Commentary on **Ned Block (1995). On a confusion about a function of consciousness. BBS 18:227–287.**

Abstract of the original article: Consciousness is a mongrel concept: there are a number of very different "consciousnesses." Phenomenal consciousness is experience; the phenomenally conscious aspect of a state is what it is like to be in that state. The mark of access-consciousness, by contrast, is availability for use in reasoning and rationally guiding speech and action. These concepts are often partly or totally conflated, with bad results. This target article uses as an example a form of reasoning about a function of "consciousness" based on the phenomenon of blindsight. Some information about stimuli in the blind field is represented in the brains of blindsight patients, as shown by their correct "guesses." They cannot harness this information in the service of action, however, and this is said to show that a function of phenomenal consciousness is somehow to enable information represented in the brain to guide action. But stimuli in the blind field are *both* access-unconscious and phenomenally unconscious. The fallacy is: an obvious function of the machinery of access-consciousness is illicitly transferred to phenomenal consciousness.

An example of access-consciousness without phenomenal consciousness?

Joseph E. Bogen

Department of Neurological Surgery, University of Southern California, Los Angeles, CA 90033-4620.

Abstract: Both Block and the commentators who accepted his P versus A distinction readily recognize examples of P without A but not vice versa. As an example of A without P, Block hypothesized a "zombie," computationally like a human but without subjectivity. This would appear to describe the disconnected right hemisphere of the split-brain subject, unless one alternatively opts for two parallel mechanisms for P?

Block (1995a) makes a clear conceptual distinction between what he calls phenomenal consciousness (P) and what he calls access consciousness (A). The former (P) he points to by saying that "P-conscious states are experiential"; he gives examples such as smells, tastes, pains, thoughts, and desires (p. 230). The latter (A) he describes as a state in which some content is "poised for use as a premise in reasoning" and "poised for rational control of action" (p. 231). A can also be "poised for rational control of speech," but for Block this is not a necessary aspect because he considers chimps to have A. Indeed, he elsewhere notes that "very much lower animals are A-conscious" (p. 238).

Block is at some pains to consider the possibilities of P without A and of A without P; in particular, he says, "It certainly seems *conceptually possible* that the neural bases of P-consciousness systems and A-consciousness systems are distinct" (p. 233). Block provides some possible examples of P without A (on p. 234 and p. 244) such as "becoming conscious" (acquiring A) of an ongoing noise (e.g., a pneumatic drill) some considerable time after one has been "aware of" or has been "experiencing" it. Although Block is reluctant to accept dreaming as an example of P without A (p. 275), some of us are inclined to agree with Revonsuo (1995) that dreaming can be a good example of subjective experience in the absence of both current perceptual input and behavioral output (see also Delacour 1995; Paré & Llinás 1995).

Block suggests a few hypothetical examples of A without P, such as a "zombie" that is computationally identical to a person but without any subjectivity. He concludes "I don't know whether there are any actual cases of A-consciousness without P-consciousness, but I hope I have illustrated their conceptual possibility" (p. 233).

If there can be A without P as well as P without A, we should probably conclude that they have distinct neural bases. However, if there can be P without A but there cannot be A without P (that is, A entails P), then there could be one neural mechanism which in the case of P without A is temporarily disconnected either from action or from ideation or both.

In my recent proposal (Bogen 1995a; 1995b) that the intralaminar nuclei (ILN) of a thalamus provide a cerebral hemisphere with *both* subjectivity and access to action and/or thought, it was explicitly assumed that a single mechanism provides both P and A as well as providing, on some occasions, only P. This assumption was criticized by Kinsbourne (1995). Using a distinction similar to that proposed by Block, but in a more anatomico-physiologic context, Kinsbourne argues that the ILN can be "attention-action coordinators" without also being "subjectivity pumps" (p. 168). At one point he suggests that without coordination of attention and action "consciousness would lapse"; that is, there would be no P without A. His main emphasis is on the possibility that A is provided by a different neural basis than P, in which case there would be the possibility of A without P. Kinsbourne does not, however, provide examples of A without P. At this point it seems that we are left with a problem: Are there actual cases of A without P? As Block (1995r) put it (p. 277), "The relative ease of finding cases of P without A as compared with A without P suggests the distinction is on to something to do with the joints of nature." He added, "If brain damage does not yield cases of A without P, this is an especially interesting fact given the fantastic wealth of variation in brain-damage cases."

Speaking of brain-damage cases, I would ask, what about split-brain humans, with whom I have had a lengthy acquaintance (Bogen 1993; Bogen & Vogel 1962)? So far as I am aware, no one has ever denied that: (1) in most of these patients speech is produced only by the left hemisphere, (2) the speech is evidence that P and A coexist in that hemisphere, and (3) verbal denial of information that has been delivered only to the right hemisphere (and rationally acted upon) reflects the existence of an independent capacity in the right hemisphere, that is, an A-consciousness different from the A-consciousness of the left hemisphere. Does the right hemisphere in that circumstance also possess its own P-consciousness? (In my scheme, this P is provided by the ILN of the right hemisphere.) Most of us with a personal experience with split-brain patients (e.g., Sperry 1974; Zaidel 1978; Zaidel et al. 1996) believe that the disconnected right hemisphere also has its own P-consciousness. The same conclusion has been recently suggested by others (Berlucchi et al. 1995; Corballis 1995). If we are wrong, and the right hemisphere of a split-brain patient does *not* have a separate P in spite of having a distinctly different A, then perhaps we have here a readily replicable example of A-consciousness without P-consciousness.

Consciousness by the lights of logic and commonsense

Selmer Bringsjord

Department of Philosophy, Psychology and Cognitive Science, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180.
selmer@rpi.edu; www.rpi.edu/~brings

Abstract: I urge return by the lights of logic and commonsense to a dialectical *tabula rasa* – according to which: (1) consciousness, in the ordinary pre-analytic sense of the term, is identified with P-consciousness, and "A-consciousness" is supplanted by suitably configured terms from its Blockian definition; (2) the supposedly fallacious Searlean argument for the view that a function of P-consciousness is to allow flexible and creative cognition is enthymematic and, when charitably specified, quite formidable.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Block's (1995t) paper, according to Warren (1995), "adds its own confusion to [the] difficult and dismal topic [of consciousness] (p. 270)." Warren proceeds to claim that the terms at the heart of the consciousness dialectic are obscure and to prescribe, therefore, that the topic should simply no longer be an object of scientific scrutiny. While Warren's view is easily refuted (if the view were correct, then given the obscurity of "proof" that reigned from Euclid to Frege, following Warren's prescription would have eventuated in a world without classical mathematics!), his attitude, I submit, *is* seductive – because, let's face it, the commentary to this point certainly at least *appears* to be a dark cacophony, with consensus, or even near-consensus, nowhere to be sensed, let alone seen. The antidote to Warren's despair is to return by the lights of logic and commonsense to a dialectical *tabula rasa* – according to which: (1) consciousness, in the ordinary preanalytic sense of the term, is identified with P-consciousness, and "A-consciousness" is supplanted by suitably configured terms from its Blockian definition; (2) the supposedly fallacious Searlean argument for the view that a function of P-consciousness is to allow flexible and creative cognition is enthymematic and, when charitably specified, quite formidable. Let's start with (1).

Consider, first, Georges Rey's (1995) *reductio ad absurdum*, summarized (and affirmed) by Block (p. 235): "Ordinary laptop computers are capable of various types of self-scanning, but . . . no one would think of their laptop computer as 'conscious' (using the term in the ordinary way)." The target of this *reductio* is the identification of consciousness *simpliciter* with internal scanning (or "monitoring consciousness," as Block calls it). But Rey's argument is of course easily adapted so as to threaten Block's fundamental distinction between P- and A-consciousness: As Lloyd (1995) points out, Block's (para. 3 of sect. 4.2.2) definition of A-consciousness[1] is satisfied (provably, I might add[2]) by "any implemented computer program" (p. 262), but no one would think of a pascal program written by a nine-year-old child and running on a humble laptop as "conscious" in any (nonweird) construal of the term.

Block is the last person on the planet equipped to dodge this argument. It was none other than Block (1980) who long ago told us that any theory of consciousness entailing that, say, a bunch of beer cans and string can be conscious is a worthless theory. But it has been known for decades that a suitably configured abacus can compute all functions a computer can compute (Lambek 1961), and beer cans and string can obviously be used to instantiate an abacus. It follows immediately from this and Lloyd's argument that an "abacused" bunch of beer cans and string, on Block's view, is conscious. And, ironically enough, it was Searle (1983) who gave us the beer can *gedankenexperiment*.

In his response to first-round commentaries, Block (1995r) registers (R3, para. 1) his observation that not only Lloyd (1995), but also Graham (1995), Natsoulas (1995), Revonsuo (1995), and the Editorial Commentary (1995) hold that A-consciousness isn't consciousness at all. Unfortunately, Block then (R3, para. 1) reminds readers that it is a non sequitur to infer from the proposition that zombies (in the technical sense of "zombie" at play in these discussions) are not conscious in any sense to the idea that A-consciousness is not a form of consciousness. As a shield against the "funny instantiation" objection, of course, this is itself fallacious reasoning: It commits the "straw man" fallacy.

So Block's attempt to disarm counterarguments that A-consciousness isn't consciousness is at best half-hearted, but what of his attempt to establish in his response that A *is* a kind of consciousness? In order to evaluate this attempt, I think it may be useful to consider (e.g.) the "average professor." About this creature we might say

(3) The average professor owns five sportcoats.

But we shouldn't infer from (3) that there really is, out there in the world, a professor picked out by (3) who owns five sportcoats. Sentence (3) is elliptical for something like

(3′) When you add together, for each professor (and there are *n* of them), the number of sportcoats owned, and then divide by *n*, the result is 5.

Is there really, out there in the world, this thing Block calls "A-consciousness"? I don't think so. In fact, all his talk of this vaporous concept is easily translated away in the manner of (3) to (3′). Consider, for example, Block's assertion (p. 275) about prosopagnosia:

(4) A prosopagnosiac "lacks A-consciousness of the information about the identity of the person."

Sentence (4), courtesy of the definition of A-consciousness Block provides, is elliptical for something like

(4′) A prosopagnosiac is afflicted by certain failures in the processing of information involved in representing and reasoning about faces.

