

Mark Engelbert and
Peter Carruthers

Descriptive Experience Sampling: What is it good for?

***Abstract:** We defend the reliability of Hurlburt's Descriptive Experience Sampling method against some of Schwitzgebel's attacks. But we agree with Schwitzgebel that the method could be used much more widely than it has been, helping to answer questions about the nature and structure of consciousness in addition to cataloguing the latter's contents. We sketch a number of potential lines of further enquiry.*

1. Introduction

What can the method of Descriptive Experience Sampling (DES) tell us about the human mind? What sorts of questions can it be used to investigate with any degree of reliability? This is the central question addressed in Hurlburt and Schwitzgebel (2007),¹ and it is the question on which we propose to focus this essay. Specifically, in addition to commenting on the reliability of the DES method for the purposes discussed in the book, we are interested in exploring ways of using DES to shed light on a wider variety of psychological and philosophical questions. We will, therefore, be sketching a number of proposals for research strategies that utilize DES, both in the method's current form and with modifications.

Correspondence:
TO BE COMPLETED

[1] Hereafter page numbers without a date refer to pages within Hurlburt and Schwitzgebel (2007), and references to either Hurlburt or Schwitzgebel by name alone should be understood to refer to their individual views as expressed in the book. When we refer to any of their other individual writings we use the usual author/date method.

One sort of modification would be as follows. Hurlburt counts among the virtues of DES its ‘ecological validity’, or the fact that it captures subjects’ experience in their natural, everyday environments. The commitment to ecological validity precludes beeping subjects in the laboratory, and also rules out instructing them to focus their reports on particular kinds or aspects of experience. While we recognize the value of this ecological validity for some purposes — particularly when forming generalizations about the prevalence and variability across individuals of different modes of experience — we feel that the DES method has the potential to help answer important psychological and philosophical questions, some of which may require dropping the commitment to ecological validity. Although these aren’t questions that have interested Hurlburt himself, they may be of interest to others, and we would urge other investigators to take them up, modifying the DES method as appropriate. As we explore various ways of putting the DES method to use, some of our suggestions will sacrifice ecological validity for the sake of other virtues that are derived from traditional psychological methods, especially the use of experimental controls.

2. Is DES reliable at all?

First, however, we address the question whether experience sampling is appropriate for the main use that Hurlburt has made of it over his career — namely, examining the patterning and diversity of conscious experience in daily life (both within and between individuals). Hurlburt (pp. 21, 27–31) provides a list of features that, to his thinking, make DES a method far superior to ‘armchair introspection’ for obtaining reliable data about conscious experience. Subjects are beeped randomly, which reduces the complications that accompany attempts to introspect with an intention to introspect (Schwitzgebel experiences precisely this difficulty when trying to capture the phenomenology of his own inner speech). Furthermore, subjects produce written notes immediately upon being beeped and are later debriefed by an experienced interviewer. In order to reduce experimenter bias, subjects are permitted and encouraged to report whatever features of their experience they find salient and important; they are allowed to decline to answer any question that they wish; and interview questions are exploratory and open-ended. Subjects are also trained by the experimenter during the first couple of days of sampling, during which time they become comfortable with the introspection process and are taught to avoid common pitfalls.

Schwitzgebel acknowledges these virtues but thinks that attempts at introspection under the DES paradigm may still be subject to systematic inadequacies. Indeed, Schwitzgebel (2008) challenges the accuracy of introspection itself. He questions, for example, whether emotional states like joy have a distinctive common phenomenology that can be identified introspectively across instances. Note that such worries are irrelevant to the reliability of the DES method, however (and to the reliability of real-time introspection generally — see Engelbert and Carruthers, 2009). This is because accuracy in this respect would require *generalizing* about one's experience over a significant number of occurrences, comparing one's experience from one instance to the next.² This isn't something that introspection can accomplish unaided. Plainly it also requires memory. But there is reason to think that memories of introspected experiences may not be preserved on a routine basis, as we explain shortly. In contrast, DES, as practised by Hurlburt, specifically shies away from such '*faux* generalizations'. It instead focuses only on subjects' statements about their experiences at the moment of a particular beep. The generalizations can then be provided by the experimenter, without needing to rely on the subject's memory.

Why should one think that memories for introspected experiences are unlikely to be retained? The most common suggestion for why we possess a capacity for introspection in the first place is that it has a monitoring function (e.g. Shallice, 1988). On this account, we monitor our own mental processes in order to intervene in them when they go awry, or to trouble-shoot when they get blocked. This plainly requires that some sort of short-term record of one's mental states should be retained, so that the monitoring mechanism can represent each stage as an event in an ongoing process. Hence subjects should be capable of reporting their immediately past mental states. But we should predict that representations of one's own mental states will *not* be stored in long-term memory, unless for some reason they are rehearsed and/or consciously attended to and revisited. For this isn't necessary to support the trouble-shooting function, and would serve no useful purpose. Rather, knowledge of our immediately past mental events should fade away rapidly, just as dreams do. And indeed, consistent with this prediction, many of Hurlburt's subjects make discoveries about the patterns in their inner experience that surprise them,

[2] By no means do all of the arguments in Schwitzgebel (2008) require generalizing across experiences, however. For some discussion and critique of his other arguments, see Engelbert and Carruthers (2009).

suggesting that long-term memories of such experience aren't routinely created.

