THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Experience and agency: Slipping the mesh

Citation for published version:

Clark, A & Kiverstein, J 2007, 'Experience and agency: Slipping the mesh', *Behavioral and Brain Sciences*, vol. 30, no. 5-6, pp. 502-503. <https://doi.org/10.1017/S0140525X07002828>

Digital Object Identifier (DOI):

[10.1017/S0140525X07002828](https://doi.org/10.1017/S0140525X07002828)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Behavioral and Brain Sciences

Publisher Rights Statement:

©Clark, A., & Kiverstein, J. (2007). Experience and agency: Slipping the mesh. Behavioral and Brain Sciences, 30(5-6), 502-503doi: 10.1017/S0140525X07002828

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

With this distinction in hand, consider Landman et al.'s conclusion:

The present data agree with the presence of two parallel types of short term memory. . . . Almost all items enter the first type of memory. It is like iconic memory, because it has a high capacity and it is maskable. . . . The second type of memory is one that resists interference by new stimuli. When new items enter the visual system, they replace the old items, except the ones that have entered the second type of representation. . . . The cue-advantage arises because the subjects selectively transfer the cued item from iconic memory to the more durable working memory. . . . (Landman et al. 2003, p.162)

Landman et al. are, then, concerned with informational persistence, not visible persistence. Their paper contains no data concerning visible persistence. Since informational persistence is consistent with no visible persistence at all, Block's appeal to Landman et al. must be somewhat indirect.

And indeed it is. Block's argument for visible persistence is based on subjects' reports: "[subjects say they are] continuing to maintain a visual representation of the whole array" (sect. 9, para. 6).

We have three points about this. First, Block needs only the weaker claim that the subjects in the Landman et al. experiment saw each rectangle as oriented horizontally or vertically, not the stronger claim that the subjects *remain* in this state after the stimulus has been replaced. The weaker claim implies Block's conclusion about inaccessibility for the same reason that the stronger one does.

We do not dispute that information about the orientation of each rectangle persists and is not as a whole accessible; we do dispute Block's claim that this inaccessible information characterizes what the subjects see. Our second point is that it is unclear that subjects' reports unequivocally support Block. Block needs subjects to agree that they saw *each* rectangle as oriented horizontally or vertically (even if they can't report which orientation each rectangle has). More precisely: for each rectangle *x*, either they saw *x* as horizontal, or they saw *x* as vertical. If the subjects merely say that they saw eight rectangles, some horizontal and some vertical, or that "they can see all or almost all the 8 to 12 items in the presented arrays" (sect. 9, para. 11), this is insufficient.

According to Landman et al., selected stimulus information is transferred from the transient iconic memory to the more durable working memory. Working memory therefore contains less information about the stimulus than iconic memory. If that is all that working memory contains, and if working memory governs subjects' reports about what they see (as Block supposes), then subjects would simply say that they saw a circle of rectangles and saw some of them as oriented horizontally/vertically. They would not, then, agree that they saw details, some of which they can't report. So our third point is this: Block must deny that the contents of working memory are simply a subset of the contents of iconic memory, which is to go beyond the results of Landman et al. If Block is right and subjects report (correctly) that they saw each rectangle as oriented horizontally or vertically, then the contents of working memory should include, not just certain information about the stimulus transferred from iconic memory, but also the meta-information that some information was *not* transferred. We are not saying that this proposal about the contents of working memory is wrong, but only that the Landman experiment does not address it.

NOTE

1. The question of the exact relationship between visible and informational persistence remains open. Loftus and Irwin (1998) argue that the many measures of visible and informational persistence pick out the same underlying process. Nevertheless, the distinction is still useful and our discussion does not rely on the assumption that it marks a real difference.

Experience and agency: Slipping the mesh

doi: 10.1017/S0140525X07002828

Andy Clark and Julian Kiverstein

School of Philosophy, Psychology and Language Sciences, George Square, Edinburgh EH8 9JX, Scotland, United Kingdom.

andy.clark@ed.ac.uk s9903600@sms.ed.ac.uk

http://www.philosophy.ed.ac.uk/staff/clark.html

Abstract: Can we really make sense of the idea (implied by Block's treatment) that there can be isolated islets of experience that are not even potentially available as fodder for a creature's conscious choices and decisions? The links between experience and the availability of information to guide conscious choice and inform reasoned action may be deeper than the considerations concerning (mere) reportability suggest.

