

The Unreliability of Naive Introspection

Eric Schwitzgebel
Department of Philosophy
University of California
Riverside, CA 92521-0201
eschwitz@citrus.ucr.edu
909 787 4288

November 11, 2002

The Unreliability of Naive Introspection

ABSTRACT

I argue that most people are unreliable introspectors of their own conscious experience or “phenomenology”. Considerations are offered for regarding as unreliable our naive introspection of negative emotions, peripheral vision, echolocation, thought, and visual imagery in particular. I conclude by suggesting not that we abandon introspective methodologies in consciousness studies but rather that we take steps to assure that introspection is done well.

The Unreliability of Naive Introspection

With the recent surge of philosophical and psychological interest in consciousness has come a renewed reliance on introspection as a means of studying the mind. After all, introspection appears to be the most direct and straightforward way to learn about consciousness. Many people find it intuitively plausible that introspection is a reliable, perhaps infallible, means of coming to know about one's own current conscious experiences.

Nevertheless, I will argue in this paper that uncritical reliance on introspective reports about conscious experience is a serious mistake. Most people are unreliable introspectors of their own conscious experience or "phenomenology". We are not just unreliable in assessing the causes of or processes generating our mental states, as Nisbett and Wilson (1977) have famously argued; and we're not just unreliable when we are distracted, or passionate, or inattentive, or self-deceived, or pathologically deluded, or when we're reflecting on minor matters, or about the past, or only for a moment, or about matters near the threshold of discriminability. In my view, we make gross mistakes about what one might think were obvious features of our phenomenology, even in favorable circumstances of extended attention.

I will devote the bulk of this presentation to describing a few cases in which I think it's plausible to regard people as apt to introspective error. Along the way, I'll ask you to do some introspection yourselves. I'll conclude with a few thoughts about what hope there might be for an introspective methodology in the scientific study of consciousness.

1. Introspecting Emotional Experience.

Let's start with introspection of emotional phenomenology. This may or may not be the same thing as the introspection of emotion itself. Perhaps to have an emotion is just to suffer a certain kind of conscious experience – or perhaps behavioral dispositions and cognitive patterns are more central. But whether one regards emotional phenomenology as constitutive, partly constitutive, or incidental to the emotion itself, we can ask, how good are we at introspecting this phenomenology? My answer is: We stink at it. Consider the stock example in discussions of the introspection of emotion: The purple-faced man banging his fist on the table, shouting that he is not angry. It is easy enough to imagine in this case that the man really does have the phenomenology of anger as well as the outward manifestations – he is not a placid lake inside. It is also easy to imagine that were he queried not about whether he was angry but simply about whether he had the phenomenology typical of anger, he would deny that too.

Of course, this is not the ideal case for my thesis, since this man is presumably not reflecting on his phenomenology in favorable circumstances of quiet attention, and my thesis requires that even people in favorable circumstances can be radically mistaken. So consider instead a more subtle case. Those of you who, like me, are prone to soul-searching conversations with a psychologically-minded spouse, will surely have entered disagreements about the presence or absence of some particular negative emotion in yourselves. Can you honestly say, even in cases in which you have reflected attentively on your own feelings, that you have consistently had the better of the argument? Now, it may be that in some of these cases the emotion and the characteristic phenomenology have come apart – that is, maybe despite really being angry (or whatever), in the sense of having enough of the physiological responses, behavioral dispositions, etc., to qualify as

angry, you genuinely lack the phenomenal feeling of anger. In such cases, you may be a perfectly adequate introspector of your own emotional phenomenology without having got it right about your overall emotional state. Maybe this is possible, but it seems to me quite as reasonable to suggest that one can sometimes have the conscious experience or phenomenology of anger without recognizing it, even on reflection – perhaps because one doesn't like to see oneself as angry. If so, one's introspective judgments about one's own current emotional experiences can be mistaken. In that case, although the emotion may be “unconscious” in one sense – in the sense that the person possessing it refuses to acknowledge its existence – it may nonetheless be conscious in the primary, phenomenal sense of ‘consciousness’ that I am employing in this paper: It seems plausible to allow that an unacknowledged emotion may be genuinely felt; it may be part of the person's stream of experience. [*]

In my estimation, cases of this sort are quite common, and not just with negative emotions. People may be quite unaware of, and apt to deny when questioned, even positive emotional phenomenology if it accompanies what they see as negative events or burdensome tasks. For example, one person has reported to me that her mother regularly complains about having to garden, describing it as an unpleasant chore, yet when she is out in the garden seems remarkably happy. I think it's plausible to imagine that if queried in the garden about her emotional phenomenology, this woman might sincerely deny the presence of the joy she actually experiences. Despite undergoing enough vivid and salient emotional experiences to have a rough sense of the phenomenology of at least the most basic emotions, I suggest that people tend to have a rather poor idea of their own emotional phenomenology in particular cases.