As regards the DES method in particular, Schwitzgebel worries considerably about the extent to which subjects' reports are influenced by their own self-theories and the metaphors they employ in describing hard-to-capture features of experience. He also raises issues about memory: subjects may forget such features as the level of detail that is present in their visual images, and whether or not there was any experience represented within a particular sense modality. Furthermore, subjects' reports may be distorted by their revisitations of their beeped experiences during the time between the beep and the follow-up interview.

Our own view is intermediate. We believe that there is every reason to trust the results obtained through DES, but only (on current evidence) when the method is addressed towards the presence or absence of fairly gross categories such as inner speech, visual imagery, and so forth, together with their approximate contents (i.e. what they are about). Since such features (or lack thereof) can be noted within seconds of the beep, and will generally be recorded in the subject's own notes, worries about the reliability of memory, situational demands, and so on, appear to be only minimally applicable.

The more fine-grained the categories employed, however, the more scope there is for scepticism of the above sorts. But relatively straightforward manipulations of the DES method itself should be able to test for such distorting effects. Schwitzgebel himself (2007) was able to address the issue of biasing by subjects' background theories and expectations through the simple expedient of asking subjects about these in advance. He was able to discover no such effects. Likewise, if one were concerned that much of the richness and detail in subjects' descriptions of their experience might be a product, either of the constructive nature of memory following a 24-hour delay (during which time subjects will no doubt have revisited the beeped moment numerous times in their thoughts), or of the pragmatic demands of extended questioning by an experimenter, then these factors could be controlled for. One might have subjects interviewed at varying intervals after the beep (one hour, four hours, eight hours, and so on) to see whether those interviewed following longer intervals tend to produce greater richness and detail in their descriptions of their experience. And one might have interviewers adopt a pair of strategies, one of which would be the usual extended probing, but the other of which would be just a single request to a subject (following initial training) to tell everything that they can remember about the experience in question, perhaps

even having them speak privately into a tape recorder rather than to an interviewer (this manipulation could also be performed with both trained and untrained DES subjects, to help answer questions about the effects of training).

Such controls and manipulations would not completely eliminate concerns about confabulation and elaboration, since there may be individual differences in how susceptible subjects are to situational and memory demands. But the experimental controls should reduce such concerns to acceptable levels. It is an empirical matter what effect such manipulations might have, but our tentative prediction after reading Hurlburt and Schwitzgebel's book is that they will not be large.

3. What do current DES findings tell us?

Hurlburt's main finding over his career has been that people's inner experiences are extremely diverse. Some people seem to spend much of their time engaged in inner speech, some in manipulating visual images; some experience unsymbolized thinking, some don't; some report emotional feelings, some don't; a handful report rich multimodal experience, most don't; and so on. Is this an additional source of scepticism about the reliability of the DES method, as Schwitzgebel alleges? He claims that the biological commonalities between all (or almost all) humans should lead us to expect that their inner lives will be basically similar, just as their digestive systems and respiratory systems are similar. We disagree. Even if there is in some sense a default preference for assuming similarity among individuals, there is more than enough evidence in the present case to indicate that individual variation should come as no surprise. Or so we will now argue.

For one thing, although this is by no means uncontroversial, most psychologists believe that the contents of consciousness are products of *attention* (e.g. Dehaene *et al.*, 2006). One form of attention is stimulus driven, or 'bottom-up'. Thus a sudden loud sound or a snake-like shape in the grass can *grab* our attention, forcing the contents in question into consciousness. But another form of attention is top-down and varies depending upon our current goals and ongoing executive functions. Attending in this latter sense is something that we *do* — although not necessarily deliberately or consciously, of course. People with exactly the same biological and cognitive systems might therefore develop very different habits of top-down attention, which might lead some to undergo frequent emotional experiences while others have frequent visual experiences. Both sets of people might equally be undergoing *emotions*, and both will possess visual

perceptions of the environment on a routine basis, just as their common biology and cognitive architecture would lead one to predict. But the groups will differ in whether states of the two kinds are likely to achieve the ‘global broadcast’ necessary for conscious status (in the sense of Baars, 1988),³ resulting from their different habits of attention.

Moreover, many of the kinds of items reported by DES subjects (particularly inner speech and visual imagery) are best thought of as belonging to so-called ‘System 2’ cognitive processes. Scientists who study human reasoning abilities have increasingly converged on the hypothesis that we possess a set of swift and automatic systems that are largely immutable and universal (and frequently shared with other animals), which produce initial intuitions and intuitive answers to reasoning questions (System 1). But we also have a slow, limited capacity, consciousness-involving system (System 2), which operates under intentional control (see Evans and Over, 1996; Sloman, 1996; 2002; Stanovich, 1999; Kahneman and Frederick, 2002; Kahneman, 2003). System 2 utilizes the centralized working memory system investigated over the years by Baddeley and colleagues (Baddeley and Hitch, 1974; Baddeley and Logie, 1999; Baddeley, 2006), and frequently employs visual and auditory imagery (especially inner speech). It is already known that the character and extent of use of System 2 varies widely between individuals (Stanovich, 1999). Moreover, Carruthers (2006; 2009a) argues that System 2 depends crucially upon mental rehearsals of action-schemata (thereby explaining how System 2 processes are under intentional control). So in our view, the discovery that some subjects spend most of their waking lives engaged in inner speech whereas others occupy most of their time manipulating visual images should be no more surprising than the fact that some people spend much of the day engaged in physical activity whereas others are mostly sedentary, or that some people gesture with their hands a lot while speaking whereas others don’t. And certainly this discovery isn’t grounds for scepticism about the reliability of the experience sampling method itself.