Who needs A-consciousness? Without it, and with some scientific toil, (4′) promises to lead the way to a complete specification (perhaps even a mathematization) of the information-processing failures at the heart of prosopagnosia. In sum, let's use "conscious" to refer to P-consciousness. And let's replace the confusing "A-consciousness" with appropriately configured terms from its Blockian definition. Then we can comfortably say that beer cans and string, calculators, and library reference systems aren't conscious – but *are* marvels of information processing. And, armed with this commonsensical view, we needn't stop there: Nearly every confusion Block seeks to disentangle, and every syndrome he seeks to analyze, can be handled quite nicely on our streamlined taxonomy.[3]

Now to the second part of my logico-commonsense view, namely, that Searle's argument for the view that a function of P-consciousness is to permit flexibility and creativity is fundamentally correct. (Despite the fact that Block tells us both that consideration of this argument is the dénouement of his paper [p. 239], and that his paper, overall, is "primarily concerned with reasoning, not with data" [p. 227], no commentator in the first round gave [printed] thought to the details of Searle's argument.) What is the argument? Block says (p. 240): "Searle argues: P-consciousness is missing; so is creativity; therefore the former lack explains the latter lack." If this *is* Searle's argument, we don't need Block's attack in the least, for this reasoning is *transparently* fallacious, as can be revealed by effortless parodies. (For example, my watch is missing [because I forgot that I left it outside in my car]; so is my car's battery [because a thief stole it]; therefore my watch's absence explains the battery's.) In fact, even the Blockian (see his note 25)[4] argument-schema of which Searle's reasoning is an instance is untenable:

(5) Person *S* loses *x* from time *t* to *t′*.
(6) *S* loses the ability to ϕ from *t* to *t′*.

Therefore:

(7) A function of *x* is to facilitate ϕing.

No logic certifies this argument-schema; once again, parodies abound. For example, suppose Jones coincidentally loses his watch over the same period of time he's afflicted by prosopagnosia. Then by the aforementioned reasoning:

(5′) Jones loses his watch from Monday to Friday.
(6′) Jones loses the ability to recognize faces from Monday to Friday.

Therefore:

(7′) A function of Jones' watch is to facilitate face recognition.

What Block is ignoring, or at least what charity in these matters dictates he consider, is a principle that unpacks the commonsense idea that if the advent of a psychological deficiency coincides with

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

a noteworthy diminution of a person's faculty, then it's a good bet that the diminution is causally linked with the deficiency. Adding (a slightly more sophisticated version of) this principle, as well as a premise that is its antecedent, to (5) and (6), turns the schema from a stunning non sequitur to a formally valid form (in first-order logic) for (7). The principle and premise, respectively, are:

(P1) If *S* loses *x* over an interval of time during which *S* loses the ability to ϕ, and there is at least a *prima facie* reason to think *x* is centrally employed when people ϕ (in part because attempts to replicate ϕing in systems lacking *x* have failed, and show no appreciable promise of succeeding in the future), then a function of *x* is to at least facilitate ϕing.

(8) *S* loses *x* over an interval . . . promise of succeeding.

Victorious instantiations of this schema seem to me to be at hand. (If *x* = "P-consciousness," and ϕ = "write belletristic fiction," then it turns out that I have elsewhere explicitly defended the relevant instantiation.[5] This defense capitalizes on [P1]'s parenthetical by including an observation that AI [Artificial Intelligence] has so far failed to produce creative computer systems.) Block would disagree. He would insist that Searle (in the Penfieldian instantiation Block intimates) gives us no reason to think that the appropriate instantiation of (5) – say, Smith loses P-consciousness during a *petit mal* seizure that overlaps his driving a car – is true. Unfortunately, this complaint is uncompelling, for at least[6] two reasons.

First, there is at least some evidence for the proposition that "normals" lose P-consciousness of, and P-consciousness arising from, behaviors that become routinized (e.g., Cooper & Shepard (1973) and Pani (1982) appear to show that when subjects become skilled enough to render imagistic tasks "automatic," they lose P-consciousness of these tasks). Given Block's undeniable mastery of the literature, I find it peculiar that such work isn't discussed in his paper.

The second reason to regard Block's objection to Searle's (5) as uncompelling is more powerful, certainly more ironic. It is that when evaluating this premise Block seems to forget his own *modus operandi:* introspection. Appeals to introspection are ubiquitous in Block's paper and in his response, but one such appeal, for autobiographical reasons, caught my attention above most others: Block's report (p. 275) – designed to counter Revonsuo's (1995) proposal that dreams constitute cases of P-consciousness without A-consciousness – that Chomsky engages in rational dreaming. If this constitutes germane evidence, then what about the ability to premeditatedly bring on what might be called quasi-*petit mal* episodes? Suppose, for example, that Brown decides, before reading to his children, that while he reads he is going to spend time reflecting upon some difficult and long-standing problem, the solution to which calls for some creativity. Brown is willing to do the reading, at least significant stretches of it, as an "automaton," while he directs his consciousness (P-consciousness in Block's scheme) toward a problem unsolvable by any familiar routine or algorithm. Is there anything it is like for Brown to read in such a scenario? Since I often do what Brown does, I can inform Block that the answer is "No." I have absolutely no memories about what I read; I have no subjective awareness of the words, sentences, themes, characters, no P-consciousness of anything related to what I have read. And yet I justifiably infer from the absence of protests from my listeners that I have performed adequately.

All of us, I venture, have experienced unplanned intervals of "automatism." To repeat the familiar example, you're driving late at night on the interstate; you're 27 miles from your exit . . . and the next thing you know, after reverie about a research problem snaps to an end, you are but *17* miles from your turnoff. Now, was there anything it was like to drive those 10 mysterious miles? If you're like me, the answer is a rather firm "No" (and I daresay the real-life cases are myriad, and not always automotive). Now, why is it that such episodes invariably happen when the ongoing overt behavior is highly routinized? Have you ever had such an episode while your overt behavior involved, say, the writing of a short story, or the proving of a theorem? These are rhetorical questions only, of course. But the point is that Block is dead wrong that there is *no reason* to think that there is nothing it's like for an epileptic driver to turn through a curve on the interstate (pp. 239–40). I conclude, then, that (5) – appropriately instantiated – is plausible; so the deductively valid reconstruction from {(5), (6), (8), (P1)} to (7) – once it too is accordingly instantiated – constitutes a formidable case for the view that a function of P-consciousness is to facilitate creative cognition.

If Searle is to emerge entirely unscathed, there is a loose end. What of Block's claim (p. 241) that Searle contradicts himself when he says both that the totally unconscious epileptic can drive home and that the car would crash if the epileptic were totally unconscious? What Searle should be read as saying here is that (1) as *a matter of contingent fact*, if the epileptic were totally unconscious his car would crash, but (2) it's nonetheless also true that it is *possible* for a driver who is totally non-P-conscious to drive home without crashing. Indeed, I fully expect successful automatic driving systems to be in place alongside humans before long; and I expect these systems to be, at heart, nothing fancier than present-day computation – computation devoid of P-consciousness. For that matter, the near-future surely also holds simple computational systems capable of reading aloud to my offspring. These systems, so exotic now, will soon become as unassuming as TVs, which are just as bereft of consciousness as beer cans.

NOTES

1. Block's definition: Something is A-conscious if it has (an) A-conscious state(s). Such a state must (p. 231) be poised (1) to be used as a premise in reasoning, (2) for rational control of action, and (3) for rational control of speech.

2. For example, the database application currently running on my laptop satisfies Block's (1)–(3) for the following reasons. One, the application is based on first-order logic, so a state of the system is nothing but a set of first-order formulas used as premises in deductive reasoning. Two, action is controlled by rational deduction from such sets. Three, "speech" is controlled by rational deduction from such sets with help from grammars designed to enable simple conversation. The application "talks" by producing text, but it could be outfitted with a voice synthesizer, and at any rate Block tells us (p. 231) that condition (3) isn't necessary, since nonlinguistic creatures can be A-conscious in virtue of their states satisfying (1) and (2). Along the same lines, it's probably worth noting that action (indicated in condition (2)) can be slight, since many paralyzed people are not only A-conscious, but P-conscious as well.

3. For example, Dennett's views, shown to be at best bizarre by Block (pp. 237–39), look that way in the commonsense view. In fact, Block's devastating critique in paragraphs 12–14, section 5, remains nearly word-for-word intact in the commonsense view.

4. I charitably neaten the schema just a bit, in part to ease analysis carried out below.

5. To put a bit of this reasoning barbarically (for lack of space), I have argued that without P-consciousness an "author" cannot adopt the point of view of his character(s), and that so adopting is a *sine qua non* for producing belletristic fiction. (See Bringsjord 1992; 1995.)

6. I say "at least" here because there may be a reason to reject Block's objection out of hand: Block's complaint has nothing whatever to do with fallaciousness. An argument is fallacious when and only when its conclusion doesn't follow from its premises. As, for example, in the famous fallacy of affirming the consequent: $p \rightarrow q$; q; therefore p. See also Harmon's (1995) commentary for sample fallacies.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Two conceptions of access-consciousness

Derek Browne

Philosophy Department, University of Canterbury, Christchurch, New Zealand. **d.browne@phil.canterbury.ac.nz**

Abstract: Block's (1995) cognitive conception of consciousness might be introduced in the service of two different projects. In one, the explanatory gap between science and folklore remains. In the other, a reductive claim is advanced, but the intuitive idea of consciousness is abandoned.

If cognitivism is true, then consciousness, if real, is really cognition. Cognitivists will reject Block's (1995t) claim that there is a distinct kind of phenomenal-consciousness, and so will reject his charge that they are guilty of conflating different properties. But many cognitivists will allow that the idea of access-consciousness is on the right track. Rather than adding to the ranks of those who query Block's defense of phenomenal-consciousness, I want instead to ask some questions about his handling of the cognitive conception. Access-consciousness rests on the idea of "poised content," that is, of semantically evaluable states that are poised for the rational control of reasoning, speech, and action. My discussion is organized around the following question: Does Block intend to introduce the concept of access-consciousness in the context of phenomenology or of cognitive theory?

Suppose the former. Then Block is saying that it is an introspectively accessible fact about our mental lives that we have thoughts. The thoughts of which we are conscious (to which we have *access*) enter our reasoning as premises and may contribute ("rationally") to the control of speech or action. The phenomenology of mind needs a concept that captures our consciousness of thought. The distinctive marker of the phenomenological conception of access-consciousness is that the subject of the access relation is the person himself. Access-consciousness, under this interpretation, picks out the set of contentful mental states in me (my "thoughts") to which I have personal-level access. This is by contrast with the hypothesized multitude of cognitive states to which my subpersonal agencies have access but to which I do not. To save the phenomena, it seems that we should say that each of us has a kind and quality of access to our own thoughts that no other person has. Notice, however, that it is, in Block's view, equally a phenomenological fact that we have immediate access to inner states that have noncognitive, experiential properties ("feelings"). Feelings are the objects of P-consciousness. If the subject of the relation of access in Block's A-consciousness is the self, then the term "access" is ill-chosen, for we equally have access (he thinks) to something other than thoughts, namely, feelings.