Reflect now on your own current emotional phenomenology. Is it completely obvious to you what it is? Does introspection reveal it to you as clearly as visual observation reveals the presence of the larger objects in your current field of view? Many philosophers, at least since Descartes, have thought that our knowledge of our own current conscious experience is secure, solid, and indubitable in a way that our knowledge of physical objects in the world around us is not. Considering cases like this, I am inclined to think that epistemically that position is quite backwards. [*]

2. Introspecting Visual Experience.

One might assert, in response to my last suggestion, that although our introspective judgments about our emotional phenomenology may not be as secure as our visual judgments about nearby physical objects, our introspective judgments about our visual phenomenology are very secure, perhaps even serving as the epistemic foundation of our ordinary visual knowledge of external objects. And indeed, barring mistakes of categorization, I must admit that normally, when I am looking directly at a bright red object in good light, it is rather hard to imagine that I could be wrong when I reach an introspective judgment that I am having the visual phenomenology of redness.

At least some cases of introspection of visual experience, then, seem to be what we might call “obvious” cases in which it would be difficult to go wrong. Of course, this by itself does not show that introspection is more reliable than perception of the external world, since the external world also yields obvious cases in which mistakes are highly unlikely. Could you be mistaken about the presence of people here in the room, or the existence of the object you have been writing with and are presently examining?

Now of course philosophers historically have been fond of pointing out that you could be wrong in such judgments about external objects, if one admits the possibility of rather far-fetched scenarios, such as one's brain having been removed at night and teleported to Alpha Centauri where it is being stimulated by genius neuroscientists so as to mimic the inputs of normal interaction with the world; and perhaps there is no equally vivacious thought experiment that reveals the possibility of your being mistaken about having reddish phenomenology. All I can offer instead is the speculation that whatever part of your cognitive system is responsible for the judgment "I'm having reddish phenomenology now" and whatever part of your cognitive system is responsible for the red phenomenology itself are conceivably separable enough that a genius scientist could somehow stimulate the first while suppressing the second and so generate in you the false judgment. [*]

Wherever those reflections might lead us, when we consider the experience of peripheral vision, a case can be made for massive introspective error. Consider your visual experience right now. Does it seem to you that you experience a fairly broad field – say maybe thirty degrees of arc – flush with color and detail, in which shapes are precisely defined, and that only outside that thirty-degree field is there a substantial decrement in clarity? When I casually introspect my own visual experience, that is the judgment I am tempted to make, and I have found others to report similar phenomenology. It is now firmly established, however, that the precision with which we see shape and color falls off precipitately outside a foveal area of approximately 1-2 degrees of arc (about the size of your thumbnail held at arm's length). Daniel Dennett has suggested one way of demonstrating this to yourself. Take the face cards out of a

normal deck of playing cards, shuffle them, and draw one without looking at it. Hold your eyes fixed on some point in front of you, take the card in one hand, and hold it at arm's length just beyond your field of view. Keeping your eyes fixed, slowly move the card toward the center of your visual field, noting how close to the center you must bring it before you can determine the color of the card, its suit, and its value. Most people are quite surprised at the result of this little experiment. They substantially overestimate their visual acuity in the central thirty or so degrees of the visual field.

Although this is what one would expect if we are mistaken about our visual experience, it does not strictly imply that we are mistaken. It is one thing to make a judgment about one's abilities and quite another to make a judgment about one's phenomenology. In particular, it seems plausible that our visual phenomenology depends not just on our visual input at any particular moment but also on our visual input in the recent past and on our general knowledge of the world. If so, then we really might experience clarity in a region broader than our foveal area. Holding one's eyes still is unnatural, and perhaps it artificially reduces the region of clarity.