Although Schwitzgebel is sceptical of the patterning of Hurlburt’s results, he is actually more interested in the question whether experience sampling can be used to answer deeper questions about the

[3] For some more recent evidence supporting the global broadcasting account of consciousness, see Dehaene and Naccache (2001), Baars (2002), Baars *et al.* (2003), Dehaene *et al.* (2003). The framework is now widely accepted among consciousness theorists. Even those like Chalmers (1996) who deny that it provides us with an *explanation* of the phenomenal properties of consciousness nevertheless accept that global broadcast is causally necessary and sufficient for consciousness.

character of consciousness, such as whether people's experience frequently has multimodal contents (tactile experience at the same time as visual and auditory experience, and so on), and whether there is generally a rich background of conscious experience at the periphery of attention. Thus Schwitzgebel wants to know whether one generally has conscious experiences of one's feet in one's shoes, or of the upper left quadrant in one's visual field beyond the window of current attention. We applaud the spirit of Schwitzgebel's approach. We, too, think that DES is an under-utilized and under-appreciated resource, which could be used or adapted to answer a much wider array of questions. But on this specific issue (richness *versus* sparseness) we are inclined to side with Hurlburt, as we will now explain.

Hurlburt tells us that the vast majority of his subjects deny the richness of their inner experience. In general, a sampled moment contains only one or a few experiences at a time. Schwitzgebel is sceptical. For the fact that randomly sampled reports are mostly sparse is consistent with richness plus quick fading. Although this is conceivable, we don't think that the suggestion is a plausible one. For in that case one would at least expect subjects to report that they have the *sense* that there was a lot more going on, but they can't remember what. Compare the way in which dreams tend to fade almost immediately upon waking. Even if one can report very little of the content of the dream, one generally has the sense that it faded rapidly, and that it actually contained much more than one can now recall. Schwitzgebel therefore needs to explain why waking experience should be so different. If experience is characteristically rich, why is it that subjects don't have an introspective sense of that richness rapidly fading away? While one might adduce other evidence for the richness of conscious experience, we think that Hurlburt's results using the DES method provide plausible *prima facie* evidence for the sparseness view.

It is important to distinguish between two different senses in which experience can be rich, however. The first is the one that we have just been discussing, which depends upon the extent to which conscious experience is multimodal in character, as well as on the extent to which there are conscious experiences outside of the window of focal attention. Let us call this 'peripheral richness', to be contrasted with what might be described as 'focal richness'. The latter concerns the richness of experience *within* the window of focal attention (which is both spatially and temporally smeared, it should be noted). The two forms of richness are independent of one another. This enables us to preserve Hurlburt's DES finding of peripheral sparseness in the face of evidence of focal richness. Subjects who briefly view a complex

array on which they focus their attention (e.g. three rows of three different numbers) can characteristically only report a small proportion of the items thereafter. But they also report having *seen* all of the items in detail. And if any one row of items is signalled immediately following the removal of the stimulus, subjects are at ceiling in their reports, suggesting that all of the relevant visual information was initially represented in consciousness, just as the subjects themselves claim, but that it fades rapidly (Sperling, 1960; Landman *et al.*, 2003).

Schwitzgebel (2007) himself set out to find evidence on the question of peripheral richness by devising an adaption of the DES methodology. What he found would support a moderate position if taken at face value, containing significantly more riches than Hurlburt's subjects characteristically report. So the question now becomes: can we trust *Schwitzgebel's* experience sampling method? We believe there are good reasons to think that one can't. The crucial feature of the new methodology is that some subjects receive instructions to report specifically on a particular aspect of their peripheral experience at the moment of the beep (e.g. the upper left quadrant of the visual field, or somatic experience in the left foot). But the problem with the new method is this: there is good reason to believe that perceptual contents that are preconscious are held briefly in an iconic memory store, and that those contents can *become* conscious if attended to (Dehaene *et al.*, 2006). Even if a subject has *no* conscious experience of his foot at the time of the beep, information from the foot will have been processed up to a certain level, and might be reverberating in a preconscious iconic memory system. As soon as the beep goes off the subject will direct attention towards that system as instructed, and its contents will thereby *become* conscious. We conclude, therefore, that Hurlburt's finding of the peripheral sparseness of subjects' experience can be allowed to stand.

4. Some possibilities for expanding the use of DES

As we noted above, Hurlburt's main interest is in the character and variety of people's conscious experience itself. But it should be obvious that the DES methodology (whether modified or unmodified) can be applied much more widely (here we are firmly in agreement with Schwitzgebel). One obvious application of the method would be to investigate correlations between patterns of conscious experience and other cognitive traits and abilities (this could in turn be a springboard to more structured sorts of investigation of the cognitive architecture underlying various forms of conscious experience, or indeed

consciousness itself). The method has obvious advantages over more traditional survey methods, such as asking people whether their experience is generally rich or not (Hurlburt is surely correct that all you can really get from such questioning are unreliable *faux* generalizations).