As that may be, and persisting for a time with this interpretation, Block does more than describe the phenomenology. He offers a hypothesis about a common feature of the set of cognitive episodes to which we have conscious access. They are those states that are (1) poised for use as premises in reasoning, (2) poised for rational control of action, and (3) poised for rational control of speech (Block 1995t, p. 231). Block's empirical hypothesis – for this is surely not "conceptual analysis" of any remotely a priori kind – is that the thoughts that are phenomenologically identified as being those of which I am conscious have a common, functional property: They are all poised for the rational control of reasoning, speech, or action. At first blush, this is a nice idea if one is both a friend of consciousness and a materialist, for it suggests a function for consciousness. The function of consciousness is to exert active, executive control, that is, to direct content into the rational control of reasoning, speech, and action. But the idea doesn't work. The concept of poised content itself has an obvious utility if one is trying to understand how the rational control of action emerges out of high-level, information-processing activities. But consciousness is no help here. The problem is to close the explanatory gap between the intuitive conception of consciousness (A-consciousness) yielded by the phenomenology and these cognitive functions. How does the fact of being accessible at the personal level *explain* the availability of content for the rational control of reasoning, speech, and action? Unless the conscious, executive self is already credited with special powers, the fact that a bit of content is present to consciousness doesn't explain how it acquires any powers that it would not otherwise have. It is unhelpful to say that consciousness can explain the mind's ability to organize its knowledge for the rational control of reasoning, speech, and action. The notion of conscious agency has no independent explanatory power. One hopes instead that the idea of poised content will help to explain the troubled phenomenon of consciousness.

Block says that his three conditions are together sufficient for access-consciousness. The sufficiency claim means that all poised content is conscious. Under the phenomenological interpretation of Block's project, the claim is that there is a uniform correlation, presumably nomological, between two "logically" distinct properties: the property of being, intuitively speaking, a conscious thought, and the property of being a state with poised content. This is worth thinking about, provided we don't succumb to the idea that by citing the former property (consciousness) we have the beginnings of a lawful explanation of the latter (poised content). But the sufficiency claim, the claim that all poised content is conscious, supports a very different interpretation of Block's project: the second alternative I mentioned at the beginning. Under this interpretation, there are not two distinct properties here, just one. Suppose (not very plausibly) that Block is not especially concerned to save the phenomena: He does not begin by accepting as given the phenomenology of consciousness and demanding that cognitive science preserve what we pretheoretically know about the conscious aspects of mind. He is instead advancing what will turn out to be a reductive hypothesis, one that saves what can be saved in the folk conception of consciousness and discards the rest. Suppose a complex behavioural control system contains some contentful states that have the following functional property: They are poised for the rational control of reasoning, speech, and action. The key idea in this line of thought is that just to be a state with poised content is to be a conscious state. Instead of a nomological connection between (conceptually) distinct properties, Block would be making a constitutive claim: The property a thought has of being conscious is identical with the property of having poised content. We are not trying to save the phenomena at all costs, so we are not committed to attributing to consciousness any properties that our cognitive theory does not recognise. Any behavioural control system of sufficient complexity to support the functions played by poised content would be conscious in the cognitive sense.

One of the tell-tale marks of this interpretation of access-consciousness is that the access relation is not a relation between the self and its thoughts. Stich's idea of "inferential promiscuity" (Stich 1978), borrowed by Block, hints that the terms of the access relation are the different thoughts that join together in inferential liaisons: They have access to each other. Perhaps a preferable reading is that the processors controlling reasoning have common access to all those contents, as have the processors controlling speech and action.

The reductionist view alarms the friends of phenomenology, however, because it replaces the first-personal character of consciousness with impersonal, computational relations that are realised in information-processing systems. Perhaps this will not alarm Block, because he also has his concept of phenomenal-consciousness, which is supposed to capture the subjective features of experience. Perhaps there is "nothing it is like to be" a conscious thinker. But if that is so, to describe a (sufficiently complex) cognitive system as "access-conscious" adds nothing to the leaner information-processing description of a system as containing "poised content."

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Availability: The cognitive basis of experience

David J. Chalmers

Department of Philosophy, University of California, Santa Cruz, CA 95064.
chalmers@paradox.ucsc.edu

Abstract: Although A-consciousness and P-consciousness are conceptually distinct, a refined notion of A-consciousness makes it plausible that the two are empirically inseparable. I suggest that the notion of direct availability for global control can play a central role here, and draw out some consequences.

Block's (1995) distinction between access consciousness and phenomenal consciousness (or experience) is very useful. There is clearly a *conceptual* distinction here, as illustrated by the fact that: (1) one can imagine access without experience and vice versa; (2) access can be observed straightforwardly, whereas experience cannot; and, most important, (3) access consciousness seems clearly amenable to cognitive explanation, whereas phenomenal consciousness is quite perplexing in this regard. But the tight *empirical* link between the two phenomena deserves attention.

Bringing access and experience closer together. Block himself notes that P-consciousness and A-consciousness often occur together. This is no accident, as one can see by noting that a P-conscious experience is usually reportable, and that reportability implies accessibility of the corresponding information. Block does not think they always occur together, but I think that with appropriate modifications they might. One of the most interesting projects in this area is that of *modifying* the concept of A-consciousness in such a way as to make it plausible that A-consciousness (in the modified sense) and P-consciousness are perfect correlates.

A good start is the modified notion of *direct availability for global control*. That is, a content is A-conscious in the modified sense when it is directly available for use in directing a wide range of behaviors, especially deliberate behaviors. I am not sure how different this is from Block's definition: it plays down the role of rationality and reasoning (after all, impairments of rationality probably do not diminish phenomenal consciousness), it relegates verbal report to the status of a heuristic (as Block himself suggests), and there is another important difference that I will come to shortly. The restriction to *direct* availability works to eliminate contents that can be retrieved with some work but that are not conscious.

To see how well this modified notion of A-consciousness correlates with P-consciousness, we need to see how it handles Block's examples in which one sort of consciousness occurs without the other. Block's examples of A-consciousness without P-consciousness are all mere conceptual possibilities (zombies and super-blindsight, for example), so they are not relevant here, but to illustrate P-consciousness with A-consciousness he gives some real-world examples. One is Sperling's (1960) example in which all nine letters in a square array are experienced, but only three can be reported at a time. In this case, only three letter-representations are *accessed*, but it is nevertheless plausible that each of the nine was *available*, until the process of access destroyed their availability. This works because the modified notion of A-consciousness is dispositional – not access, but *accessibility* is required. And it is plausible that all nine letter-representations are A-conscious in the modified sense. So even in this case, P-consciousness and modified A-consciousness occur together.

The case of the drilling noise in the background can be handled similarly. Here it seems reasonable to say that the information was directly available all along; it simply wasn't accessed. The case of experience under anesthesia (if it is actual) is trickier, but we might handle it by saying that in these cases the corresponding contents are *available* for global control; it is just that the control mechanisms themselves are mostly shut down. We might say that the information makes it to a location where it could have been used to direct behavior, had the motor cortex and other processes been functioning normally.

Other cases could be considered and further refinements could be made. A fuller account might flesh out the kind of availability required (perhaps a kind of high-bandwidth availability is required for experience, or at least for experience of any intensity) and might specify the relevant kind of control role more fully. Counter-examples are not threatening but helpful; they allow us to refine the definition further. The details can be left aside here; the point is that this project will lead to a functionally characterized property that might correlate perfectly with P-consciousness, at least in the cases with which we are familiar.

This property – something in the vicinity of direct availability for global control – could then be thought of as the information-processing correlate of P-consciousness, or as the cognitive basis of experience. There are some interesting consequences for the issues that Block discusses.

Empirical work on consciousness. Block notes that researchers on consciousness often start with an invocation of phenomenal consciousness but end up offering an explanation of A-consciousness and leaving P-consciousness to the side. The tight link between the two suggests that a somewhat more charitable interpretation is possible. If experience correlates with availability for global control, much of this work can be interpreted as seeking to *explain* A-consciousness, but trying to find a *basis* for P-consciousness. For example, Crick and Koch's (1990) oscillations are put forward because of a potential role in binding and working memory; that is, in integrating contents and making them available for control (working memory is itself an availability system, after all). If both the empirical hypothesis (oscillations subserve availability) and the bridging principle (availability goes along with experience) are correct, then the oscillations are a neural correlate of experience, which is just what Crick and Koch claim.

The same holds elsewhere. Shallice's "selector inputs" for "action systems" (1972) and his "supervisory system" (1988a; 1988b) are clearly supposed to play a central role in availability and control; if the empirical hypothesis is correct, they could reasonably be regarded as part of the basis for conscious experience. Similarly, the "global workspace" of Baars (1988), the "high-quality representations" of Farah (1994), the "temporally-extended neural activity" of Libet (1993), and many other proposals can be all be seen as offering mechanisms in the process whereby some contents are made available for global control. The common element is striking. Of course, it is an empirical question which of these proposals is correct (although more than one might be, if they offer accounts of different parts of the process or descriptions at different levels). But insofar as these mechanisms play a role in the availability/control process, they are candidates to be neural or cognitive correlates of experience, which is often what the authors suggest (correlation is all that Farah and Libet claim; Shallice and Baars oscillate between "correspondence" and explanation).

The picture is this: (1) we know that availability goes along with experience; (2) we discover empirically that some mechanism plays the central role in the availability process. We may then conclude that the mechanism is part of the explanation of A-consciousness and part of the basis of P-consciousness. Of course, the story about the mechanism alone does not explain P-consciousness, as we still have not explained why availability always goes along with experience; we have simply taken for granted that it does. But if we are prepared to take the link between availability and experience as a kind of background assumption (perhaps for later explanation), this can provide a useful partial explanation of the contents of experience.

A phenomenal consciousness module? Interestingly, this analysis allows us to make some sense of the idea of a phenomenal consciousness module. *If* it turns out that there is a single system responsible for mediating the availability of certain contents for global control – something like Baars's global workspace or Shallice's supervisory system – then it might be plausible that the contents of that system correspond precisely to the contents of experience, and maybe we could call it a P-consciousness module. I do not think it is probable that there is such a module – more

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

likely there are many different mechanisms by which contents become available for a control role – but at least the idea makes sense. But the *only* way there could be a "P-consciousness" module would be for it to be an availability/control module. If a module were dissociable from the relevant role in availability and control, the considerations above suggest that it would be dissociable from P-consciousness too.

In particular, there is something very strange about the idea of an "epiphenomenal" P-consciousness module (Block's Fig. 3). The main motivation for epiphenomenalism is surely that experience seems superfluous to any information-processing; but Block's idea suggests an implausible epiphenomenalism *within* the information-processing story. Indeed, if the module has no effect on other processes, then we could lesion it with no external change (same reports, even), and no empirical evidence could support the hypothesis. Perhaps Block means to allow that the module has the very limited function of causing phenomenal reports, so that lesioning it eliminates remarks such as "I am having a blue sensation." But now either (1) remarks such as "There is a blue object," confident blue-directed behavior, and so on are all eliminated too – in which case the module had an important function after all – or (2) they are preserved (a kind of ultra-superblindsight), implying an extraordinary independence between the pathways responsible for phenomenal report and those responsible for visual descriptions and normal visual processing. Given the remarkable coherence between visual descriptions and reports of visual experience, one presumes that they are tied more closely than this.

The function of consciousness? The link between P-consciousness and (modified) A-consciousness makes the search for a function for P-consciousness even more hopeless. Given the correlation, any purported function for P-consciousness can be attributed to A-consciousness instead.