In this elegant and tightly argued target article, Ned Block seeks to persuade us that phenomenal consciousness routinely "overflows" cognitive accessibility. By this he means that we can (and do) have experiences even in cases where we lack the kind of access that would yield some form of report that such and such an experience had occurred. The case Block makes for such an apparently hard-to-support judgment rests on a "mesh" between psychological results and work in neuroscience. The psychological data seem to show that subjects can see much more than working memory enables them to report. Thus, in the Landman et al. (2003) experiments, for instance, subjects show a capacity to identify the orientation of only four rectangles from a group of eight. Yet they typically report having seen the specific orientation of all eight rectangles. Working memory here seems to set a limit on the number of items available for conceptualization and hence report.

Work in neuroscience then suggests that unattended representations, forming parts of strong-but-still-losing clusters of activation in the back of the head, can be almost as strong as the clusters that win, are attended, and hence get to trigger the kinds of frontal activity involved in general broadcasting (broadcasting to the "global workspace"). But whereas Dehaene et al. (2006) treat the contents of such close-seconds as preconscious, because even in principle (given their de facto isolation from winning frontal coalitions) they are unreportable, Block urges us to treat them as phenomenally conscious, arguing that "the claim that they are not conscious *on the sole ground of unreportability* simply assumes metaphysical correlationism" (sect. 14, para. 9; italics in original). That is to say, it simply assumes what Block seeks to question – that is, that the kind of functional poise that grounds actual or potential report is part of what constitutes phenomenology. Contrary to this way of thinking, Block argues that by treating the just-losing coalitions as supporting phenomenally conscious (but in principle unreportable) experiences, we explain the psychological results in a way that meshes with the neuroscience.

The argument from mesh (which is a form of inference to the best explanation) thus takes as its starting point the assertion that the *only grounds* we have for treating the just-losing back-of-the-head coalitions as non-conscious is the unreportability of the putative experiences. But this strikes us as false, or at least premature. For underlying the appeal to reportability is, we suspect, a deeper and perhaps more compelling access-oriented concern. It is the concern that any putative conscious experience should be the experience *of an agent*. The thought here is that we cannot make sense of the image of free-floating experiences, of little isolated islets of experience that are not even potentially available as fodder for a creature's rational choices and considered actions. Evans (1982) rather famously rejects the very idea of such informationally isolated islands of experience. According to Evans, an informational state may underpin a conscious experience only if it (the informational state) is in some sense input to a reasoning subject. To count as a conscious experience an informational state must:

[serve] as the input to a thinking, concept-applying and reasoning system: so that the subject's thoughts, plans, and deliberations are also systematically dependent on the informational properties of the input. When there is such a link we can say that the person, rather than some part of his or her brain, receives and processes the information. (Evans 1982, p.158)

The real point here is (or should be) independent of Evan's appeal to conceptualization. What matters, rather, is that the information must be available to the agent qua "reasoning subject," where this may be unpacked in many different ways, not all of them requiring full-blown concept-use on the part of the agent (see, e.g., Bermúdez & Macpherson 1998; Hurley 1997). Evans' insight is that the notions of conscious experience and reasoned agency (here very broadly construed) are deeply intertwined: that there are non-negotiable links between what is given in conscious awareness and the enabled sweep of deliberate actions and choices available to a reasoning subject. Such a story opens up a different way of interpreting the Sperling (1960) and the Landman et al. (2003) results. In these cases (we suggest) subjects report phenomenally registering all the items because information concerning each item was, at that moment, available to be deployed in the service of deliberate, reasoned, goal-directed action. Such momentary potentiality is not undermined by the (interesting and important) fact that the selection of a few items to actually play that role then precludes the selection of the rest.

Contrariwise, Block argues that a subject like G.K. can be having an experience of a face and yet it be impossible for him to know anything of this experience. Block takes G.K.'s phenomenal experience to be constituted by recurrent processing in the fusiform face area. We believe that G.K. can be consciously experiencing a face only if this experience is at least momentarily poised for use in reasoning, planning, and the deliberate selection of types of action. Recurrent processing in the fusiform area will no doubt prove to be among the conditions necessary for realizing a state that plays this causal role.