For example, Hurlburt himself (Hurlburt *et al.*, 2002) has recently investigated whether faster speakers tend to report more detail in their inner experience generally, finding that they do. This is an intriguing result. It is natural to wonder whether the discovered correlation might be mediated by individual differences in the capacity for fast attention switching. This would have an obvious effect on speed of speech generation. And it would likewise enable multiple contents to enter the temporally extended window of conscious experience. There is reason to think that conscious visual experience, in particular, is built up serially over a period of a second or two through multiple saccades. But something similar might also be at work across different sensory modalities, enabling people to juggle with visual and auditory imagery within the same time-window, while also devoting attention to their emotional feelings, and so on. One can imagine various ways in which the hypothesized differences in fast attention switching might be tested.

In addition, we think that experience sampling might be used to investigate whether the richness and detail of people's reported images correlates with ability in some of the standard psychological tasks requiring imagery, such as mental rotation. Schwitzgebel (2002) assumes, very naturally, that there should be a correlation, and Hurlburt (p. 275) seems inclined to agree. But since the existing studies have for the most part failed to find any, Schwitzgebel (2002) uses this as grounds for suspecting the reliability of introspection itself. The studies in question did not use DES, however. They relied upon more traditional forms of introspective report, with all of the latter's drawbacks. If one can find a way to quantify the richness in imagery reported by DES subjects, then it would be worth re-examining the issue by conducting batteries of imagery-related tasks with individuals who report varying levels of richness and detail in their images.

We are not entirely convinced of Schwitzgebel's (2002) assumption that differences in experience should correlate with differences in task performance, however. For whether there should be such a correlation depends, at least in part, on whether images are constructed serially and actively or whether they can, rather, spring into consciousness fully-formed (this question is raised at various points in Hurlburt and

Schwitzgebel's book).⁴ If the former is the case then we might expect to find a correlation with image manipulation tasks, since there will be close parallels between the processes employed in subjects' everyday imagistic thinking and deliberate image manipulation. But if images can arrive instantaneously, then there would be no reason to expect such a correlation. For it would then seem that the capacity to *manipulate* images is something quite different. Anecdotally, at least (warning: we are about to make a *faux* generalization), it appears that images can often arrive fully-detailed, especially in the case of memory images. It seems that an evocative smell or word can sometimes bring to mind an image of a scene from one's past in a single flash.

More ambitiously, then, one might try to adapt the experience sampling method to address this very question. We could (in randomly distributed trials) beep at specified intervals of time following the experimenter's instruction to imagine a particular object. If the object has a known number of basic features, we could compare the results against a baseline for each subject. If images arrive fully formed then we should expect to find no correlation between richness and elapsed time, or even perhaps a negative correlation, as aspects of the initial image are forgotten. If images are constructed serially, in contrast, then we should expect to find a positive correlation with elapsed time. Note that the methodology needed here might be burdensome, however. For the test trials would need to be distributed with sufficient infrequency that subjects aren't continually 'on the watch out' for the occurrence of the beep. And it might be hard to combine the method with probes of *evoked* memory images, as opposed to ones that are formed on command. Our larger point, however, is that until researchers interested in the structure and functioning of cognition take up the DES method and adapt it for their purposes, it is hard to know what its limits might be.

5. DES and the introspectability of attitudes

We think that DES could also be used to investigate a quite separate issue, which isn't discussed directly by Hurlburt and Schwitzgebel. This is the question whether there is any such thing as introspection for attitudes. There is an extensive literature in psychology suggesting that we often engage in self-interpretation when attributing propositional

[4] Note that serial construction of images might provide an explanation of the otherwise puzzling phenomenon of images that violate rules of perspective (p. 86). If it turns out that images are *constructed*, and constructed serially, then it wouldn't be surprising if subjects might sometimes combine together image components that presuppose different and conflicting perspectives.

attitudes to ourselves (see Festinger, 1957; Bem, 1967; 1972; Wicklund and Brehm, 1976; Nisbett and Wilson, 1977; Eagly and Chaiken, 1993; Gazzaniga, 1995; 2000; Wegner, 2002; Wilson, 2002; Briñol and Petty, 2003). While most agree that such confabulation data are successful in showing that we *sometimes* engage in self-interpretation, there is considerable debate over whether this is the *only* method of self-attribution, or if we sometimes have direct introspective access to our own attitude states as well.

In recent contributions to this debate, Carruthers (2009b; 2010) has argued that there is no such thing as introspection for propositional attitudes. Perceptual experiences (e.g. visual, auditory, proprioceptive, and somatosensory) and quasi-perceptual experiences (such as visual imagery and inner speech and/or hearing) can be introspected, on Carruthers' model. This is because (and to the extent that) they are globally broadcast, hence being made available as input to the mindreading faculty. But propositional attitudes like beliefs, desires, judgments, decisions, wonderings, and supposings must be self-attributed via processes of self-interpretation, similar to the processes that are employed when we attribute attitudes to other people. But in one's own case the mindreading faculty can not only use information about one's own behaviour and circumstances, it can also utilize 'internal' data in the form of visual imagery, inner speech, and so on, enabling it to make an (unconscious) inference as to the current mental state of the agent even in the absence of overt behaviour.

Here we offer some suggestions for how the DES method might be used to gather additional evidence for or against the existence of introspective access to attitudes. All of these suggestions are tentative, and none of them would be expected to settle the issue decisively. Rather, data from the experiments we propose would be useful supplements to the large corpus of experimental work on the prevalence of confabulation. We focus first on two areas where robust confabulation effects have been reported: split brain subjects and subjects who have undergone hypnosis.

