision is larger than the normal one. In the former I appear to see everything from some commanding point of view, which at once embraces every object and all sides of every object" (1880, p. 314). Galton also says that he knows

> many cases of persons mentally reading off scores when playing the pianoforte, or manuscript when they are making speeches. One statesman has assured me that a certain hesitation in utterance which he has at times is due to his being plagued by the image of the manuscript speech with its original erasures and corrections. He cannot lay the ghost, and he puzzles in trying to decipher it (1883/1907, p. 67; Titchener evidently made similar claims: see Sommer 1978, p. 44-45).

Other respondents say: "My powers are zero. To my consciousness there is almost no association of memory with objective visual impressions. I recollect the

breakfast table, but do not see it" (1880, p. 306); "No power of visualizing" (ibid.); "My impressions are in all respects so dim, vague and transient, that I doubt whether they can reasonably be called images" (ibid.). William James (who, in his classic *Principles of Psychology,* leans heavily on Galton's treatment of imagery) claims that his own powers of visual imagery are so weak that he "can seldom call to mind even a single letter of the alphabet in purely retinal terms. I must trace the letter by running my mental eye over its contour in order that the image of it shall have any distinctness at all" (1890/1981, p. 708).

One of Galton's subjects, a scientist, embarks on a critique of Galton's questionnaire itself:

> These questions presuppose assent to some sort of a proposition regarding the "mind's eye" and the "images" which it sees.... This points to some initial fallacy.... It is only by a figure of speech that I can describe my recollection of a scene as a "mental image" which I can "see" with my "mind's eye".... I do not see it ... any more than a man sees the thousand lines of Sophocles which under due pressure he is ready to repeat (1880, p. 302, ellipses Galton's).

In fact, Galton says that "the great majority of men of science" with whom he interacted at the start of his investigations "protested that mental imagery was unknown to them, and they looked on me as fanciful and fantastic in supposing that the words 'mental imagery' really expressed what I believed everybody supposed them to mean" (ibid.). Failing to find such skepticism among non-scientists – finding, instead, a general willingness to declare their imagery distinct and full of detail, even in the face of feigned

skepticism from him – Galton concludes that, contrary to what one might have expected, scientists tend to "have feeble powers of visual representation" compared to non-scientists (1880, p. 304). (In recent studies, however, Isaac and Marks [1994] and Brewer and Schommer-Aikins [2006], have failed to replicate Galton's result, finding undergraduate science majors and practicing scientists to report imagery about as vivid as that of non-scientists. This may reflect a change in culture or a subject-pool difference, or possibly a theory-driven misinterpretation or distortion by Galton of his own data, as suggested by Brewer and Schommer-Aikins and contemplated by Burbridge [1994].)

Although Galton and James assume that these self-reports accurately reflect a surprising variation in the quantity and quality of visual imagery, I'm inclined to view the reports suspiciously. You may not share such suspicion, but maybe you'll grant this: Before accepting the existence of such wide variability in the imagery of normal people, we should ask whether those with self-reported high and low imagery powers differ significantly in their success on cognitive tasks that are plausibly aided by the use of visual imagery. In this vein, James Angell (1910), discussing the imagery literature of the time, stresses the importance of looking for correlations between what he calls "objective methods" of measuring imagery, in which success or failure on a task depends on the nature of a subject's imagery, and "subjective methods" of self-report. If the correlation between subjective and objective methods is poor, Angell suggests, the differences in subjective report might be differences in *report only,* not reflecting real differences in imagery experience. And furthermore, if differences in imagery ability are as vast as they would seem from the reports of Galton's respondents, we should presumably expect vast corresponding differences in performance on imagery-involving

cognitive tasks – differences like that between a prodigy and a normal person or between a normal person and one with severe disabilities. For example, we might expect Galton's erasure-plagued statesman to show stupendous memory of the look and layout of his manuscript. Antecedently, it seems plausible to doubt that such differences will be prevalent in normal populations – but let's look at the data.

(Established prodigies do, by the way, sometimes claim to have detailed imagery of the sort that could explain their special talents [e.g., in Luria 1965/1968; Coltheart and Glick 1974; Grandin 1995; Sacks 1995]. Temple Grandin, famous for her self-description of life with Asperger's autism, claims for example that she can visually imagine and design entire slaughterhouses in her head – and consequently, once the imagination is complete, that she can transfer the design quickly and effortlessly to paper as no ordinary architect could do. In such cases, the subjective reports have at least prima facie plausibility – although it is also possible that some prodigies are confabulating details of their imagery to explain what they know to be unusual performances.)

v.