Nevertheless, I doubt that the level of precision and detail in our visual experience is anything like what people naively take it to be. As others have suggested, the mistake leading to the naïve view of visual experience seems to be this: When the thought occurs to someone to attend to some part of her visual phenomenology, she normally moves her eyes (or "foveates") in that direction. The consequence is that wherever she thinks to attend, within a certain range of natural foveal movement, she finds the clarity and precision of foveal vision. I'd like you again to introspect your own visual phenomenology. To block the error just described, I ask you to observe that you can in

fact attend to parts of your visual experience that are not in your immediate foveal region. It is easiest to achieve this by fixating on some particular point in the distance, holding your eyes steady while your introspective attention wanders. However, since this artificial fixation may reduce the region of clarity, it would be even better if you can turn your attention away from the foveal area while your eyes continue to move around (or “saccade”) normally. I have found that with a little practice it is easy enough to attend to non-foveal regions while saccading. For example, you can attend steadily to one non-foveal part of your visual field while you allow your eyes to foveate anywhere but there.

Performing such introspective experiments on themselves, most of the people I have spoken to eventually conclude, to their surprise, that clarity does decrease sharply even a few degrees from the foveal region. I suggest, then, that it is possible introspectively to discover that rather than a broad, stable field of clarity with fuzziness only at the borders, visual experience consists of a clear focal center that bounces rapidly around a rather indistinct background. If I am right about this, and if I am right that most naive introspectors think otherwise, then most naive introspectors are grossly mistaken about their own visual phenomenology, even though they may be patiently reflecting on it as it is happening.

Even if you think I am wrong in my view of visual experience, I have at least established this: Sincere and reflective people can introspect their own current visual experience and come to vastly different conclusions. Unless the visual experience itself is radically transformed by the theoretical commitments or introspective techniques used by the different parties, or unless there are huge individual differences in visual

phenomenology that underwrite the different conclusions, then at least one of the disagreeing parties must be seriously mistaken.

3. Introspecting Echolocation.

Another sensory experience we all have, but tend to be badly mistaken about when we introspect, is echolocation. Can you hear how far away a wall is? Can you hear the shape of a podium when you are lecturing at it? Research in auditory perception suggests that the answer to these questions is yes. Human beings have non-trivial capacities at echolocation – that is, at detecting the properties of silent objects by noticing how they reflect sound generated by other sources. Although blind people tend to be more practiced, both blind and sighted people can use echolocation to avoid colliding with objects in novel environments; blind people have been shown to be able to detect the direction of target objects through echolocation with substantial accuracy; both blind and sighted people can detect differences in shape between targets of equal surface area – such as a triangle versus a square; and some blind children can even echolocate well enough to ride a bicycle around a neighborhood street, avoiding parked cars, poles, and even curbs.

Furthermore there is something it is like to echolocate. Echolocation is not merely a nonconscious capacity. To see this, try the following experiment. Close your eyes and hold your hand a few inches from your face (a friend's hand would be even better, but I'll assume that's not convenient right now). While making a "shhh" sound, move your hand around. You can hear where your hand is. I hope you'll agree as you do this that you are not merely hearing the sound of your own voice, or even hearing that sound as modulated

by the position of your hand, but actually hearing the position of your hand, just as you might hear the position of a sound source as to the right or left, closer or farther away.

If colleagues in philosophy and cognitive science with whom I've discussed these matters are any guide, most people will deny that they have any substantial echolocatory capacity or experience, unless particularly vivid examples such as this are pointed out. Besides the hand-in-front-of-the-face test, most people recognize that there is a substantial change in their auditory experience when they are driving on the freeway and a wall by the side of the road comes to an end, or when they sing in the shower, then try it again in the hallway. Showers and hallways sound different. But not because they produce different sounds, rather because they differently reflect sound produced from other sources. But if echolocatory phenomenology is vivid in cases of this sort, it is not implausible to suppose that it is also present, more subtly and pervasively, in much of our auditory experience. And now that I am attuned to it, my own introspection seems to reveal this to be the case. Try it yourself later as you are moving around, going about your normal day: You can hear that the door is near your left ear as you pass through it; from the sounds of your reflected voice, you can hear that the room you just entered is full of soft, sound-absorbing objects rather than just a few hard tables on a tile floor.