Only those who implausibly identify the *concept* of P-consciousness with that of (modified) A-consciousness have a way out. If one accepts the conceptual distinction, one will accept the conceivability of zombie functional isomorphs (made of silicon, say). To be consistent, one must then accept the conceivability of zombie *physical* isomorphs, as there is no more of a conceptual entailment from neural stuff to consciousness than there is from silicon stuff. From here, it is easy to see that P-consciousness gives me no functional advantage. After all, I am different from my zombie twin *only* in that I have P-consciousness and he does not, but we are functionally identical.

Block suggests that P-consciousness might "grease the wheels" of A-consciousness, but this cannot work. P-consciousness is redundant to the explanation of the physical mechanisms of A-consciousness, as the conceivability of the zombie shows: same physical mechanisms, same explanation of A-consciousness, no P-consciousness. The remaining option is to "identify" P-consciousness with modified A-consciousness (empirically but not conceptually), solving the problem by fiat. I think this sort of identification without explanation misunderstands the way that scientific identification works (see Chalmers 1996), but in any case it still leaves the concept of P-consciousness with no *explanatory* role in cognitive functioning. The independent concept of A-consciousness does all the work. I think it best to accept instead that phenomenal consciousness is distinct from any physical or functional property, and that it does not need to have a function to be central to our mental lives.

P-Consciousness presentation/ A-Consciousness representation

Denise Gamble

Department of Philosophy, The University of Adelaide, Adelaide 5005, South Australia. **dgamble@arts.adelaide.edu.au**

Abstract: P-Consciousness (P) is to be understood in terms of an immediate fluctuating continuum that is a presentation of raw experiential matter against which A-consciousness (A) acts to objectify, impose form or make determinate "thinkable" contents. A representationalises P but P is not itself representational, at least in terms of some concepts of "representation." Block's arguments fall short of establishing that P is representational and, given the sort of cognitive science assumptions he is working with, he is unable to account for the aspect of phenomenal content that he thinks goes beyond "representational" content. *BBS* discussion reveals the need for greater analysis and justification for a representationalist thesis of P.

An important question arising from discussion is whether phenomenal consciousness (P) is, itself, wholly or partly, representational. Block (1995r, p. 273) rejects the view, proposed by Armstrong (1995, pp. 247ff) and others, that P is entirely representational, that is, merely a matter of the degree of detail, specificity, and informational richness of representational content. Block agrees, however, that P *is* representational (see, pp. 278, 280). It is just that "phenomenal content" transcends "representational content." Block is careful to point out that by "representational" he does not mean "propositional." If he had, he would have used the term "intentional" (answering Tye; pp. 268–69, p. 278). It is not clear what Block thinks the phenomenal content that transcends representational content actually *is*.

An argument for P representationality is given (Block 1995r, p. 278). Explaining phenomenologically either (1) the difference between perceiving a thing from different orientations (e.g., with left versus right ear) or as located differently in space, or (2) "seeing-that" the squares but not the circles are packable, *must* appeal to the "representational features" of phenomenal consciousness itself. The argument is inconclusive. One could exploit Block's view (p. 274) that access-consciousness (A) is "parasitic on," or can come and go against a background of, P. In seeing that the squares are packable an A state occurs against a P background. This A state is a second-order, intentional–representational state. The occurrence of the second-order state presupposes *some* first-order state. This first-order state does not itself have to be "representational."

Block speaks of "seeing through" (p. 279) our P (perceptual) states to the world. "Seeing through" is not an image of representing. If the foregoing A "story" were right, there must be something transpiring internally in virtue of which we can "see through" and something transpiring to which we bring A awareness in "seeing-that." It is likely that these somethings transpiring (or ways of being self-modified) are one and the same, but one needs a substantial argument to establish that they themselves are representational somethings. In virtue of what would an internal activation or disturbance of sensory substratum instantiating phenomenal awareness constitute a *representation?*

Cognitive science *must* treat P as representational if it hopes to say anything about it. Representation in cognitive science usually means concrete particulars having representationality in functional, hence relational, properties. Representations are arbitrary mediating vehicles of content and bring with them a whole set of machinery for moving them around (manipulable in virtue of formal properties). Concrete particulars with shared formal–relational properties count as tokens of the same type. You and I can both have the same representations, but we can't have the same qualia. We can't share awarenesses because awarenesses are subject-specific, immediate material instantiations of "content" (in C. I. Lewis's [1929] original notion of "quales," cited by Lycan [1995, pp. 262–63]: i.e., "the introspectible monadic qualitative property of what seems to be a phenomenal individual"). Representations work by standing for what they are not. Awarenesses work, insofar as they do work, by being what they intrinsically are.

Block grants (p. 278) that P is "relatively" intrinsic compared to A, which is "relatively" relational. But he continues to treat P as significantly "representational." His image of A interfacing with the more fundamental P is of one parquet floor over another floor (p. 274). The suggestion is of one level and/or type of representation meeting another. In the instance of seeing-that, concerning the packability of squares, Block holds that P-contents do not represent extrinsically (R9) but are themselves "intrinsically packable," hence "represent per se." But what exactly does this amount to? A sand dune can intrinsically change its shape – does that mean

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

it represents-per-se changeability? The squares, because of their spatial properties, only actually represent packability via some propositionalizing act of A.

An ontology of representations is a powerful tool for explaining some types of content. But not every internal stimulation or activation in mentality need be a representation. Is there no other conceptual framework for understanding phenomenology? Armchair metaphors and analogies go some way to satisfying introspective intuition but fall far short of empirical or conceptual rigour. However, I will indulge in one briefly (Levine [1995, p. 261] spoke of P as a kind of *presenting to* the self, and that is basically the idea I would like to see explored). Consciousness is like a pond enclosed by the inner skin or membrane of the person. Things can be *in* the pond ("presenting") without actually impinging on the membrane. An interplay of internal and external factors determines what impinges on the membrane, and where and when. Being at the surface, membrane permits "presentings" to become representated by and to the system as a whole. Being at the surface allows representation but does not *necessitate* representation. *How* phenomenal content in its intrinsic nature gets to instantiate or effect information flow and where it does (I agree with Levine, p. 261), remains a mystery because we are nowhere near understanding what consciousness in its phenomenal intrinsic nature *is.* The point of representation is to objectify, make determinate, focus, extract for use, and integrate latent information. Apart from representationalization, latent "content" of P manifests in a four-dimensional fluctuating continuum of activations in content-sensitive vicinities of the pond of consciousness – experienced by subjects as the substratum of experiential sensitivity or awareness.

Block wonders (p. 281) whether he is up against a new version of the old debate about functionalism and qualia. The representationalists claim to have a position stronger than functionalism. That is, two tokens of the one P-state type can differ not just in functional role but in virtue of representational content. What such theorists will say is the vehicle of representational content, that is, the "occupier" of the functional role? And how does this vehicle have its representational content, since the latter is now distinguished from the functional role? It is in the very nature of the concept "representation" that the answer must be a relational, mediational one. Insofar as it is, it fails to satisfy the intuitions that phenomenal consciousness is fundamentally nonrelational and immediate: moments presenting in a fluctuating continuum, not a series of inert objects for manipulation.

Is the objection to representationality really only of determinate conceptual or propositional representation? Bachmann (1995, p. 251) suggests that basic sensational, intuitively "nonrepresentational," states are really low-level embryonic representational states that represent, for example, the bare categorial fact of sensing something as "existing"; or perhaps the fact that events are occurring in parts of one's body (what Block suggests is "me-ish" representation; pp. 275, 281). Are activations/disturbances/presentations that reach the surface, then, representations just *because* of this latent, vague existential or "selfy" informing import? But of what kind? Representations to whom? How are any kinds of representation instantiations of subjective awarenesses? There is an ambiguity in the Bachmann view. From *our* theoretical perspective someone's activations represent some state of their body, whereas that person simply by means of those activations/presentations *feels* some state of their body.

The problem might just be lack of clarity in the claim that P is "representational," in which case that point ought to be cleared up before debate is taken much further. Maybe some representationist construal of consciousness is right, but the thesis is doomed to failure in a cognitive science dominated by classical computation. That paradigm, having long dissociated consciousness from operations over object-like representations, seems now to be trying to put them back together again. However it may not matter what implementation of representations you propose. An unbridgeable gap exists between specification of representation in terms of intersubjective, external, and determinate relations and an understanding of the intrinsic, often indeterminate nature of subjective awareness. This latter is the essence of phenomenal consciousness.

Consciousness and mental representation

Daniel Gilman

Department of Humanities, College of Medicine, Penn State University, Hershey, PA 17033. **djg3@psuvm.psu.edu**

Abstract: Block (1995t) has argued for a noncognitive and nonrepresentational notion of phenomenal consciousness, but his putative examples of this phenomenon are conspicuous in their representational and functional properties while they do not clearly possess other phenomenal properties.

Block (1995t) has argued for a nonrepresentational, nonfunctional notion of phenomenal consciousness, or "P-consciousness." A mental state might be both P-conscious and A-conscious (representational, among other things). But it can neither be phenomenal in virtue of having a particular sort of content, nor in virtue of playing a certain functional role in perceptual or other cognitive processing. Several commentators – notably Armstrong (1995), Lycan (1995), Harman (1995), and Tye (1995) – have taken issue with this position. Rightly so, for it is a mistake to conceive of phenomenal consciousness as being intrinsically nonrepresentational. It is a mistake because Block is right that "the paradigm P-conscious states are sensations" (p. 232) and because we ought not to divorce study of sensation and perception from consideration of their central function, that is, to detect, discriminate, and represent sensible features of the environment (including features of the organism itself).

Block depends on our recognizing several sensory examples as manifestly non- or suprarepresentational; namely, the experience of orgasm and the sensation of intermodal differences in perception. (Presumably, we are to read "orgasm" as "male orgasm" throughout.)

First, Block claims it obvious that "the phenomenal content of orgasm" is not representational at all; then, in response to Tye's objection that we delineate the content of orgasm just as we delineate the contents of other bodily sensations (a location, an intensity, a quality, etc.). Block allows that the experience has a limited representational content but one that does not begin to explain its phenomenal content. What is missing? Block raises several possibilities. Sensations might vary phenomenally while their representational contents do not. Is that clear? Phenomenal contents may vary in a more fine-grained way than natural language labels for those contents, but is such variation obviously nonrepresentational and nonfunctional? Block admits that not all representations are expressible in English. Tye has suggested that the representational content of orgasm is complex and dynamic. Many subtly variable physiological phenomena attend orgasm. If space and decorum do not permit a more comprehensive articulation of a particular male orgasm, sensory contents might nonetheless differentiate particular orgasms just as other sensations differentiate particular pains in the foot. We have complex capacities for recognizing and discriminating states of the body, among other things. These are capacities subserved by representations; this claim is not confounded by the fact that dynamic variation within phenomena needs to be represented, nor by the fact that variation across phenomena needs to be represented. Notice, too, that there is no *in principle* problem with representational or functional accounts of the evaluative part of an experience. Magnitudes and varieties of pleasantness might be, for example, input as variables in some sort of dispositional calculus.