Split brain subjects have had their corpus callosum (which is the structure connecting the two brain hemispheres) severed as a treatment for epilepsy. Because the hemispheres of split brain patients cannot communicate with one another, stimuli presented in the right visual field (which feeds into the left hemisphere) are not perceived in the right hemisphere and *vice versa*. Since language is usually lateralized to the left hemisphere, any decisions taken or intentions formed in response to information presented to the right brain will not be available to the subject's language comprehension and production

systems. When queried about actions they have performed in response to information presented to the right brain, these subjects often offer verbal reports that suggest they are engaging in unconscious self-interpretation. For example, Gazzaniga (1995) reports an episode in which the instruction 'Walk!' was presented in a subject's left visual field (and hence to the right brain), which the subject proceeded to follow: he got up and began to leave the testing van. When asked what he was doing, the speech-controlling left brain replied, 'I'm going to get a Coke from the house'. This report is plainly confabulated, but seemingly with all of the apparent-introspective-obviousness as normal.

Similarly, Wegner (2002) documents a number of instances of post-hypnotic behaviour to show that when asked about their intentions in performing an action for which they were given instructions under hypnosis, subjects will often confabulate a reason. For example, one subject was hypnotized and instructed to pick up a book from the table and place it on a shelf, and when asked why she did it she responded, 'I don't like to see things untidy; the bookshelf is the place for books, so that is why I am placing it there' (Wegner, 2002, p. 149). A plausible interpretation of such cases is that since the subject has no knowledge of her *real* intention (that is, to follow the hypnotist's instructions, which she does not remember), she interprets her own behaviour to devise a likely explanation: that she is tidying the room.

Some have argued against drawing strong conclusions about the prevalence of self-interpretation on the basis of such experiments. Critics point to a variety of issues with these experiments that purportedly weaken the force of their results. One worry is that interrupting subjects mid-action and requesting an explanation of their behaviour introduces a pragmatic demand on them to *justify* their behaviour or make it appear *rational* (Rey, 2008). Subjects in such situations might thereby be pushed into self-interpretative mode as a result of the situational demands, which wouldn't impugn the introspective nature of spontaneous, unsolicited self-attributions. The DES method could therefore be used in the context of such an experiment to reduce or eliminate such pragmatic demands by providing a more indirect way of having subjects report their intentions. It could be arranged so that the beep goes off just as a subject is performing an action prompted by post-hypnotic suggestion or by an instruction flashed to the right brain of a split brain subject (it would be best if this could be done using already-trained subjects as part of a larger sampling exercise, for purposes of ecological validity). Subjects could then note the contents of their experience, either to an experimenter present in the laboratory or using a notebook or tape recorder.

The conditions of this experiment could vary, with some subjects left simply to report whatever is ‘in their awareness’ at the time of the beep and others being explicitly asked to report their intentions. Those in the former condition may well not report any intentions at all, in which case their responses would not be relevant to the hypothesis under investigation. For the subject’s intentions might well be *introspectable* even if they aren’t actively being introspected at the time of the beep. Instructing subjects beforehand to report their intentions at the time of the beep, however, runs the risk of introducing the very pragmatic demands that the proposed use of the DES method seeks to avoid. For subjects may feel pressured to confabulate an intention that makes their behaviour appear rational. This risk might be mitigated, however, by using trained DES interviewers together with subjects who have already undergone DES training. For Hurlburt is adamant that the open-ended and unstructured nature of DES questioning makes it transparent to subjects that it is perfectly OK to report nothing if there is nothing to report.⁵

Experiments with hypnotized subjects would be especially informative if they could be designed to involve *ambiguous* actions. For example, subjects might be given a post-hypnotic suggestion that when a designated individual (who is known to the subject) walks into the room, they are to move their right hand back and forth in the air, palm outward, to wave away any bugs. Subjects could then be beeped just as they begin to wave, to see what sort of intention is reported. If intentions are introspectable, subjects in such circumstances should report that they are intending to move away any bugs. But if, as Carruthers suggests, intentions are attributed to oneself by interpretation, then subjects ought to reach for the most reasonable interpretation of their own behaviour — which in this instance is that they are waving to greet the person who has just walked into the room. Using the DES beeper to solicit subjects’ reports would, again, remove most of the pragmatic and memory demands that are alleged to have contributed to confabulation effects in previous experiments.⁶

-
- [5] Note that if subjects are operating under a general instruction to introspect their intentions at the time of the beep, then it would have to be made clear to them whether they are supposed to report a local intention, like placing a book on a shelf, or a more global one, like tidying the room (depending on the details of the experiment).
- [6] We admit that the example described here is not ideal. For even if there is introspection and subjects do have access to their bug-removing intentions, they might get pushed into self-interpretation mode when they find themselves waving as if to deter bugs and yet they can plainly see *that there are no bugs*. The example is merely meant to illustrate the general strategy of using ambiguous actions in hypnotism cases. Since subjects differ in their degrees of suggestibility, moreover, we would need to take care that they have no recall of

Another worry about many confabulation experiments is that subjects are often asked to report their intention in performing an action *retrospectively* (i.e. they are asked, ‘why *did* you do that?’). One might, therefore, argue that subjects’ real intentions were *introspectable* at the time but didn’t happen to be *introspected*. That is, it could be that the introspective faculty is only capable of detecting occurrent or very recently past attitudes, and that unless one directs one’s attention to an attitude while or shortly after it occurs, no record is thereafter stored in memory. Beeping subjects in the midst of an action for which a confabulated explanation might be expected would allow us to test this hypothesis. If subjects still offer confabulated explanations even when they are beeped *during* an action, it would suggest that confabulation data cannot be dismissed by invoking the introspectable/introspected distinction.