In the past hundred years, many researchers have, as Angell advised, compared subjective and objective measures of visual imagery. The results are discouraging.

Research has focused mainly on three subjective measures, more readily quantifiable descendants of Galton's original questionnaire: Betts's (1909) Questionnaire upon Mental Imagery (and Sheehan's [1967] shortened version of that questionnaire), Gordon's (1949) Test of Visual Imagery Control, and Marks's (1973) Vividness of

Visual Imagery Questionnaire (VVIQ). The early returns were very bad. Through the 1970s, attempts to correlate these subjective measures with anything objective failed so regularly that most prominent reviewers denied the existence of a relationship between objective and subjective measures of imagery (e.g., Ernest 1977; J. Richardson 1980). Even Allan Paivio, otherwise a great defender of the importance of visual imagery, concluded that "self-report measures of imagery tend to be uncorrelated with objective performance tests" (1986, p. 117). Self-reports and objective measures did not line up.

More recent reviews have been slightly more optimistic. The most thorough treatment is Stuart McKelvie's (1995) review and meta-analysis of the literature on Marks's VVIQ, the most extensively studied of all the visual imagery questionnaires. The VVIQ prompts subjects to form visual images (such as of a relative or of a rising sun), then asks them to rate the vividness of those images on a scale from 1 ("perfectly clear and as vivid as normal vision") to 5 ("no image at all, you only 'know' that you are thinking of the object"). McKelvie finds strong relationships between the VVIQ and tests of hypnotic suggestibility, tests involving the Gestalt completion of incomplete figures (e.g., the speed at which someone would recognize a fragmentary stimulus as a "circle"), and tests of motor and physiological control; he finds spotty relationships between the VVIQ and tests of visual memory; and he finds no relationship between the VVIQ and tests of visual creativity (for people of normal I.Q.[7]) or tests of ability at spatial transformation or "mental rotation". (A mental rotation task might involve judging whether the line drawings of two three-dimensional objects are such that one object would be a simple rotation of the other. Another type of spatial transformation task might involve something like judging how a sheet of paper will look when fully

unfolded.) McKelvie concludes, primarily on the basis of this pattern of results, that "[o]n balance... the evidence favors the construct validity of the VVIQ, with a more definitive conclusion awaiting further research" (p. 93).

Suppose we accept McKelvie's tentatively positive assessment. Though that might cheer us a bit, it would still seem to follow that researchers generally have failed to find the dramatic performance differences that seem implied by the wide spread of narrative reports in Galton, and consequently that the reports of Galton's subjects remain to a significant extent unjustified.[8]

My own take, however, is that McKelvie's conclusion is too sanguine. To start with, it's a bit odd to suggest that further research is needed to establish the VVIQ's validity. McKelvie's bibliography contains over 250 publications, many of which are reports of multiple VVIQ studies. The VVIQ is not a complex instrument: It has sixteen questions total, and it fits on one page. If there's any hope of establishing its validity as a measure of imagery vividness, shouldn't several hundred studies be enough to do the trick? Also, judging by where the bulk of research has been done, the four kinds of tasks that psychologists historically expected good visualizers to excel at were visual memory, visual creativity, mental rotation, and other sorts of spatial transformation tasks like unfolding. McKelvie finds no relationship between the VVIQ and any but the first of these – and even there, the relationship is at best partial and disorganized. The three tasks McKelvie reports as showing the most robust relationship between objective performance and the VVIQ have not been nearly as thoroughly studied, and are tasks whose connection with visual imagery seems prima facie more tenuous: hypnotizability, motor and physiological control, and Gestalt figure completion. Furthermore, more recent

research has called into doubt the robustness of two of these three relationships: Crawford and Allen (1996), Kogon et al. (1998), Sebastiani et al. (2003), and Gemignani et al. (2006) find no relationship between the VVIQ and hypnotic suggestibility (though Santarcangelo et al. [2005] does suggest a relationship); and Eton et al. (1998) find no relationship between the VVIQ and motor control. Although McKelvie describes the Gestalt completion findings as derived from four different studies, in fact all four of these studies are reported by a single author in a single six-page journal article (Wallace 1990). The result does not appear to have been replicated across time or laboratories. And although reports of correlations between the VVIQ and performance on various cognitive tasks presumably involving visual imagery have continued to appear since 1995, so also, perhaps even more often, have studies reporting no significant correlation between the VVIQ and visual or imagery-related tasks (for details see this note[9]).