Block anticipates the response that contents might be specified in terms of recognitional capacities but says "that runs into the problem that recognitional capacities can work without P-content, as in blindsight." (p. 281) The theory that phenomenal differences are representational does not say that all representational differences are phenomenal. Of course, if our full discriminative and recognitional capacities, and all the functions they subserve, are

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

realized without phenomenal experience, then the theory is in trouble. But why suspect that? Such a conclusion would not strictly follow even from Block's fantasy case of "super blindsight" and, as Farah (1995) points out, there is no such thing as super blindsight (see also Gazzaniga et al. 1994).

Block also considers both an auditory and a visual experience of something overhead. This is supposed to be informationally impoverished perception: "I'm imagining a case where one has an impression only of where the thing is without an impression of other features" (p. 14). We are to conclude, from our imagining, "that there is a modal difference that isn't at all a matter of representation, but rather is a matter of how those modes of representation feel" (p. 14). So all we get is: (1) something overhead, heard; and (2) something overhead, seen.

Tye suggests that the *gedankenexperiment* fails because there will be extra visual or auditory information (size in the visual case, loudness in the auditory case) that differentiate (1) from (2). Block surrenders in the auditory case but not in the visual case, where, he says, we cannot track size with peripheral vision.

I think Block mistakes impoverished scale coding for no scale coding. But suppose he is right about size. There are further representational differences. Some object flies across the periphery of my visual field too fast for mechanisms of attention to direct eye movements for foveal scanning. So I cannot see what it is, be it bird, plane, or superman. As peripheral vision is rod-rich and cone-poor, I fail to discern the color of the stimulus. But even peripherally I see dark or light against a field. This picks out a feature of the distal stimulus and it carries information about the mode and media of transmission of information. These, collectively, represent features of the environment. Trying to "replicate" Block's *gedankenexperiment,* I fail to achieve his reported "results." But his example is more puzzling still. Presumably, Block picked the subject of orgasm because of his conviction that the phenomenon is, well, phenomenal. He takes pains to point out that "there can be no doubt that orgasm is 'phenomenologically impressive.'" (p. 273) Presumably this was supposed to encourage an intuition that there is so much to the experience of orgasm that one couldn't possibly exhaust "all that" with a representational or functional account. But what's curious about the modal tracking is that it seems so simple; it is just a way of tagging the active sensory system, and surely there is no problem imagining how a representational system might simply tag, as opposed to describe, something. What could be more representationally primitive?

On the relation between phenomenal and representational properties

Güven Güzeldere[a] and Murat Aydede[b]

[a]Duke University, Department of Philosophy, Durham, NC 27708 and [b]The University of Chicago, Department of Philosophy, Chicago, IL 60637.
[a]guven@aas.duke.edu; www.duke.edu/philosophy/faculty/guzeldere.html; [b]maydede@midway.uchicago.edu; tuna.uchicago.edu/homes/murat/index.ma.html

Abstract: We argue that Block's charge of fallacy remains ungrounded so long as the existence of P-consciousness, as Block construes it, is independently established. This, in turn, depends on establishing the existence of "phenomenal properties" that are essentially not representational, cognitive, or functional. We argue that Block leaves this fundamental thesis unsubstantiated. We conclude by suggesting that phenomenal consciousness can be accounted for in terms of a hybrid set of representational and functional properties.

Block (1995t) thinks there is a widespread confusion in the recent philosophy and neuropsychology literature regarding the function of consciousness. This confusion manifests itself in "a persistent fallacy involving a conflation of two very different concepts of consciousness" (p. 228): Phenomenal-consciousness and Access-consciousness.

According to Block, the (target) reasoning commits the fallacy of equivocation in concluding that consciousness has the function of initiating voluntary action based on the phenomenon of blindsight. The blindsight patients, under forced-choice conditions, succeed in making simple visual judgments in their blind fields accurately, all the while insisting that they are only guessing to please the experimenter, hence they never initiate relevant actions themselves.

On the basis of these facts, the two parties reach two different conclusions. The target reasoning concludes that "blindsighted patients never initiate activity toward the blindfield *because* they lack subjective awareness [phenomenal consciousness] of things in that field" (p. 242). In contrast, Block's conclusion is that it is Access-consciousness that is missing in blindsight patients (and, as such, responsible for the lack of voluntary action). Phenomenal-consciousness may or may not be missing (but that is irrelevant), and the fallacy lies in "*sliding* from an obvious function of A-consciousness to a nonobvious function of P-consciousness" (p. 232).

The fallacy claim. Clearly, the validity of Block's charge of fallacy depends critically on the validity of his distinction. Unless it is established independently that Block's distinction between A-consciousness and P-consciousness must be accepted by all, including proponents of the target reasoning, all Block's argument shows is that there is a *disagreement* between the notions of phenomenal consciousness he and proponents of the target reasoning use. And from a mere disagreement, a charge of fallacy does not follow.

Block discusses the work of Schacter (1989) as representative of the target reasoning. The notion of phenomenal consciousness that Schacter uses, however, happens to be much closer to Block's A-consciousness, not his P-consciousness. Schacter uses the term "phenomenal consciousness" to mean "an ongoing *awareness* of specific mental activity." Schacter's fundamental distinction is presented in terms of "implicit" versus "explicit" knowledge, where the former is "knowledge that is expressed in performance without subject's phenomenal awareness that they possess it," and the latter, which occurs as a result of *access to consciousness,* "refers to *expressed* knowledge that subjects are phenomenally aware that they possess" (Schacter 1989, p. 356, emphasis added).

However sketchy it may be, it is worth noting that Schacter's notion of "phenomenal consciousness" involves the sort of cognitive elements that belong to Block's A-consciousness, most notably, verbal expressibility. Block's notion of P-consciousness, on the other hand, has no counterpart in Schacter's framework. But then, Block's argument that P-consciousness does not play any role in voluntary behavior runs orthogonal to the target reasoning, since the target reasoning makes *no claim* vis-à-vis *Block's sense of* P-consciousness.

Put differently, Block's fallacy charge has some validity only when coupled with the assumption that his distinction is already established, and that his P-consciousness is the same as the target reasoning's phenomenal consciousness. Pointing out Schacter's work was one way of demonstrating that the target reasoning does not necessarily share this conceptual starting point with Block. In any case, our argument stands independent of this demonstration. Until it is established that it is Block's P-consciousness that provides the right starting place, Block and the target reasoning could only beg the question against one another on what they take "phenomenal consciousness" to be, and any charge of fallacy remains ungrounded.

Phenomenal versus representational properties. Does Block establish the validity of his distinction between A- and P-consciousness? We think not. Block tries to provide support for his distinction by presenting a number of cases that are purported to demonstrate how P-consciousness can exist in the absence of A-consciousness, and conversely. But he takes for granted a more fundamental distinction on which the plausibility of his cases rest. This is the distinction between phenomenal properties (P-properties or P-content) and representational/functional properties (R/F-properties or content). In the rest of this commentary, we will show that Block's distinction between P- and A-conscious-

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

ness is not established because the distinction between P-properties and R/F-properties is left unsubstantiated.

Block's starting point is to take "P-conscious properties distinct from any cognitive, intentional, or functional property" (p. 230). For Block, "P-consciousness, *as such*, is not consciousness *of*" (p. 232). By this, Block means that P-conscious properties are *in essence* not representational. They intrinsically constitute a *kind*, or *type*, in themselves. Echoing Kripkean intuitions, Block asserts, for example, that, "the feel of pain is a P-conscious type – every pain must have that feel" (p. 232).

But these claims are far from constituting a neutral starting point. They are rather substantially controversial philosophical theses that need to be established at the end of argument, not taken for granted at the beginning. We thus fail to see how a proponent of the target reasoning who thinks that P-properties are exhausted by R/F-properties could be expected to accept Block's fallacy charge.

In other words, the issue ultimately comes down to whether the phenomenal character of mental states can or cannot be accounted for in representational and causal/functional terms. Needless to say, there are many accounts that purport to show that it can (e.g., Dennett 1991; Dretske 1995; Tye 1995). Block thinks otherwise, especially vis-à-vis the distinction between phenomenal and representational properties (or content). (Here, we should state that by "representational content" we intend to cover both conceptualized and nonconceptualized content. We will use "intentional" to indicate conceptualized content. Thus, an R-property may be intentional or not. Roughly speaking, such a property, if intentional, is possessed by thought-like mental states; otherwise, it is a property of sensory states and the like.)[1]

Now, what precisely is Block's position on the relationship between R- and P-properties? He thinks that P-properties are essentially nonrepresentational (and noncognitive/nonfunctional), but nonetheless, "P-conscious contents can be representational" (p. 232). In other words, "P-conscious contents often have an intentional aspect, and also that P-conscious contents often represent in a primitive, nonintentional way" (p. 245, n. 4). However, "P-conscious content cannot be reduced to intentional content." That is, Block maintains (p. 234) "phenomenal content need not be representational at all (my favorite example is the phenomenal content of orgasm)."

By this, we take Block to mean that certain phenomenal properties, even though they are *in essence* phenomenal, can *contingently* be representational as well. To clarify, consider the set, **P,** of all P-properties that can be associated with a conscious mental state. Consider, also, the set **R** of all representational properties. Now, some (e.g., Dretske 1995 and Tye 1995) think that **P** is just a subset of **R** – that is, any P-property is also an R-property (but the converse does not have to hold). Perhaps some others think that **P** and **R** are mutually exclusive (cf. Katz 1995). In contrast, Block seems to think that certain P-properties may also be R-properties, but there are (can be) also certain other elements of **P** that are not elements of **R.** That is, what Block seems to have in mind here are "purely phenomenal" properties that are not representational (not cognitive/functional) at all. Call these properties P*-properties, and their set **P***. It is this set we are interested in.[2]

Block seems committed to the existence of such a set. In his reply to Lycan and Harman, he actually takes it as obvious commonsense that such a set exists: "As reflection on the example of the phenomenal content of orgasm should make clear, the idea that there is more to phenomenal experience than its representational content is just common sense from which it should take argument to dislodge us" (p. 279). But not everyone thinks so. Dretske and Tye would presumably think of **P*** as the empty set, for example. So our point, once again, is that so long as Block's fallacy charge fundamentally relies, as it does, on an unsubstantiated thesis on the relation between P- and R-properties, it remains ungrounded.[3]

A further problem: What is access to P-properties? There would remain a further problem, even if Block could convince everyone that there was indeed a nonempty set, **P***, of nonrepresentational phenomenal properties. This problem, as a number of commentators also point out (Church 1995; Kobes 1995; Levine 1995; and Rey 1995) concerns specifying the nature of "access" to such "purely phenomenal" properties. Block talks about access to P-content/P-properties. But it is not clear if the notion of "access" used here is, or can be, the same as his technical notion that is definitive of Access-consciousness.