Beyond these suggestions for further research, we should briefly address evidence from existing DES studies that bears on the issue of introspection for attitudes. Two closely related bodies of data are the existence of what Hurlburt calls ‘partially unworded speech’, on the one hand, and ‘unsymbolized thinking’, on the other (p. 141). The former is comparatively rare, Hurlburt tells us, whereas the latter occurs commonly in some subjects (while being absent in others). Nevertheless, the mere existence of either one might seem to count against Carruthers’ self-interpretational account of self-knowledge. We will discuss them in turn.

In cases of partially unworded speech, subjects report fragmented inner speech tokens at the time of the beep, but nevertheless feel that they know the complete content of the underlying thought or speech intention (i.e. they have a conscious sense of what words belong in the missing parts of the token). For example, when one of Melanie’s inner-hearing tokens was interrupted by a beep, she reported having a clear sense of how the sentence was going to end. She had just realized that she had forgotten to take off the parking brake and heard, ‘Why can’t I...’ just at the moment of the beep. She reported knowing that the sentence was going to end with ‘...remember about the parking brake?’ (pp. 135–136). Do cases like this suggest, as Schwitzgebel (p. 137) claims, that there is introspective access to one’s thoughts beyond the imagery that one experiences?

Not necessarily. Carruthers argues that self-interpretation is a global affair, drawing on evidence not just from internal imagery but

the conditions involved in their hypnosis. Otherwise we will get the contrary problem: subjects may report their real intention from memory rather than from introspection.

also facts about the agent's behaviour and situation (in the latter regard operating much like third-person mindreading). Thus, just as a third party observing Melanie might predict, given her situation, that her inner speech token would involve the parking brake, so Melanie is able to make that prediction about her own experience. Hence as long as the content of an inner speech token could be inferred from a subject's situation, partially unworded speech poses no threat to Carruthers' model. And indeed, consistent with that model, Hurlburt reports that 'the most frequent experience of inner speech involves simply the speech itself with no conscious sense of what is about to be said' (p. 138). Moreover, he tells us that when people are beeped in the midst of speaking they generally have no conscious awareness of what they are intending to say.

In cases of unsymbolized thinking, in contrast, subjects report propositional thoughts in the absence any mental symbols like visual imagery or inner speech. It happens that Melanie is one of the many subjects who report no instances of this phenomenon. But unsymbolized thinking is by no means uncommon, and Hurlburt and Akhter (2008) describe a number of cases in some detail. Many instances can be handled by Carruthers in the manner just outlined above. For there will often be features of the context, and/or the agent's own behaviour, that would make self-attribution of the thought in question entirely natural (and something that a third party, too, might hit upon as an interpretation), independently of any imagistic cues. But this is by no means true of all cases of unsymbolized thinking. For example, Hurlburt and Akhter describe how Abigail was wondering, at the time of the beep, whether her friend Julio would be driving his car or his pickup truck when he came to collect her. Not only was this thought not symbolized in inner speech or other imagery, but there was nothing in Abigail's physical or behavioural context to suggest it.

Hurlburt himself (2009) suggests that unsymbolized thinking is consistent with Carruthers' self-interpretational model, however.⁷ For to say that someone is undergoing unsymbolized thinking is to say that

[7] Indeed, Hurlburt (2009) goes further, claiming that DES data actually *supports* Carruthers' self-interpretational model. But here we think he oversteps the mark. His reasoning is that DES subjects never report any awareness of an attitude at the moment of the beep (except in the early stages of training, in which case they are inclined to back off such claims in discussion with the interviewer). But in making this claim Hurlburt must have in mind medium-term or standing attitudes like *intending to go out to dinner this evening*, or *believing that the economy will soon turn around*. For DES subjects *do* report *momentary* attitudes in cases of unsymbolized thinking, and they *don't* back off these claims. One will report *wondering* something (as in the case of Abigail, described above), whereas another will report *urging* herself to do something, and so on. But Carruthers (2009b) doesn't claim only that *standing* attitudes are self-attributed through self-interpretation. On the

there is no sensory awareness of any imagistic symbols at the time of a self-attributed thought. But ‘sensory awareness’, for Hurlburt, is a technical term, referring to sensory information that is at the focus of attention. He therefore suggests that ‘the apprehension of an unsymbolized thought may involve the apprehension of some sensory bits, so long as those sensory bits are not organized into a coherent, central, thematized sensory awareness’ of the sort that would be revealed in a standard DES interview (2009, p. 30). It is quite possible, therefore, that unsymbolized thinkers do have fragmentary imagistic awareness at the moment of the beep that could aid in a process of self-interpretation, leading to the attribution to themselves of a particular thought. Since subjects are unaware of the self-interpretation process, but find themselves inclined to attribute a specific thought to themselves, they will have the sense that they are consciously thinking that thought in an unsymbolized way.