There's reason to expect *some* positive findings even if the VVIQ doesn't accurately measure visual imagery. For one thing, as Paul Meehl (1990) has stressed, psychological variables tend to correlate, sometimes robustly, for a whole variety of reasons apart from those hypothesized by the experimenter. For example, low (vivid) VVIQ scores and good performance on a cognitive task might both be influenced by some feature of one's social background or one's personality: Upper-middle class Caucasians or extraverts (completely hypothetically), or people who tend to give excessively positive self-descriptions (somewhat less hypothetically: see Allbutt et al. 2008) may tend both to say they have vivid imagery and to do well on certain laboratory tasks, independent of any actual difference in imagery. One possibility that strikes me as especially likely (and which, it seems to me, receives not nearly enough skeptical

attention from psychologists in the imagery literature and in most other psychological subliteratures) is a possible *reactivity* between the measures.

Let me explain. Suppose you're a participant in an experiment on mental imagery – an undergraduate, say, volunteering to participate in some studies to fulfill psychology course requirements. First, you're given the VVIQ, that is, you're asked how vivid your visual imagery is. Then you're given a test of your visual memory – for example, a test of how many objects you can correctly recall after staring for a couple of minutes at a complex visual display. Now if I were in such an experiment and I had rated myself as an especially good visualizer when given the VVIQ, I might, when presented with the memory test, think something like this: "Damn! This experimenter is trying to see whether my imaging ability is really as good as I said it was! It'll be embarrassing if I bomb. I'd better try especially hard." Conversely, if I say I'm a poor visualizer, I might not put too much energy into the memory task, so as to confirm my self-report or what I take to be the experimenter's hypothesis. Reactivity can work the other way, too, if the VVIQ is given second. Say I bomb the memory (or some other) task, then I'm given the VVIQ. I might be inclined to think of myself as a poor visualizer in part *because* I know I bombed the first task. In general, participants are not passive innocents. Any time you give them two different tests, you should *expect* their knowledge of the first test to affect their performance on the second. Exactly how subjects will react to the second test in light of the first may be difficult to predict, but the probability of such reactivity should lead us to anticipate that, even if the VVIQ utterly fails as a measure of imagery vividness, some researchers should find correlations between the VVIQ and performance on cognitive tasks. (To the extent there is a pattern in the relationship between the VVIQ

and memory performance, the tendency is for the correlations to be higher in free recall tasks than in recognition tasks, as noted by McKelvie [1995]. Free recall tasks [like trying to list items in a remembered display] generally require more effort and energy from the subject than recognition tests [like "did you see this, yes or no?"] and so may show more reactivity between the measures, as well as more mediation by differences in personality or motivation.)

Psychologists have also tended to find that the experimenter's own expectations often influence the outcome of experiments – sometimes through subtle or non-verbal communications between the experimenter and the subject (Rosnow and Rosenthal 1997). Margaret Intons-Peterson (1983) found such experimenter effects specifically in imagery studies. Using advanced undergraduates as her experimenters, she led some of the experimenters to expect subjects to do better on a mental rotation task under one condition, and she led other of the experimenters to expect the reverse pattern of performance. The experimenters then read instructions to subjects from a typewritten sheet and all stimuli and responses were presented and recorded by computer, minimizing the most overt sources of experimenter influence on results. Despite these precautions, Intons-Peterson found subjects' responses tending to conform to the experimenters' (presumably) subtly communicated expectations. It's also widely recognized in psychology that positive findings, whether they arise spuriously or due to the experimenters' hypothesized causes, are more likely to be pursued and published than negative findings – the so-called "file drawer effect". Paul Chara (1992) stresses the importance of this source of distortion in imagery research.

For all these reasons, then, we ought to expect some reports of a relationship between subjective measures of visual imagery, like Galton's questionnaire or the VVIQ (or Betts's 1909 QMI or Gordon's 1949 TVIC), and performance on cognitive tasks. The question is, what do the positive findings look like? Are there mostly positive relationships between the subjective reports of imagery and skills that would theoretically be aided by imagery? Are there mostly weaker or negative relationships with skills that would presumably not be aided by imagery? Or – as we should expect if there is in fact no substantial relationship between subjective measures of imagery and actual patterns of imagery use in cognitive tasks – are the positive findings a disorganized smattering, with frequent failures of replication? To me it seems much more like the latter.

vi.