Access, as defined by Block in the technical sense, is essentially access to only R-properties: "A state is access-conscious (A-conscious) if, in virtue of one's having the state, a representation of its content is (1) inferentially promiscuous . . . , that is, poised for use as a premise in reasoning" (p. 231). The notion of access involved in A-consciousness is thus meant to be introduced as a technically restricted notion: The content of a certain state may be accessed in this sense insofar as the content is representational.

But what about access to nonrepresentational P-content or P-properties? It cannot be access in the technical sense. It does not suffice for Block to say that *some* P-properties are also representational, for here we are interested in the nonrepresentational P-properties that belong to the set **P***. Perhaps, then, we can resort to access to nonrepresentational properties in some undefined yet intuitive sense. But what exactly is the nature of such access?

So far as we can see, this issue remains unexplicated in Block's account. Given that access in the technical sense is ruled out, the idea of "access" to P-consciousness remains mysterious. This seems to be the underlying worry Rey (1995), Shepard (1995), Harman (1995), and Lycan (1995) express in their commentaries, and it explains, for instance, why Rey thinks that if the essence of a P-property is neither representational nor functional, we cannot, even in our own case, come to know whether we have P-conscious states at all.[4]

Final remarks. In closing, we would like to leave open the question of whether all P-properties are, in fact, representational properties. But this does not necessarily leave the door open to the existence of "purely phenomenal" properties. For it may be that a hybrid set of representational, functional, and cognitive properties actually account for the phenomenal character of any given mental state.

In experiences like pain, in particular, there seems to be a natural place for each of the three kinds of properties to account for the different dimensions of its phenomenology. Roughly speaking, the representational properties can provide one with a sense of some particular type of damage occurring in a certain part of one's body (incision in the foot, burning on the fingertips), whereas the functional properties (and, in the cognitively manipulable cases of pain, cognitive properties as well) can account for the reactive/motivational aspects and the affective/emotional tone of the experience. In other words, causal/functional properties, which can account for the attractive/aversive dimensions of certain experiences in terms of an organism's special "pro-" or "con-" reaction to incoming sensory information, can, when coupled with representational and cognitive properties, constitute just the right candidate for appropriately capturing its phenomenal aspects, without leaving any peculiar and mysterious "phenomenal residue" behind.[5]

ACKNOWLEDGMENTS

We would like to thank David Chalmers and Jesse Prinz for their helpful suggestions.

NOTES

1. We would like to think of this position as being in accord with Block's, but his position with respect to conceptualized/nonconceptualized content is not all that clear. On the one hand, he seems to think that nonconceptualized content (as well as conceptualized content) can be representational, as, for example, in: "A perceptual experience can represent space as being filled in certain ways without representing the object perceived as falling under any concept. Thus, the experiences of a creature that does not possess the concept of a donut could represent space as being filled in a donutlike way" (p. 245, n. 4).

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

On the other hand, in the Response (Block 1995), Block seems forced to reject that there can be any nonconceptualized content at all: "On the substance of Tye's argument: How do we know if P is preconceptual? I used the phrase "representational" to describe P-content instead of "intentional" to allow for that possibility, but I have seen no convincing argument to the effect that P-content is preconceptual" (p. 278).

All in all, however, we think there is good reason not to think of Block as being committed to representational content as being only conceptualized content. As regards the convincing argument he is seeking, we would like to suggest Dretske's long-standing work on the nonconceptual nature of (nonepistemic) perception, which is fully representational (Dretske 1969; 1981; 1995).

2. Block sometimes talks as if R-properties are properties of P-properties (i.e., second-order properties), or vice versa. This, we think, is suggested by his use of such predications as the intentional aspects of P-content or P-properties (p. 245, n. 4). We do not think this is his real intention, but if it is, it is not altogether clear how he would work out the details of the ontology this would commit him to.

3. Actually, things take an unexpected turn during the rest of Block's reply, as he goes on to say: "Furthermore, why should believing in phenomenal contents that are *partly* nonrepresentational commit one to *wholly* nonrepresentational phenomenal contents (of the sort Katz advocates)? Perhaps Harman and Lycan think that if P-content is partly nonrepresentational, one can simply separate off the nonrepresentational part and think of it as a separate realm. But once the argument is made explicit it looks dubious. Consider the examples I used in my reply to Katz, say, the example of packability in the case of experiences as of squares contrasted with circles. Is it obvious that there is any separable phenomenal content of that experience that is phenomenal but not representational? I don't think so" (p. 280).

This is surprising. Could Block really be denying that "there is any separable phenomenal content of [an] experience that is phenomenal but not representational"? This would amount to claiming that there are no P-properties that make a state P-conscious without thereby making it a representational state – that is, that are not also R-properties. But if all P-properties are representational, why would Block think that P-consciousness is mysterious to the extent that "no one really has any idea about what P is" (p. 279), or that current research programs "contain no theoretical perspective on what P-consciousness actually is" (p. 231). We remain puzzled.

4. Some notion of "access" to nonrepresentational P-properties seems to find its analog in sense-data theories – perhaps we simply "behold" P-content with an inner mental eye. But Block cannot possibly be a friend of such ideas, as he says: "I am grateful to Lycan for explicitly not supposing . . . that the advocate of qualia is committed to sense-data or 'phenomenal individuals.' If any of us is committed to sense data, it is Lycan, Armstrong, Church, Kitcher, (and perhaps Harman) and other advocates of monitoring. The rest of us can agree with Harman (1990) that we look *through* our experiences, and that the experiences do not need to be *observed* in order to be phenomenally conscious" (p. 279). But then how does Block account for his access to **P***? Nothing in his account caught our (mind's) eye as a promising answer.

5. See Güzeldere (1997) for the development of a similar view, the "bundle thesis of qualia." See Aydede (1995) for an analysis of pain and pleasure experiences along these lines.

Empirical status of Block's phenomenal/access distinction

Bruce Mangan

Institute of Cognitive Studies, University of California, Berkeley, Berkeley, CA 94720-3020. **mangan@cogsci.berkeley.edu**

Abstract: P/A (Block's phenomenal/access) confounds a logical distinction with an empirical claim. Success of P/A in its logical role has almost no bearing on its plausibility as an empirical thesis (i.e., that two kinds of consciousness exist). The advantage of P/A over a single-consciousness assumption is unclear, but one of Block's analogies for P (liquid in a hydraulic computer) may be used to clarify the notion of consciousness as cognitive "hardware."

Block (1995t) is certainly right about one thing: Two different *concepts* of consciousness now prowl the cognitive landscape. The reaction of two of Block's referees (p. 235) is, I can attest, all too indicative: One referee thought that only Block's access-consciousness "deserves the name 'consciousness,'" yet the other wondered "why access is called . . . access-consciousness? Why isn't access just . . . a purely information processing (functionalist) analysis?"

Block tries to give both sides their due and work out a *modus vivendi* between (roughly) functionalist and antifunctionalist concepts of consciousness. P-consciousness is the robust, phenomenal, what-it-is-like concept; A-consciousness "is the information processing image of P and thus a good candidate for what P is in information processing terms" (p. 277). But while I find Block's general program refreshing, I am still confused about the precise interpretation of A and P, especially at the empirical level.

Block argues in detail for the conceptual possibility of the P/A distinction (e.g., p. 231) in order to "reveal the fallacy in the target reasoning" about a function of consciousness. But he also uses the P/A distinction to frame the empirical hypothesis that there are two different kinds of consciousness in the world: P and A are said to "interact" with one another (p. 231), to be distinct cognitive systems with presumptively different loci in the brain (p. 233), to have "something to do with the joints of nature" (p. 277), and so on.

The P/A distinction, then, looks as if it plays two very different roles – one narrowly logical, the other broadly scientific. Apparently Block thinks these roles dovetail: If the concepts of P and A are logically possible and help clarify a line of reasoning about consciousness, then we have plausible grounds to believe that two different "consciousnesses" exist.

But this is a problematic transition, open, first of all, to a purely formal objection: A concept can help clarify a line of scientific reasoning and yet refer to almost anything – to a completely imaginary entity, for instance. Block himself uses concepts about a Martian experiment on Pentagon drinking fountains to help clarify a related problem in reasoning about the function of consciousness (Note 25). But I doubt Block thinks that Martians exist because the concept of Martians is logically possible and can help isolate a formal problem in a scientific argument.

Of course the great practical difficulty with the thesis that A and P are separate kinds of consciousness is that Block cannot show us in any straightforward way how to tease them apart. Even in extreme cases such as blindsight, we are told that A and P are both absent. At one point Block straightforwardly concedes that "perhaps P-consciousness and A-consciousness amount to much the same thing empirically even though they differ conceptually" (p. 242).

But even conceptually, I am not sure that the P/A distinction is viable. One example: At first P and A seem to incorporate a clear-cut phenomenal/functional split: for example, "A-consciousness is a functional notion . . . P-consciousness is not a functional notion" (p. 232). Yet at this point, in a footnote, Block begins to pull back: "I acknowledge the empirical possibility that the scientific nature of P-consciousness has something to do with information processing" (note 10). So Block's notion of P-consciousness will bear functional attributes after all. This becomes increasingly clear in later sections, for instance, when Block says that P-consciousness could function in the senses proposed by Marcel (1986; 1988) or Schacter (1989, p. 242), or when he concludes that "learning something about the function of P-consciousness may help us in finding out what it is" (p. 245). Whatever distinguishes P from A, it is not functionality per se. So, to the degree that P's functionality can be captured in information processing terms, P collapses into A.

I do not see why Block maintains that there are two distinct kinds of consciousness. Certainly we do not need a "two consciousnesses" premise to (1) identify the logical limitation in the target reasoning or (2) distinguish, for scientific purposes, phenomenology from those cognitive function(s) consciousness may execute.

The old fashioned "single-consciousness" assumption will do much the same work as Block's A and P duo. It, too, is compatible with the view that phenomenology and cognitive function have no *necessary* connection, and this is enough to show the logical gap in

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

the target reasoning (with, say, arguments similar to those used for epiphenomenalism). And, at the empirical level, there is nothing in a single-consciousness assumption to prevent us from either distinguishing cognitive function from phenomenology, or looking for systematic links between them. In particular the single consciousness assumption is able to handle some of the more puzzling phenomena Block himself mentions – imageless thought, "just knowing," feelings of relation – when considering (without resolution) the possibility of A without P (p. 275). Both phenomenological and functional analysis of these puzzles are already underway using experimental support and standard information processing notions (Mangan 1993b) without the need for Block's more radical option.

Finally, I have a question about P and A that I don't believe Block addresses. At one point he speculates that "perhaps P-consciousness is like the liquid in a hydraulic computer, the means by which A-consciousness operates" (p. 242). Now if A "is the information processing image of P and thus a good candidate for what P is in information processing terms" (p. 277), it looks as if we have the following consequence: A could be instantiated in an indefinite number of information-bearing media without loss, but P, qua P, cannot. For P can only be a liquid or it isn't P. P-consciousness is, by analogy, "part of the hardware," while A retains the classic functionalist indifference to its particular physical manifestation. This captures one crucial feature of the functionalist/antifunctionalist dispute about consciousness (Mangan 1993a, pp. 10–14), though probably not as Block intends, since he generally takes A and P to constitute an interacting system. A-consciousness captures those features of P-consciousness that can be instantiated (functionalist "consciousness"), but A cannot capture *everything* that it is to be P (antifunctionalist "consciousness"). Or, for Block, can A *completely* instantiate P?