Carruthers (2009b) proposes an alternative way of responding to the challenge posed to his views by instances of unsymbolized thinking.⁸ This involves noticing that even if there is *no* imagery or other sensory information available in consciousness at the time of the beep, such information may nevertheless have been available just a moment previously. Since the DES method focuses only on the moment of the beep, and subjects are urged to ignore everything that came before it or after it, we (presently) have no way of knowing. So for all we know, there may have been visual or other images occurring shortly before any reported instance of unsymbolized thinking, which would have been sufficient for the mindreading faculty to construct the reported (unsymbolized) thought through the usual process of self-interpretation. And again, since subjects are unaware of the self-interpretation process, but find themselves inclined to attribute a specific thought to themselves at the time of the beep, they will have the sense that they are thinking that thought in an unsymbolized manner.

This leads us to our final experimental suggestions, the first of which is to adapt DES to probe the explanations of unsymbolized thinking outlined above. Subjects who are already known to have a high proportion of unsymbolized thoughts could be asked to sit quietly in a dimly lit room allowing their minds to wander, thus ensuring that there will be no overt behaviour or aspects of the context

contrary, the thesis extends also to momentary ones like wondering, supposing, and urging.

[8] The two suggestions are consistent with one another, it should be stressed. Some instances of unsymbolized thought might be arrived at in the manner suggested by Hurlburt (2009) and others in the manner outlined by Carruthers (2009b).

sufficient to issue in self-interpretation of specific thoughts. Intervals between beeps would have to be adjusted so that subjects aren't continually 'on the lookout' for them. In these circumstances one might hope to get a high number of reports of unsymbolized thinking. The background instructions would deviate from normal DES protocols, however, and could take one of two forms. In one condition subjects could be asked to note any sensory aspects at the moment of the beep, no matter how peripheral. If these are always found to be paired with reports of unsymbolized thinking in ways that are relevant to the content of the latter, then it would support Hurlburt's (2009) account. In another condition subjects could be asked to note, not just what is at the focus of their attention at the time of the beep, but to try to recall as much as they can of their experience in the immediately preceding moments. Naturally, this task might be quite demanding. But if even some subjects report relevant imagery in the moments before the occurrence of an unsymbolized thought, then it would provide some support for Carruthers' (2009b) proposal.⁹

Another sort of suggestion would be to probe the correlates of individual differences in reports of unsymbolized thinking. If Carruthers is right that self-attributions of unsymbolized thought are the result of swift and unconscious processes of self-interpretation, then one might predict that subjects who score high in unsymbolized thinking should perform better on other measures of mindreading capability, especially those requiring swift and intuitive (unreflective) forms of interpretation. Such data would need to be interpreted in the light of individual differences in rates of reporting visual imagery and inner speech, however. For while one reason why some subjects report unsymbolized thinking more often than others could be that they interpret current and recent sensory and imagistic data more swiftly and smoothly, another might be that some subjects simply have more such data to interpret.

6. Conclusion

We have sought to address two main issues. The first is the extent to which DES is a sound and reliable method for gathering data about ongoing conscious experience. While we share many of Schwitzgebel's doubts about the ability of subjects to report accurately on finely-detailed features of their experience, we also share Hurlburt's optimism about the reliability

[9] Ideally, of course, one would like to combine the two conditions into one, since the proposals made by Hurlburt and Carruthers are consistent with one another. But we fear that this would almost certainly exceed what it is possible for a normal individual to notice and describe in the space of just a few seconds.

of DES in allowing subjects to accurately apprehend broad, coarse-grained, features. The second issue builds upon the first: given that DES is a reliable source of data about at least many features of ongoing conscious experience, how can we put this method to use in answering outstanding psychological and philosophical questions? We suggest that DES embodies a useful method for clustering subjects into groups based on shared traits of inner experience (in terms of frequency and detail of particular sensory modalities, say), who can then be subjected to testing for correlations with other cognitive abilities. In addition, we are hopeful that the DES method might be adapted to help settle the question whether propositional attitudes can be directly introspected, or whether they must always be self-attributed based on a process of self-interpretation.¹⁰

References

- Baars, B. (1988) *A Cognitive Theory of Consciousness*, Cambridge: Cambridge University Press.
- Baars, B. (2002) The conscious access hypothesis: Origins and recent evidence, *Trends in Cognitive Science*, **6**, pp. 47–52.
- Baars, B., Ramsoy, T. & Laureys, S. (2003) Brain, consciousness, and the observing self, *Trends in Neurosciences*, **26**, pp. 671–675.
- Baddeley, A. (2006) *Working Memory, Thought, and Action*, Oxford: Oxford University Press.
- Baddeley, A. & Hitch, G. (1974) Working memory, in Bower, G. (ed.) *Recent Advances in Learning and Motivation*, vol. 8, Maryland Heights, MO: Academic Press.
- Baddeley, A. & Logie, R. (1999) Working memory: The multiple-component model, in Miyake, A. & Shah, P. (eds.) *Models of Working Memory*, Cambridge: Cambridge University Press.
- Bem, D. (1967) Self-perception: An alternative interpretation of cognitive dissonance phenomena, *Psychological Review*, **74**, pp. 183–200.
- Bem, D. (1972) Self-perception theory, in Berkowitz, L. (ed.) *Advances in Experimental Social Psychology*, vol. 6, Maryland Heights, MO: Academic Press.
- Briñol, P. & Petty, R. (2003) Overt head movements and persuasion: A self-validation analysis, *Journal of Personality and Social Psychology*, **84**, pp. 1123–1139.
- Carruthers, P. (2006) *The Architecture of the Mind*, Oxford: Oxford University Press.
- Carruthers, P. (2009a) An architecture for dual reasoning, in Evans, J. & Frankish, K. (eds.) *In Two Minds*, Oxford: Oxford University Press.
- Carruthers, P. (2009b) How we know our own minds: The relationship between mindreading and metacognition, *Behavioral and Brain Sciences*, **32**, pp. 121–138.
- Carruthers, P. (2010) Introspection: Divided and partly eliminated, *Philosophy and Phenomenological Research*, **79**.
- Chalmers, D. (1996) *The Conscious Mind*, Oxford: Oxford University Press.