We habitually attend to objects in the world and only rarely to the experiences those objects produce in us. Nonetheless, people's obliviousness to echolocatory experience is not wholly a consequence of the fact that people generally don't pay attention to their sensory phenomenology. Even when people's attention is drawn to their auditory phenomenology, they generally fail to notice this aspect of their experience, unless their attention is drawn first to some really salient piece of echoic phenomenology, as in the

hand-in-front-of-the-face demonstration. This fact is most striking in the case of blind people who, though often actively using echolocation to avoid colliding with objects in novel environments, historically tended to attribute this ability to feelings of pressure on the face rather than auditory sensations. They called the capacity “facial vision”. In some experiments, blind subjects were asked specifically to reflect on the sensory modality used in avoiding collisions, and they were advised that echolocation was one possibility. Even in these circumstances, in which subjects expert at echolocation were actively and repeatedly attending to their auditory experience, aiming to discern anything echolocatory, still not all subjects reported echolocation. It might be suggested that if these subjects reported no echolocatory phenomenology, they must not have had any; but it seems to me rather more plausible to suggest that they had the same echolocatory phenomenology as the rest of us, perhaps even more vividly than most sighted people, yet completely failed to notice it.

4. Introspecting Visual Imagery.

Another favorite case of mine is visual imagery. Form a visual image of some familiar scene – the front of your house, say, as viewed from the street. If you are able to do this (some people claim not to be able) then presumably you are having a conscious experience of visual imagery. Now consider the following questions. How stable is the image? Do all its elements endure fairly steadily over time or do they fade quickly without renewed attention? Does the chimney fade as your attention shifts to the front door? Do all the parts of the image have color at once, or do you have to work to assign color to them a bit at a time? Exactly how much detail is there in the focal part of the

image (if there is one), and how swiftly, if at all, do color and shape fade toward the periphery? How similar is your current imagery to standard visual perception, or to a dream or daydream, or to phosphenes or afterimages? Many people, myself included, find these questions and others like them, rather difficult to answer. Despite their being questions about basic features of current conscious experience, they are hard questions, questions about which it seems reasonably likely that a well-meaning, reflective person could go wrong.

It turns out that when asked to describe such features of their visual imagery as its vividness, its similarity to vision, and the breadth of the imagistic field, people give very different answers. In samples of apparently normal people, answers have been found to range all the way from complete denial of the existence of visual imagery to reports of imagery as vivid and as detailed as ordinary vision, or even more so. Some respondents have even reported being able to visualize objects simultaneously from multiple angles.

Yet all these vast differences in report correspond very little with differences in performance on tasks that psychologists have thought to involve visual imagery, such as mental rotation tasks, visual creativity tasks, and tests of visual memory. Although several explanations could be offered for this finding, such as that these tasks are not actually aided by visual imagery or that conscious visual imagery is epiphenomenal, I would suggest that the best explanation is one that challenges the veracity of the subjects' reports. This should not seem too implausible if you accepted the conclusion, based on the questions I asked a minute ago, that accurate introspection of visual imagery is difficult and if you find yourself warming to my general thesis that introspective reports

of conscious experience are not sacrosanct and unchallengeable but rather quite prone to error.

Emotion, sensory experience, and imagery constitute a substantial portion of our phenomenology. I suggest that naive introspective judgments about each of these aspects of our phenomenal life are prone to error, that even in circumstances of attentive reflection, we easily make gross mistakes in introspecting basic features of these experiences. Introspection is harder than inspection. Discerning the gross features of your house is easier than discerning the gross features of your visual image of your house, or your visual experience as you look at your house, or even the sometimes complex emotional phenomenology you may experience upon arriving home after having been away.

5. Introspecting Thought.

Some people regard the experience of thinking as another major portion of our phenomenology, distinct from the types so far mentioned. Others have argued that there is no phenomenology of thought distinct from the phenomenology of imagery, perhaps supplemented with feelings like discomfort, familiarity, or confidence. David Chalmers polled seminar participants (all professional philosophers) after a weeklong series of sessions devoted in part to this topic, and found 8 respondents to say that there is no phenomenology of thought, or that its phenomenology is exhausted by the phenomenology of imagery, and 17 respondents to say that there is a distinct phenomenology of thought. Before reflecting ourselves on the existence or absence of a phenomenology of thought, it is worth remarking that unless the disagreement here is just

a disagreement over the use of words, which I think is not the case, or unless experiential differences between people are so large that some have a phenomenology of thought and some do not, this disagreement implies that some of these philosophers are profoundly mistaken about their own stream of conscious experience. If there is such a thing as a conscious thought, then presumably we have them all the time. How could someone go looking for them and simply not find them? Conversely, if there is no distinctive phenomenology of thought, how could someone introspect and come to believe that there is – that is, invent a whole category of distinctive conscious experiences that simply do not exist? Such fundamental mistakes seem to beggar the imagination and plead for interpretation as something other than what they initially appear to be.