Perception and content

Alva Noë

Department of Philosophy, University of California, Santa Cruz, Santa Cruz, CA 95064 **anoe@cats.ucsc.edu; www.ucsc.edu/people/anoe/**

Abstract: It is argued that to have an experience is to be in a phenomenal state with A-conscious content. Perceptual contents are always both P-conscious and A-conscious.

The dubious line of reasoning about the function of consciousness which Block (1995t) criticizes concerns the phenomenon of blindsight. Blindsight patients are said to acquire perceptual contents in a P-unconscious manner. Since they are unable to use them to guide action or reason, it is hypothesized that the function of P-consciousness is to allow rational use of content. The argument is fallacious, however, because these patients also lack A-consciousness; without a demonstration of the dependence of A-consciousness on P-consciousness, no conclusions can be drawn about the latter's function. It may turn out that there could be no A-consciousness without P-consciousness, but as a conceptual matter, Block argues, the two kinds of consciousness are independent (pp. 233–35).

It is this conceptual point I want to question. For Block, the question whether a given content is P-conscious is comparable to the question whether a given sentence is true. It either is, or it is not, and it could just as well turn out to be either. P-consciousness is in this way taken to be an accidental or external property of perceptual contents. One can "have" the content – it can be available as input to the action-guiding, reasoning systems (A-conscious) – even though one has no experience as of that content. This is essentially what is imagined in the superblindsight scenario (p. 233).

But there are compelling reasons to think that P-consciousness is an internal property of perceptual contents, comparable not so much to the truth or falsity of a sentence, as to the validity of a proof. A proof which is not valid is not a proof; in just this way, a content which is "had" but not experienced (which is not P-conscious) is not a possible content of experience.

Here is my reasoning: Visual perceptual content (for example) is representational content, and the way it represents things as being is, *inter alia*, as colored. But colors are phenomenal qualities *par excellence*. As far as I can see, there is no difference between the perceptual awareness of phenomenal qualities and the experience of them. This bears on the question of whether it is correct to describe superblindsight patients as A-conscious of P-unconscious contents. I would say that any evidence that they acquire perceptual knowledge of the scene (what they say and do) would also count as evidence that they have perceptual experience. Superblindsight patients either have experience as of how things are in their environment, or they fail to acquire perceptual knowledge of that environment altogether. There is no third way.

In this last point, I agree with Dennett (1995), who thinks that the differences Block is trying to get at with his A-consciousness/P-consciousness distinction really have to do with the relative richness of the content of experience. We find it plausible to suppose that superblindsight patients have rational access to contents of which they are P-unconscious, only because the contents in question are highly impoverished. Once one controls for this variable richness of content, Dennett argues, it is *much* more difficult to imagine that the two kinds of consciousness can come apart. This is right as far as it goes, but it misses what I take to be a more fundamental point. Where Dennett sees degrees of richness, I see the absence of perceptual content altogether. To see is to learn how things are by dint of their looks, just as to hear is to acquire knowledge from the sounds of things. The sense-datum theory was mistaken to suppose that we have an immediate and certain awareness of sense-data. But the truth in that famous version of empiricism is the fact that any form of perceptual awareness of states of affairs is also always an awareness of phenomenal properties, such as colors, looks, sounds, smells, and tastes. Even the barest representation of the orientation of a line gradient will represent the lines as having some distinctively apparent qualities. Perhaps it represents the lines as colored (using "color" broadly to include shades of gray), or as having some apparent shape. In the case of a patient who can detect only motion, but no color or static form, the motion must at least be represented as having a certain characteristic "look." For perception is access to the world from a point of view, and that point of view incorporates not only one's relative position, but also the limitations and characteristics of sense modalities. What Dennett views as the impoverishment of content is, in fact, the evaporation of *perceptual* content altogether.

The claim that superblindsight patients, no less than blindsight patients, lack perceptual contents, can be supported in a different way. Perceptual contents are always, at least at some level of description, demonstrative and particular see McDowell 1994). One sees, for example, that *that* tree over *there* is in full bloom. The demonstrative identification of the object of sight enters into the content itself. The patient's contents, on the other hand, lack this direct contact with their objects. The content that just "pops into the head" is general at best, for example that there is a tree of such-and-such a description in a certain location. Block refers to superblindsight patients as acquiring knowledge, but it should be clear that the knowledge thus acquired is inferential; they must infer the presence of *a* tree from the independently known correlation between "contents popping into her head" and states of affairs. The patient lacks *perceptual* awareness; that is why super-blindsight, as much as blindsight, is a form of blindness, not a form of sightedness.

According to Block, P-consciousness is experience. I have argued that perceptual contents are intrinsically phenomenal. But there are also reasons to think that perceptual content is intrinsically A-conscious. First, experience, by its very nature, has a rational bearing on action and judgment. To have a visual experi-

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

ence that things are thus-and-so is for it to look to one as if things are thus-and-so. The fact that things look thus-and-so to one is compatible with their not being that way, and also with one's not being disposed to believe that they are. Other things being equal, however, their looking thus-and-so gives one a reason for so judging. We could not have such a reason if we lacked access to the content of experience. Second, I doubt that one would credit someone (for example) with P-consciousness as of a line gradient who was unable to say something fairly elaborate about what the gradient looks like (Dennett 1995).

A-consciousness and P-consciousness, then, are not separable properties of experience. Experience, content, and the two different kinds of consciousness are internally related; *to have an experience is to be in a phenomenal state with A-conscious content.*

Consciousness versus states of being conscious

Ernst Pöppel

Forschungszentrum, 52425 Julich, Germany. **e.poeppel@kfa-juelich.de**

Abstract: States of being conscious (S) can be defined on the basis of temporal information processing. A high-frequency mechanism provides atemporal system states with periods of approximately 30 msec to implement the functional connection of distributed activities allowing the construction of primordial events; a low frequency mechanism characterized by automatic temporal integration sets up temporal windows with approximately 3 seconds duration. This integration mechanism can be used to define S. P-consciousness and A-consciousness as conceived of by Block can be mapped onto these neuronal mechanisms.

Let us refer to particular phenomenal states as "states of being conscious" (S), and let us try to define these states operationally. In so doing I believe that what Block (1995t) refers to as P-consciousness and A-consciousness can be mapped onto neuronal mechanisms; thus, semantic confusions about "consciousness" may be (at least partly) prevented.

Equivocations are difficult to avoid if it is taken for granted that "there is consciousness" (similar equivocal problems arise with other general concepts in philosophical or psychological discourse). The locution "there is . . . " compells us to search for its "whatness," but as history shows, answers to questions about the "whatness" of general concepts like "what is time" apparently never converge. To an experimental scientist not directly involved in the traditional search in philosophy for answers to questions of the "what is . . . ?" variety, it appears that the assumption of the existence of consciousness necessarily leads to semantic difficulties. The reason for this may be that the starting point of reasoning or the underlying assumptions define so many implicit constraints for the trajectories of reasoning that these trajectories never or hardly ever cross. I believe that one can stay clear from equivocations about "consciousness" by operationally defining S. If one refers to a phenomenal state such as "I am conscious now" and not to an abstract concept such as "my consciousness in this moment," one avoids the dualistic trap, because assuming the existence of consciousness one might be inclined to treat consciousness as an independent substance in contrast to the Cartesian "*res extensa*," that is, the brain.

S can be reconstructed if one looks at the mode of temporal information processing in the brain. (Here only a short outline can be given; for more extensive discussions, see Pöppel 1994.) Each mental act is implemented by simultaneous neuronal activities in spatially distributed areas. This produces a logistical problem for the brain: How are neuronal activities linked together? One answer is that the brain endogenously creates system states within which specific neuronal algorithms bind spatially distributed activities together. These system states provide a formal basis for linking operations; they are atemporal in nature because all information processed within one such state independent of its location is treated as contemporaneous.

Converging experimental evidence indicates the existence of such system states – of approximately 30 msec duration (SS-30). It has been suggested that SS-30 msec can be used to define "primordial events." There are at least two reasons for this idea: (1) The before–after relationship of stimuli is not defined for shorter intervals; only if there is temporal independence is it possible to talk about events (i.e., a before–after relationship is a necessary condition for separate events). (2) Neuronal activities from different areas in the brain in which different elementary functions are represented are linked together; in this way supramodal states are created, comprising several attributes that are necessary to constitute events.

Neuronally, SS-30 appears to be provided by oscillations that can be observed in the midlatency response of the evoked potential and derived from experiments on temporal order threshold, multimodal distributions of reaction time (Pöppel 1970), or other experimental paradigms (1994). Direct evidence that oscillations provide the formal background for SS-30 comes from anesthesiological experiments. Only if these oscillations disappear can one be sure one has sufficient anesthesia (Madler & Pöppel 1987). If under anesthesia such oscillations are preserved, sensory information can still be processed (Schwender et al. 1994). Patients who have undergone anesthesia in which the oscillatory activities are completely suppressed often report spontaneously that no time at all has elapsed while in the anesthetic state (statements may be: "nothing has happened"; "when does the operation start?"); this phenomenal state is qualitatively different from regular sleep. These reports can be taken as one support for the hypothesis that SS-30 is a necessary condition for the availability of primordial events.

Whereas SS-30 provides elementary building blocks by throwing a discrete temporal net over the cortical mantle (and presumably other brain regions), an independent mechanism of temporal integration links successive building blocks of primordial events together. Substantial experimental evidence indicates that temporal integration is limited to 2 to 3 seconds (3sec-int). An essential aspect of 3sec-int is that it is not determined by content, that is, by what is processed. Because of the independence of what is processed, the term "linking" or "binding" for 3sec-int may be misleading, because these terms imply an integrative activity secondary to content analysis. The 3sec-int is automatic and presemantic.

What is the experimental support for automatic presemantic integration? In tasks of sensorimotor coordination such as in synchronizing regularly occurring sensory stimuli with simple movements, such stimuli can be anticipated up to approximately 3 seconds and not beyond (Mates et al. 1994). A similar time window in the motor domain has been observed in studies on the duration of intentional movements (Schleidt et al. 1987). Speech appears to be segmented in a similar temporal fashion (e.g., Kowal et al. 1975; Vollrath et al. 1992), that is, movement and speech patterns are preferentially implemented in 3-second windows. Other examples come from the perceptual domain (e.g., Gerstner & Fazio 1995; Pöppel 1994). If the duration of visual or auditory stimuli has to be reproduced, veridical reproduction is only possible up to 3 seconds; longer lasting stimuli are reproduced as shorter. Another example comes from the time course of the perception of ambiguous figures and stimuli inducing binocular rivalry. The spontaneous reversal rate for both auditory and visual stimuli indicates that each percept can only be kept for 3 seconds on average; after an exhaust time of 3 seconds, the alternative interpretation of the stimulus automatically takes over and dominates perception for the subsequent integration time. Even mnemonic processes are temporally structured in this way; if rehearsal is prevented, short-term storage is limited to the same temporal interval (Peterson & Peterson 1959).