Eidetic imagery – sometimes popularly (but in the view of many theoreticians inaccurately) called "photographic memory" – has also been widely studied with the aim of finding correlations between subjective imagery report and performance on cognitive tests. Eidetic imagery is imagery of previous but now absent visual stimulation (such as of a witnessed scene or a viewed page) that is in some respects like an afterimage, but with two crucial differences: While most afterimages have colors complementary to the colors of the objects perceived (e.g., a red object will normally leave a green afterimage), eidetic images retain normal color, and while afterimages follow the eye's movement, eidetic images are motionless and scannable. Eidetic images may also be more under voluntary control than are afterimages (Jaensch 1930; Haber and Haber 1964).[10] Eidetic

imagery is measured primarily by subjective report (though some researchers, following Haber and Haber 1964, also check that direction of gaze corresponds with the relative location of the details being reported) and is attributed primarily to children. Often, eidetic images are described as being very detailed (e.g., Allport 1924; Jaensch 1930; though see Leask et al. 1969). Unlike "photographic memory" as the phrase is commonly used, eidetic imagery is defined purely phenomenologically, in terms of the imagery experience resembling (in certain ways) looking at a photograph. One might expect – researchers did expect – that people with imagery experiences of that sort would tend to have excellent visual memories; but that would be an empirical question, not something true by definition.

Early researchers on eidetic imagery sometimes claimed to find a variety of differences between "eidetikers" and non-eidetikers in personality, perception, and cognition – including of course visual memory – but the methodology was often obscure and inconsistent between laboratories (for critical reviews see Allport 1928; Klüver 1933; Gray and Gummerman 1975). For example, Gray and Gummerman (1975) state in their review that frequency estimates of eidetic imagery among children span the full range from 0% to 100%, depending in part on the methodology of the study. Later, more careful research begun and inspired by Ralph Haber and his colleagues in the 1960s resolved some of these methodological inconsistencies, but at the price of most of the positive results – so much so that in 1979 Haber was forced to concede that "extensive research has failed to demonstrate consistent correlates between the presence of eidetic imagery and any cognitive, intellectual, neurological, or emotional measure" (p. 583).

Soon thereafter, mainstream psychologists largely abandoned the study of differences between people who report and people who do not report eidetic imagery.[11]

vii.

Back to our main question: Are people, as Galton and James assumed, accurate judges of their own imagery experiences? I've offered some reasons for pessimism. To summarize: There's the ease with which most people can be brought to confusion or uncertainty about substantial features of their imagery experiences when confronted with questions like those in section ii. There's the incredible diversity of imagery reports, even among apparently normal people without unusual skills or deficits. And there's the apparent lack of any systematic relationship between differences in imagery report and differences in performance on any sort of objective cognitive test, especially tests that psychologists have historically thought likely to be enhanced by imagery, like tests of visual memory, mental rotation tasks, and tests of ability at other sorts of spatial transformation.

Of course, all this evidence is indirect. It's impossible – at least it is right now, in the current state of neuropsychology – to measure people's imagery directly. (Whether advanced neuroscience will ever definitively settle questions about imagery and other aspects of the stream of consciousness, I'm not sure. I don't dismiss the possibility, but in Chapter 6 I'll offer one reason for doubting that neuroscience will be a panacea.) So the argument I offer here has limited force: I recommend pessimism only as the most plausible interpretation of the evidence. There's ample room for the determined optimist

to stage a response. (This is, I find, true in general for claims about the accuracy or inaccuracy of judgments about conscious experience – the claims at the heart of this book. The indirectness of the evidence makes fertile ground for disagreement.) Let's consider a few avenues for the optimist, focusing on gaps in my last argument, the one that turns on the lack of relationship between subjective report and objective performance. The first three of these potential responses I'll address quickly. The fourth I'll dwell on a bit more and devote a section to, since it easily generalizes into a concern for every chapter of this book.

First, it could be that quantified versions of Galton's questionnaire like the VVIQ, and other related measures, don't really capture the aspects of imagery relevant to performance on cognitive tests. Akhter Ahsen (1985, 1986, 1987), for example, suggests that vividness is often irrelevant or even detrimental to cognitive tasks. The view has some plausibility: In rotating an imagined figure, for instance, to see if it matches another figure on the page, what would seem to matter is its gross morphology, not its vividness. Of course, not all imagery questionnaires center on vividness. Gordon's (1949) Test of Visual Imagery Control, which has also been extensively studied, simply asks respondents *whether* they form certain visual images, such as an image of a car crashing through a house. It correlates no better than the VVIQ with cognitive performance measures (recent studies, mostly negative, include Antonietti, Bologna, and Lupi 1997; González, Campos, and Pérez 1997; Lequerica et al. 2002; MacIntyre et al. 2002; Burton and Fogarty 2003; Guillot et al. 2007). This objection also doesn't address non-imagers or people with imagery as feeble and apparently useless as James claims his to be. Of course it could be that there are systemic difficulties with *all* the major visual imagery

questionnaires, but someone who wishes to mount this sort of argument at least owes us a story about why a century of trying hasn't yielded anything demonstrably better.