They do, that is, until one attempts the introspection oneself. Think of the Prince of Wales. Now consider: Was there something it was like to have that thought? Perhaps you experienced a visual image of the Prince, or an auditory image of yourself saying the words “Prince Charles” in silent soliloquy. Neither of these is the thought. You might have the same visual image while considering the (nonexistent) twin of Prince Charles; you might have the same verbal image while thinking of some other person named “Prince Charles”. The question is whether there is something further in your phenomenology, beyond the auditory and visual imagery, that is the distinctive phenomenology of thinking. Regardless of whether you answer yes, no, or maybe (my preferred answer), I hope you’ll agree that it is not introspectively obvious. People might go wrong about it. Disagreements about this subject need not be limited to disagreements about the use of words or about philosophical abstracta – despite the fact that getting it

wrong about the presence or absence of a phenomenology of thought would be a fundamental mistake about one's own stream of consciousness.

Such mistakes needn't imply that we get it wrong about what we are thinking. We may have some introspective access to, or non-introspective awareness of, our thoughts without those thoughts having any distinctive phenomenology. Reports of current thoughts may even be self-verifying: If you say that you are thinking of some topic, the mere fact that you comprehendingly utter that assertion may make it true, even if it wasn't true a moment prior. But let me emphasize that even if our reports of our current thought contents are generally accurate, it by no means follows that our reports on the phenomenology of thinking are likewise accurate.

6. Introspecting Attitudes.

Usually when someone sincerely says she believes something, she believes it. So also for desires, fears, intentions, doubts, and most of what philosophers of mind have called the "propositional attitudes": People generally don't get it wrong about their attitudes. [*] This is not to say that people aren't sometimes self-deceived or possessed of "unconscious" attitudes the presence of which they are apt to deny. I have argued elsewhere for a dispositional approach to belief in particular, and I see no reason to think it impossible for someone sincerely to judge that she believes that P while otherwise so thoroughly lacking the dispositional structure appropriate to the belief that P that we rightly deny her the belief. But for whatever reason, this sort of thing happens rarely enough that we are generally safe accepting someone's sincere word about her attitudes. (Framing this observation within a dispositional account, I am inclined to suggest that the

stereotypical dispositions associated with the belief that one believes that P contain as a subset the stereotypical dispositions associated with the belief that P.)

Consequently, it seems to me that our judgments about our attitudes are more secure and reliable than our judgments about our phenomenology, including our current phenomenology. This seems to be just the reverse of the standard opinion. Of course, it is not that I have an unusual view of the reliability of our judgments about our attitudes, but rather that I have an unusual view of the reliability of our judgments about our phenomenology.

I see three ways to reconcile the reliability of the one sort of judgment with the unreliability of the other. One would be to say that judgments about attitudes are typically introspective judgments about phenomenology, but the phenomenology is so overwhelmingly obvious that it is difficult to be mistaken about it, despite our general ineptitude at phenomenological introspection. This I highly doubt. Perhaps sometimes the judgment that one believes or desires something is a consequence of introspection of an unmistakable auditory verbal image – something uttered silently to oneself, like “Paris is the capital of France” or “Chocolate ice cream, that sounds good”; but this cannot serve as a general explanation of our attitude judgments, since (among other difficulties) our judgments about our attitudes are often quicker than that. There is no time for us to withhold judgment, pausing to “listen” to our auditory images and then reach conclusions about our attitudes based on what we “hear”. But other than the phenomenology of auditory imagery, I find no overwhelmingly obvious phenomenology that comes anywhere close to reliably attending my judgments about my attitudes.