[10] We are grateful to Russ Hurlburt and Eric Schwitzgebel for their respective comments and corrections on an earlier draft of this essay.

- Dehaene, S. & Naccache, L. (2001) Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework, *Cognition*, **79**, pp. 1–37.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J. & Sergent, C. (2006) Conscious, preconscious, and subliminal processing: A testable taxonomy, *Trends in Cognitive Sciences*, **10**, pp. 204–211.
- Dehaene, S., Sergent, C. & Changeux, J.-P. (2003) A neuronal network model linking subjective reports and objective physiological data during conscious perception, *Proceedings of the National Academy of Science*, **100**, pp. 8520–8525.
- Eagly, A. & Chaiken, S. (1993) *The Psychology of Attitudes*, San Diego, CA: Harcourt Brace Jovanovich.
- Engelbert, M. & Carruthers, P. (2009) Introspection, in Nadel, L. (ed.) *Wiley Interdisciplinary Reviews: Cognitive Science*, Hoboken, NJ: John Wiley & Sons.
- Evans, J. & Over, D. (1996) *Rationality and Reasoning*, London: Psychology Press.
- Festinger, L. (1957) *A Theory of Cognitive Dissonance*, Palo Alto, CA: Stanford University Press.
- Gazzaniga, M. (1995) Consciousness and the cerebral hemispheres, in Gazzaniga, M. (ed.) *The Cognitive Neurosciences*, Cambridge, MA: MIT Press.
- Gazzaniga, M. (2000) Cerebral specialization and inter-hemispheric communication: Does the corpus callosum enable the human condition?, *Brain*, **123**, pp. 1293–1326.
- Hurlburt, R. (2009) Unsymbolized thinking, sensory awareness, and mindreading, *Behavioral and Brain Sciences*, **32**, pp. 149–150.
- Hurlburt, R. & Akhter, S. (2008) Unsymbolized thinking, *Consciousness and Cognition*, **17**, pp. 1364–1374.
- Hurlburt, R. & Schwitzgebel, E. (2007) *Describing Inner Experience? Proponent Meets Skeptic*, Cambridge, MA: MIT Press.
- Hurlburt, R., Koch, M. & Heavey, C. (2002) Descriptive experience sampling demonstrates the connection of thinking to externally observable behavior, *Cognitive Therapy and Research*, **26**, pp. 117–134.
- Kahneman, D. (2003) A perspective on judgment and choice: Mapping bounded rationality, *American Psychologist*, **58**, pp. 697–720.
- Kahneman, D. & Frederick, S. (2002) Representativeness revisited: Attribute substitution in intuitive judgment, in Gilovich, T., Griffin, D. & Kahneman, D. (eds.) *Heuristics and Biases*, Cambridge: Cambridge University Press.
- Landman, R., Spekreijse, H. & Lamme, V. (2003) Large capacity storage of integrated objects before change blindness, *Vision Research*, **43**, pp. 149–164.
- Nisbett, R. & Wilson, T. (1977) Telling more than we can know, *Psychological Review*, **84**, pp. 231–295.
- Rey, G. (2008) (Even higher-order) intentionality without consciousness, *Review Internationale de Philosophie*, **62**, pp. 51–78.
- Schwitzgebel, E. (2002) How well do we know our own conscious experience? The case of visual imagery, *Journal of Consciousness Studies*, **9** (5–6), pp. 35–53.
- Schwitzgebel, E. (2007) Do you have constant tactile experience of your feet in your shoes? Or is experience limited to what's in attention?, *Journal of Consciousness Studies*, **14** (3), pp. 5–35.
- Schwitzgebel, E. (2008) The unreliability of naïve introspection, *Philosophical Review*, **117**, pp. 245–273.
- Shallice, T. (1988) *From Neuropsychology to Mental Structure*, Cambridge: Cambridge University Press.

- Sloman, S. (1996) The empirical case for two systems of reasoning, *Psychological Bulletin*, **119**, pp. 3–22.
- Sloman, S. (2002) Two systems of reasoning, in Gilovich, T., Griffin, D. & Kahneman, D. (eds.) *Heuristics and Biases*, Cambridge: Cambridge University Press.
- Sperling, G. (1960) The information available in brief visual presentations, *Psychological Monographs: General and Applied*, **74**, pp. 1–29.
- Stanovich, K. (1999) *Who is Rational? Studies of Individual Differences in Reasoning*, Mahwah, NJ: Lawrence Erlbaum Associates.
- Wegner, D. (2002) *The Illusion of Conscious Will*, Cambridge, MA: MIT Press.
- Wicklund, R. & Brehm, J. (1976) *Perspectives on Cognitive Dissonance*, Mahwah, NJ: Lawrence Erlbaum Associates.
- Wilson, T. (2002) *Strangers to Ourselves*, Cambridge, MA: Harvard University Press.