Or it could be that visual imagery is useless in most of the cognitive tasks we've been examining. There'd be no reason, then, to expect subjective measures, even if accurate, to correlate with the cognitive tests. An extreme version of this view would treat imagery as completely cognitively epiphenomenal (that is, causally inert): Although some people have powerful, vivid, and lifelike imagery and others have very little imagery at all, the actual mechanisms they deploy to manage cognitive tasks and to solve problems are the same. The cost of this position is that it seems to posit a major faculty with a fairly obvious range of purposes but in fact with little purpose at all, and little effect on behavior apart from the power to generate reports. The strangeness of this view is compounded if one treats subjective reports of imagery with the uncritical credulity of Galton and James, since people will often claim to have *used* imagery in a particular way to solve a problem. To the extent an advocate of this line of response mitigates extreme epiphenomenalism by allowing visual imagery to serve *some* important cognitive functions, it becomes mysterious why correlations haven't been found between measures like the VVIQ and success on any but a suspiciously desultory sprinkle of tasks.

Or it could be that both self-reported good and poor visualizers use imagery, but only good visualizers experience that imagery consciously. This position is a variation of the previous one, except that what is epiphenomenal is not the imagery itself but the conscious experience of it. People seem ordinarily to think of imagery as consciously experienced, as part of the flow of phenomenology, but maybe a suitably functional approach can give some sense to the idea of an unconscious image (as in Paivio 1971).

However, unless conscious experience is epiphenomenal, people whose imagery is mostly conscious ought to perform somewhat differently on cognitive tasks than people whose imagery is largely unconscious; so it remains strange that such differences have not been found. Maybe consciousness *is* epiphenomenal, or at least largely so, but such a view faces the challenge of explaining why whatever biological or functional facts permit some cognitive processes but not others to be conscious seem to have so few other ramifications.[12] Locating the top of the scale also creates challenges for this view. To fully credit subjects' reports, we would have to take reports of extremely detailed and vivid imagery as the benchmark of fully conscious imagery and assume that every subject has imagery (perhaps partly unconscious) at roughly that level of detail. Otherwise, one must either grant that there are substantial differences in the degree of detail in subjects' imagery after all, thus rekindling the original problem of explaining the lack of correlation between subjective report and cognitive test, or grant that the subjects at the top of the scale have overestimated their imagery, which would mean granting just the sort of error for which I'm arguing. But if everyone's imagery has the level of detail described in the grandest self-assessments, then it's surprising that we don't all perform substantially better on mental rotation tasks, visual memory tasks, and the like.

viii.

Even if Galton's (and others') uncritical acceptance of such reports is folly – or, to put the point more politely (since I really do enjoy Galton), even if people's differences in

self-report generally fail to reflect underlying differences in their imagery – we may still know our imagery perfectly well. Maybe the problem is mostly one of verbal expression.

Galton asks his respondents "Is the image dim or fairly clear? Is its brightness comparable to that of the actual scene?" (1880, p. 302). Marks tops the VVIQ scale with the phrase "perfectly clear and as vivid as normal vision" (1973, p. 24). What do such phrases mean? In interpreting them, at least two problems arise. (For similar concerns see J. Richardson 1980; Cornoldi 1995.) First, it's not entirely obvious what it is for an image to be "clear" or "vivid". Is the question whether the images are *detailed?* Sharply contoured? Salient? Full of saturated color? Lively? Forceful? Some or all of these? Respondents may understand the question very differently. Second, there's the problem of comparing clarity across different types of experience. When I visit the optometrist and she asks if what I see through one lens is as clear as what I see through another, I understand the question. Since I'm comparing two sensory visual experiences, what it is to be "clear" remains the same across the cases. But if I'm asked to compare the clarity of my vision without glasses to the clarity of an orchestra heard through a wall, the matter isn't so straightforward. Although visual imagery and visual sensation presumably have some phenomenological commonalities, they also seem to differ significantly, making it unclear what the criteria are for saying that a visual image is as clear and vivid as normal vision. (Exactly how sensation and imagery differ phenomenologically is a matter of dispute. [Need I add "of course"?] Hume notoriously suggests that images differ from sensory experiences mainly in being fainter [1740/1978, p. 1-2], while most scholars seem to hold that imagery experiences differ from sensory experiences in kind, not just intensity. Titchener describes images as also having a "textural difference from

sensation... more filmy, more transparent, more vaporous" [1910/1915, p. 199]. McGinn says that compared to percepts, images are indeterminate, "gappy, coarse, discrete" [2004, p. 25]. This of course doesn't exhaust the alternatives.) So, respondents interpreting "vividness" differently or using different standards for the comparison of clarity may have similar imagery experiences, which they accurately apprehend and yet describe differently.