The contents of conscious phenomenal experience, if all this is on track, must be at least potentially available for use in the planning and selection of deliberate, stored-knowledge-exploiting, and goal-reflecting and goal-responsive, actions. Block's just-losing coalitions fail to trigger winning frontal coalitions and hence fail to be in a position to contribute their contents in this manner to the full sweep of the agent's deliberate acts and choices. It is this fact (rather than the more superficial indicator of unreportability) that should motivate our treating the contents of the just-losing coalitions as non-conscious. If this is correct, then the staging post for the argument from mesh is called into question. Until the considerations concerning links between experience and rational agency are more fully addressed, it remains unclear whether the kind of "fit" to which Block appeals can really favor his conclusion over our own.

ACKNOWLEDGMENTS

Thanks to Matthew Nudds, Tillmann Vierkant, and the participants in the Edinburgh Philosophy, Psychology, and Informatics Reading Group for stimulating discussion of the target paper. This commentary was prepared thanks to support from the AHRC, under the ESF Eurocores CNCC scheme, for the CONTACT (Consciousness in Interaction) project, AH/E511139/1.

Why babies are more conscious than we are

doi: 10.1017/S0140525X0700283X

Alison Gopnik

Department of Psychology, University of California at Berkeley, Berkeley, CA 94720.

gopnik@berkeley.edu

ihd.berkeley.edu/gopnik.htm

Abstract: Block argues for a method and a substantive thesis – that consciousness overflows accessibility. The method can help answer the question of what it is like to be a baby. Substantively, infant consciousness may be accessible in some ways but not others. But development itself can also add important methodological tools and substantive insights to the study of consciousness.

Infants and young children cannot report their phenomenology. This has led some philosophers to argue that babies' consciousness must be limited. Even if we think that babies are conscious it might seem impossible to recover the particular character of their experience. A version of Block's abductive method can at least partially solve this problem. We can consider a wide range of functional and neural correlates of conscious experience in adults and then look to see similarities and differences in babies.

In adults attention is highly correlated with vivid consciousness. This consciousness has a particular subjective quality – "the spotlight" – with a defined – "brightly lit" – focus and surrounding darkness. Both phenomenologically and functionally, attention to one event seems to inhibit consciousness of other events – as in inattention blindness.

There is a distinction between exogenous and endogenous attention. Exogenous attention is driven by information-rich external events. These events may be intrinsically salient. But exogenous attention may also be driven by subtle unexpectedness. Exogenous attention is marked by characteristic event-related potential (ERP) signatures, eye movements, decelerating heart rate, and parietal activation. Exogenous attention and vivid consciousness characteristically fade as information is obtained, a process of habituation.

Endogenous attention is the sort of top-down "paying attention" that is motivated by specific goals rather than by intrinsic interest. It has been the focus of the adult literature for methodological reasons, and endogenous and exogenous attention are often not distinguished in discussions of consciousness. Frontal activation seems to be particularly important for endogenous attention.

These functional features of attention correlate well with neurological patterns. In adult animals, endogenous attention leads to the release of cholinergic transmitters to some parts of the brain and inhibitory transmitters to other parts. Attention increases both the efficiency of a particular part of the brain and its plasticity; and it inhibits activation and plasticity in other brain areas.

So for adults there is an elegant if undoubtedly oversimplified story about how consciousness works. In the canonical case, the goal-directed control systems in frontal cortex indicate that a particular kind of event is important. The perceptual system, guided by attention, zooms in on just that event and the brain extracts information about the event and modifies itself, that is, learns accordingly. Significantly, though, this whole process is highly focused; other parts of the brain may actually be shut down in the process. Vivid spotlight consciousness is the phenomenological result.

Even very young infants have extensive exogenous attention capacities. When they are presented with even highly subtle and conceptually unexpected novel events, they immediately focus their gaze on these events, and show similar heart rate deceleration and ERP signatures to those of adults. Indeed, this is the basis for the habituation technique that is our principal source of information about infant's minds. However, infants develop endogenous attention much later, and it is still developing during the preschool years. Moreover, and probably correlated with this fact, infants and young children appear to have less focused attention than older children – for example, they show better incidental memory.

The neurology suggests a similar picture. The parietal and sensory systems involved in exogenous attention are on line at an early age. The top-down frontal regions and connections that control endogenous attention only mature later. Young animals' brains are far more plastic than adult brains and this plasticity is much less focused and attention-dependent. Cholinergic transmitters are in place early, while inhibitory transmitters emerge only later.

So again, an undoubtedly oversimplified but suggestive picture emerges. When infants and young children process information