If our judgments about our attitudes are not based on phenomenological introspection, then two possibilities remain: They are based on introspection, but not introspection of phenomenology, or they are not based on introspection at all. In the normal case, I suspect things work somewhat as Richard Moran describes: Reflection on one's attitude toward some proposition P – whether it is believing, desiring, intending, or whatever – is simultaneous with reflection on the truth, desirability, or whatever, of P itself, such that the judgment one reaches about P serves as a kind of update of one's attitude toward P. In thinking about whether you believe Bush will be re-elected, you think about whether Bush will be re-elected. The judgment you reach then both serves as an update of your attitude about Bush's re-election and as a report of your attitude. In abnormal cases, you might fail to bring your overall standing attitude into line with your current judgment, and then you will be mistaken in your attitude report. There will also be times one takes a more clinical look at oneself, assessing one's genuine attitude by looking, non-introspectively, at one's patterns of behavior, which may not accord with one's current reflective judgment. Also, as Victoria McGeer has emphasized, we might work to make the reports true by acting in accord with our previously avowed attitudes. In none of these cases is the accuracy of the report due to introspection.

But possibly also some cognitive mechanism exists that allows us to detect our attitudes directly. If so, it should reveal itself most clearly when one is pondering oneself in a somewhat clinical mood, bracketing the question of the appropriateness of the attitude, and furthermore basing one's judgment not wholly on recollecting externally observable evidence about one's behavior. I'm not invested in denying that such introspectively informed assessments do occur, and perhaps they involve to some extent

the introspection of phenomenology – but the accuracy of attitude reports is plausibly explainable without reference to that sort of introspection. What I would suggest then is this: To the extent self-reports of attitudes are more accurate and reliable than judgments about phenomenology, they are grounded in something other than phenomenological introspection.

7. Introspective Methodology in the Study of Consciousness.

I promised to conclude with some remarks about introspective methodology in consciousness studies. In psychology, two attitudes toward introspective methodology dominate. The first tends to take subjects' introspective reports at face value. The second rejects introspective methodologies altogether, or regards them as substantially less valuable than standard behavioral or physiological measures. Given my remarks so far, it would be natural to conclude that I advocate the second of these two positions. Unfortunately, however, I think that despite the unreliability of naive introspection, it is impossible to avoid giving introspection a central role in the study of conscious experience. Without introspection, we might not even know that we are conscious in the relevant sense – that a stream of phenomenology accompanies our outwardly observable behavior. Behavioral and physiological measures alone tell us nothing about consciousness unless it is established that those measures correlate with conscious experience; and introspection is the most straightforward way to establish those correlations. All of our tools for understanding conscious experience are problematic in their application. (Maybe this is why consciousness studies has been so slow to find a firm scientific footing.) It is not as though in the face of the unreliability of introspection

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

Our situation is in some ways analogous to that of a foreign intelligence agency that must depend for its information on a network of unreliable, double-crossing spies. Just as the reports of spies can to some extent be corroborated or cast into doubt by such independent means as satellite photos and bank records (which by themselves may say little), so also can introspective reports be to some extent checked against behavioral, physiological, and cognitive measures; and just as consistency or inconsistency between the reports of independent spies provides at least prima facie reason to accept or doubt the reports, so also consistency or contrast between independent introspective reports, when there is no reason to suspect corresponding differences in conscious experience, may justify tentative acceptance or rejection of the reports.

Besides introducing external checks of this sort on introspective reports, steps can also be taken to assure that individual introspective judgments are as well justified as one can reasonably expect. For example, the introspector can work to guard against common sources of introspective error – as in my instructions to you earlier regarding the introspection of visual experience outside the foveal region. In some cases, a certain amount of introspective training may be called for, as Wundt and Titchener of the introspectionist school of psychology insisted a century ago. I have been struck by how much introspective training can help one discern features of one's auditory experience, such as the experience of echolocation and of overtones. Of course, risks accompany such training. The experience after the training may be in important respects different

from the experience prior to the training. One probably does not hear in exactly the same way once one is able to discern overtones. Also, excessive theoretical bias may be introduced – a fault that helped sink the introspectionist school, where differences in the reports of introspective experts sometimes correlated suspiciously well with geographical location or the theoretical position of one’s teacher. However, the desire to avoid these difficulties should not lead us to the other extreme, embodied by the Gestaltist maxim that the naive introspector is generally the best subject.