Consider also Galton's skeptical scientist who finds a fallacy in supposing the existence of a "mind's eye" that "sees" images. Taking "see" literally, the scientist is surely right: There is no homunculus who literally sees your images. Yet there also seems to be a looser or metaphorical sense in which it's permissible to say that we see our visual imagery. Maybe, then, the difference between the scientists' and non-scientists' responses to Galton's questions reflects neither differences in their imagery (as Galton supposes) or epistemic failure (as I suggest) but only differences in how strictly they interpret the word "see".

These verbal difficulties are mostly what we might call *between-subjects* difficulties. *Within-subjects* measures – that is, measures that compare different instances of imagery within individuals over time – should avoid at least some of these problems, as long as the subject interprets the questionnaire items consistently over time. When someone reports that one of her images is different from another, that report, even if incomparable with the reports of other subjects, may reflect a real difference between those two instances of her imagery, and thus possibly a real cognitive difference. Unfortunately, few researchers have explored such within-subjects differences in imagery (exceptions include Bower 1972; Walczyk 1995[13]), and the methodological difficulties

are daunting. (For example, if a subject rates an image as vivid and then remembers it better than an image rated as less vivid, there are many possible explanations for this relationship, only one of which is that the memory performance shows she had accurately discerned the phenomenological vividness of the image.)

I accept that the lack of clear reporting standards accounts for some of the variation in reports; in fact, I believe that it's a huge problem. However, I don't think that can be the whole story. Although the standards of "vividness" may be particularly murky, especially across different experience types, other questions involve no such problematic language, unless even the term "image" is problematic. (And if the term "image" *is* so problematic that its use invalidates all imagery questionnaires, I doubt the optimist about introspective accuracy will find much consolation in that fact.) Gordon's Test of Visual Imagery Control, as I've mentioned, simply asks respondents whether they can form visual imagery of particular situations, like a car crashing through a house. And although Galton's skeptical scientist may be a mere quibbler over words, if a respondent *can* form an image of his breakfast table, it would be perverse to deny that fact – the fact Galton is clearly after – based on concerns about how the question is phrased. Are we to suppose that all Galton's imageless scientists were so perverse? Likewise, the disputes between Locke and Berkeley, between Titchener and Külpe, don't *seem* to be disputes merely over the use of words. There are substantive phenomenological issues in the vicinity – issues the resolution of which, I find, is not entirely obvious, judging by my own experience. Barring a compelling reason to suppose otherwise, I recommend taking the disputes at face value, rather than recasting them as miscommunication. And if they are genuine disputes, and if we assume that the disputants did not differ radically in their

imagistic phenomenology, then some of the parties must have been pretty far wrong about their own experience.

ix.

It's extreme, of course, to suppose that there's absolutely *no* correlation between what people say or think about their visual imagery and their actual experienced visual phenomenology. In particular, I see nothing that inclines me to doubt that people have at least a bit of a handle on *what it is,* roughly, that they are visually imagining, when they report visually imagining something. (Though I wonder: Why not doubt even this? Is it my own and others' seemingly unshakeable confidence about such matters? Could that confidence be misplaced?) However, regarding more general features of imagery experience – structural or presentational features, we might say – such as its vividness, its degree of indeterminacy, its color saturation, its spatial location or flatness, its picturelikeness, the empirical evidence is discouraging. If there's any relationship between our subjective judgments about such matters and our actual phenomenology, for some reason our experimental techniques don't seem to be getting at it.

The explanation I'm drawn to is that our judgments about our experience just aren't in fact very well aligned with the experiences we actually have: We tend, simply, to get it wrong, to be captured by our own assumptions, our metaphors, by what it seems appealing to say in the face of strange questions. This is, I think, a natural interpretation of the experimental evidence. And it harmonizes well, I think, with the introspective and anecdotal considerations offered in the first few sections of this chapter. I find in my

own case – and so also apparently do many (but not all) other people, when I've interviewed them about such matters – that it's not entirely obvious how vivid and detailed my visual imagery is, how determinate or indeterminate, how narrowly it confines itself to my scope of attention, how richly colored, etc. I feel in myself, and I think you the reader might also feel in yourself, the potential for error, the liability to be swept up by a theory or a picture or a set of background assumptions. I find the introspection of visual imagery *difficult,* if I set about it conscientiously. So I think we should be unsurprised if people – including maybe you and me – can go badly awry.

[1] On such "non-imagers" see Faw (1997, 2009) and Thomas (2008 [http://www.imagery-imagination.com/non-im.htm]). Thomas is more skeptical than Faw about purported non-imagers, though he doesn't dismiss the possibility (personal communication 2009). I recommend his 1989 account of the behaviorist John Watson's apparently theoretically-driven shift from claiming that he had vivid visual imagery to claiming that he had none. Russ Hurlburt (personal communication 2002) says that several self-described non-imagers have reported visual imagery when given beepers and interviewed about randomly sampled moments of ordinary experience.

[2] If you aren't familiar with phosphenes, press gently on the corner of one eye. A little spot should appear in the opposite corner (perhaps a gray spot with a bright ring). Wiggling your finger a little to make the spot move can help bring it out.

[3] The view that there is often no determinate fact about the extent to which imagery is indeterminate has not been widely discussed, though some readers of this chapter have urged it upon me. A theoretical attraction to meta-level indeterminacy may arise from a view on which talk about visual imagery is a useful fiction and no fact of the matter whether a fiction that posits a determinate number of stripes is more useful than a fiction that does not; or it might arise from a view on which visual images are insufficiently stable to support predications about the determinacy or indeterminacy of their features over even the smallest duration of attention; or it may grow from an account of the nature of vagueness or from some other motivation. If one accepts some such species of higher-level indeterminism about visual imagery, one might deploy it to explain why people are often baffled by the kinds of questions posed in the second

paragraph of this section – but this explanation must be handled delicately if it is meant to preserve the view that introspective judgments about visual imagery are largely accurate, since many people are confident in their diverse, and on this view not determinately true, assessments of their imagery experiences.

[4] Reisberg, Peason, and Kosslyn (2003), however, argue, partly based on a retrospective survey, that at least in the early period of the debate around 1980, differences in researchers' own imagery experiences partly predicted their theoretical stances, with theoreticians who had more vivid imagery (as measured by the VVIQ: see below) being more attracted to the pictorial view.

[5] Monson and Hurlburt (1993) argue that the actual experiential reports of the subjects on both sides of the debate were very similar, only interpreted differently by the disputing parties. If so, that tangles the issue at hand; and yet I think it remains the case that – regardless of what their subjects may have said – Titchener and Külpe genuinely disagreed about the phenomenology of thinking.

[6] Since Locke says "idea", not "image", it is possible to interpret him as thinking of the idea of the triangle as non-imagistic. However, the standard interpretation of Locke seems to be that ideas, in his view, are always imagistic (though not always visually imagistic). In any case, James (1890/1981), Huxley (1895), and others clearly acknowledge the possibility of images with vague or indeterminate features, so they could substitute for Locke as opponents to Berkeley if necessary.

[7] McKelvie finds three studies (all reported in one article: Shaw and DeMers 1986) suggesting a relationship between VVIQ scores and visual creativity for subjects of high I.Q. Interestingly, he finds a parallel result for *verbal* creativity: No relationship to

the VVIQ unless participants are specially selected for high I.Q. What to make of these results is unclear. In light of the vast number of studies McKelvie considers and the various sources of spurious correlation (discussed below), I wouldn't put too much weight on these results.

[8] McKelvie does find less variability in imagery reports than one might expect from reading Galton: McKelvie's meta-analysis yielded a mean VVIQ score of 2.307 and a standard deviation of 0.692 on the 5 point scale. Demand characteristics of the survey may explain some of this trend toward the low (vivid) end of the scale. As Ahsen (1990) notes, the survey begins by asking the subject to "think of some relative or friend" and then to "consider carefully the picture that comes before your mind's eye". The latter phrase implies that a picture-like image will be experienced. However, Galton's survey employs similar language. It's possible that the narrative format of Galton's questionnaire was more encouraging of extreme responses than is the five-point scale of the VVIQ. Or maybe cultural or subject-pool differences explain the apparent decline in the variability of self-reports of imagery. Loaning beepers to thirty people and asking them to report on ten randomly sampled moments of experience, Heavey and Hurlburt (2008) found very high variability in rate at which imagery was reported – ranging from 0% to 90% of the sampled experiences.

[9] Reports of correlations between the VVIQ and visual or imagery-related tasks include: Wallace, Allen, and Propper 1996; Campos and Fernández 1997; Campos and Fernández 1998; Keogh and Markham 1998; Riske, Wallace, and Allen 2000; Lanca and Bryant 2001; Petrova and Cialdini 2005; Keng and Lin 2006; Hills, Elward, and Lewis 2008; and Lobmaier and Mast 2008.

Reports of no statistically detectable relationship include: Antonietti, Bologna, and Lupi 1997; Campos and Pérez 1997; Campos, Pérez, and González 1997; Eton et al. 1998; Antonietti 1999; Heaps and Nash 1999; Kunzendorf et al. 2000; Lewis and Ellis 2000; Tomes and Katz 2000; Kilgour and Lederman 2002; Laeng and Teodorescu 2002; Pérez-Mata, Read, and Diges 2002; Burton and Fogarty 2003; Dean and Morris 2003; Sebastiani et al. 2003; Kozhevnikov et al. 2005; Gemignani et al. 2006; Wyra, Lawson, and Hungi 2007; Gyselinck et al. 2009.

Studies with mixed results include: Crawford and Allen 1996; Campos and Fernández 1997; Tomes and Katz 1997; Winograd, Peluso, and Glover 1998; Walczyk and Taylor 2000; Eberman and McKelvie 2002; Blajenkova, Kozhevnikov, and Motes 2006; Mast, Merfeld, and Kosslyn 2006; and Carretti, Borella, and De Beni 2007.

I exclude from these lists studies looking only at the relationship of the VVIQ to other self-report measures (e.g., personality measures or other imagery measures). Although I cannot claim that the above constitutes a complete list of published studies after 1995 that attempt to relate the VVIQ to visual or imagery-related abilities, it should include most of the work published in ISI-indexed journals.

A related issue is whether there is a relationship between VVIQ score and brain activity while engaging in visual imagery. Amedi et al. (2005) and Cui et al. (2007) instructed participants to form visual images while in an fMRI machine. They found that participants scoring toward the vivid end of the VVIQ exhibited greater differences between activity in visual and non-visual areas of the brain during the imagery than those reporting less vivid imagery on the VVIQ. Although these results are encouraging, Ganis et al. (2004, clarified by a personal communication) and Schienle et al. (2008) fail to find

such a relationship, and Kosslyn et al. (1996) report mixed results. Given all the potential sources of spurious correlations, I think we have to regard the issue as open.

[10] This section concerns what has come to be called "typographic eidetic imagery" as opposed to "structural eidetic imagery" of the sort posited (but not in my view very clearly characterized) by Ahsen (1977), Marks and McKellar (1982), and Hochman (2002). It's only for the former that individual differences have been broadly studied.

[11] A few studies continued to be done. Kaylor and Davidson (1979), Paine (1980), and Miller and Peacock (1982) report somewhat better memory performance among self-described eidetikers, while Wasinger, Zelhart, and Markley (1982) report no difference, and A. Richardson and DiFrancesco (1985) report a non-significant trend. Kunzendorf (1984) reports electroretinogram differences and differences in physiological control between eidetikers and non-eidetikers; Matsuoka (1989) finds eidetikers to report more absorption in sensory and imaginative experiences; and Glicksohn, Steinbach, and Elimalach-Malmilyan (1999) suggest a connection between eidetic imagery and synaesthesia. For retrospective personal reports of frustration in searching for relationships between self-reported eidetic imagery and performance on objective tasks, see Furst 1979 and Sommer 1980. Also, as mentioned at the end of section iv, work is occasionally done on established prodigies, who sometimes claim high imagery capacities.

[12] Some recent philosophical discussions of epiphenomenalism include Flanagan 1992; Chalmers 1996; Nichols and Grantham 2001. Note that the purest kind of metaphysical epiphenomenalism is not what's at issue here. Consciousness could be metaphysically epiphenomenal in the sense that it itself has no causal power while still

being nomologically connected to some causally efficacious process such that we could always discern, by appeal to the causally efficacious process, whether the accompanying conscious process was present. Such metaphysical epiphenomenality would not explain the lack of correlation between imagery self-reports and performance in cognitive tasks.

[13] I exclude from discussion here the literature on the "bizarreness" of imagery and its memorability (e.g., Einstein and McDaniel 1987), since bizarreness seems be more closely related to the strangeness of the situation depicted than to the phenomenological features of the image itself.