Because of its universality in cognitive activities, it is suggested that we use the 3sec-int for a formal definition of S (Pöppel 1994). An endogenously generated presemantic integration process sets up a temporal window within which phenomenal awareness can

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

be realized. Whatever gets access to phenomenal awareness (to P-consciousness as described by Block) is represented within this temporal window of approximately 3 seconds. Thus, mental activity is segmented in time characterized by the succession of logistically independent integration intervals, each interval implementing S. Successive intervals are connected with each other on the basis of asemantic nexus, that is, by what is represented within each S. If one needs a verbal marker, the term "consciousness" could refer to the sequence of logistically independent but semantically dependent states, but this would be only a way of speaking.

The next question is: What gets access to S and what are the mechanisms that control this access? To approach this problem it is useful to ask two questions: (1) What is the potential repertoire of mental phenomena that might get access to S? and (2) What could be the functional use of S? In an attempt to develop a taxonomy of subjective phenomena (Pöppel 1989), it has been argued that four distinct domains of subjective phenomena define the content of S. For each of these domains a modular representation is suggested by neuropsychological and neuroanatomical evidence. These domains comprise in everyday language perceptions, memories, evaluations, and volitions (resulting in actions). Particular functions from each of these domains may get access to S (P-consciousness), and particular neuronal algorithms control this access. Without the operation of such a control mechanism, S may be "blind," that is, specific percepts, memories, emotional evaluations, or intentional acts may not get access to S.

The phenomenon of residual vision or blindsight (Pöppel et al. 1973) can be interpreted within this framework. The access mechanism to S has been disrupted because of a specific lesion; through an experimental trick it is still possible to demonstrate residual visual capacities, but they cannot reach those neuronal processes that set up S. Similarly, other subjective phenomena may be available in principle, but because of deficiencies in the access mechanisms (A-consciousness according to Block) they are absent. The dissociation of emotional evaluation and perceptual registration often observed in schizophrenic patients is another case in point, that is, the emotional evaluations do not get access to S. Another demonstration would be the tip-of-the-tongue phenomenon; the speaker knows that he knows something, but his knowledge cannot reach S, because the access is disrupted.

What could be the functional use of S, if any? An answer to this question could lead to those mechanisms controlling access. I believe that S plays a basic role in communication (Pöppel 1988). The selection pressure for S was such that individual states could be made available to others. To have functional communication, one needs a temporal framework with interindividual constancy. S is, thus, the expression of a logistical solution by the brain to ensure interactions by communication, the latter only being possible because of an interindividual temporal match of S.

Phenomenal consciousness and what it's like

David M. Rosenthal

Ph.D. Program in Philosophy and Concentration in Cognitive Science, City University of New York, Graduate School, New York, NY 10036-8099.
drosenth@broadway.gc.cuny.edu

Abstract: Even if A-consciousness and P-consciousness were conceptually distinct, it is no fallacy for researchers relying on a suitable theory to infer one from the other. But P-consciousness conceptually implies A-consciousness – unless one or the other is mere ersatz consciousness. And we can best explain mental states' being conscious, in any intuitively natural sense, by appeal to higher-order thoughts.

1. Conceptual versus empirical connections. Block (1995a) concedes that it's hard to come by actual cases of P-conscious states that are not A-conscious, or A-conscious states that are not P-conscious. Indeed, it's "plausible," he says, that "A-consciousness and P-consciousness are almost always present or absent together." Still, he insists, the two "differ conceptually." He concludes that even if "P-consciousness and A-consciousness . . . amount to much the same thing empirically" (p. 242), it's fallacious to infer facts about the one from facts about the other.

Few inferences rest solely on conceptual connections, however, especially in scientific investigations; typically they also rely on theoretical and empirical assumptions. So even if Block is right that A- and P-consciousness are conceptually distinct, theorists who regard them as empirically connected in suitable ways may reasonably infer one from the other. Thus Block's concession that Schacter's alleged conflation of P-consciousness with A-consciousness doesn't "cause any real problem in Schacter's theorizing" (p. 237).

Block's seventeenth century Florentine experimenters got incompatible results from measuring "degree of heat" in two distinct ways, because heat and temperature diverge empirically. If the results had coincided empirically, inferring one from the other would have been warranted despite their differing conceptually; it is the actual conflicting results that show that heat and temperature differ. Block offers nothing parallel for A- and P-consciousness; the divergences he considers between them are all science fiction or highly speculative. Conflating the two, Block thinks, closes off possibilities for theorizing; but if the two coincide, those possibilities are dead ends.

2. The pretheoretic tie. Indeed, there is, *pace* Block, a robust *pretheoretic* tie between P- and A-consciousness. Block sees all P-conscious states as having a characteristic kind of content – call it phenomenal content. And he holds, conversely, that every state with such content is P-conscious. So perhaps Block's P-conscious states are just states that have phenomenal content. Armstrong (1995) and Lycan (1995) adopt this reading, which Block (1995a) sometimes encourages by talking interchangeably of P-consciousness and a state's having P-content.

But being P-conscious is distinct from having content. What P-conscious states all have in common is that they are conscious; they differ in respect of content. Phenomenal content consists of the properties we use to sort P-conscious states into types. And states may exhibit these very content properties without in any *intuitive* way being conscious states, for example, in subliminal perception. We fix the extensions of terms for these properties by way of the conscious cases, but those terms apply equally to nonconscious states.

Since states with phenomenal content are not all conscious, if "P-consciousness" means simply having such content, P-consciousness can occur without A-consciousness – indeed, without consciousness of *any* sort. But P-consciousness would then be mere ersatz consciousness. So Block must mean more by P-consciousness, and he does: A state is P-conscious if there is something it's like to be in that state. This helps, since whenever there is something it's like to be in a state, that state is, intuitively, a conscious state.[1]

What it's like be in a state depends partly on its distinguishing content properties. What it's like to be in pain, for example, depends on properties in virtue of which pains are all alike but differ from other states, whether conscious or not. But if one is in no way conscious *of* these properties, there can be nothing it's *like for a subject* to be in that state. Phenomenal content can occur without being conscious, but what it's like to be in a state with such content cannot.

This explains why it is so hard to find convincing cases of P-conscious states that aren't A-conscious.[2] A state is A-conscious if it is poised for use as a premise in reasoning, or for the rational control of action or speech. That is because these things involve one's having access to the state in question; intuitively, A-consciousness is having access to one's own states – that is, one's being conscious *of* those states. Much in Block's discussion relies on this pretheoretic notion of A-consciousness, rather than the official connection with inference and the control of speech and action.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Some states we are conscious of are not conscious; I may think I am in a state because of a theory or what another person says. But when one is conscious of a state in a way that *seems* to one immediate, that state is, intuitively, a conscious state. (Being able to guess successfully makes the blindsighter conscious of visual states, but because this access is not intuitively immediate, we don't count the states as conscious.)

A state cannot be P-conscious unless one is conscious *of* it, and that means being A-conscious of it. So P-consciousness always involves A-consciousness. This is neither an empirical discovery, for example, about interactions between the two phenomena nor a theoretical hypothesis. It is part of how we think, pretheoretically, about consciousness. Perhaps A-conscious states occur that are not P-conscious, but P-conscious states are always A-conscious.

3. A-consciousness and higher-order thoughts. Block would resist construing A-consciousness in terms of one's being conscious of a state. His official account of A-consciousness allows for P-conscious states that are not A-conscious, that is, P-conscious states not poised for use as premises in reasoning nor for the rational control of action and speech.

But no intuitive notion of consciousness corresponds to this official account. Even if a state is poised for such use – indeed, even if it is actually so used – it is not intuitively a conscious state if the subject is not conscious of it. Indeed, though the intentional states our speech acts express are always conscious,[3] many nonconscious thoughts rationally influence what we say, and how. And many provide premises in nonconscious reasoning leading to nonconscious conclusions, which in turn may rationally influence our actions. Since states that control speech, action, and inference needn't in any intuitive way be conscious, these roles define no intuitive notion of a state's being conscious. Block finds in common sense a "notion of access" corresponding to A-consciousness (1995, p. 277), but common sense has no corresponding notion of *consciousness.* At best, Block's official account is a theoretical proposal about what it is for certain states to be conscious.[4]

But higher-order thoughts (HOTs) explain more successfully our pretheoretic notion of a mental state's being conscious.[5] Having a thought about something is one way of being conscious of it; so I am conscious of whatever states I have thoughts about. When those thoughts rely on no *conscious* inference, my being conscious of those states seems to me unmediated; so we count those states as conscious.[6] Indeed, such HOTs would result in conscious states' being suitably poised in respect of speech, action, and reasoning, even though being thus poised cannot itself secure consciousness for a state.

Block's monitoring consciousness is introspective consciousness, and so outstrips the ordinary way states are conscious. A state is introspectively conscious if it is conscious and, in addition, one is conscious of being conscious of it; thus Block's (1995) identification of monitoring consciousness with attention. Block is surely right that monitoring in this sense need not figure in either P- or A-consciousness.

But it distorts things to see HOTs in terms of monitoring consciousness. If a state is accompanied by a HOT that is itself conscious, one is introspectively conscious of the state. But HOTs, like other intentional states, need not be conscious; when they are not, the target states are conscious, but not introspectively so. Block notes (p. 234) that monitoring consciousness is somewhat intellectualized; that's because in the HOTs those cases are conscious, whereas the HOTs that accompany nonintrospectively conscious states are not.

That is why, as Block (1995) objects, some commentators simply assumed that A- and P-consciousness involve monitoring. We need not be in any way conscious of A-conscious states, on Block's official account, nor of P-conscious states if they are simply states with phenomenal content.[7] So more is needed for A- or P-consciousness to be genuine forms of consciousness. Although we are conscious of our conscious states, we normally are not conscious that we are. So monitoring, as Block construes it, is too strong; nonconscious HOTs are just right.

Conscious states are mental states we are conscious of *as mental.* When I am conscious, apparently without mediation, of my veins throbbing, I am conscious of two things: states of my veins, and a certain bodily sensation. Being conscious of the sensation as such results in its being conscious, but being conscious of the veins, as such, results in no conscious state. That is why, as Block notes (1995), HOTs about states of one's liver (as such) don't result in conscious liver states.[8]

NOTES

1. As Kitcher (1995) notes; though she also assumes, wrongly as I argue below, that this involves monitoring consciousness in Block's sense.

2. Many of Block's ostensible cases of P-consciousness without A-consciousness are really just cases of diminished or indistinct A-consciousness. Thus he speculates that the Sperling (1960) experiment may exhibit P-consciousness of all the letters jointly without A-consciousness of all of them jointly. But even in Block's own experience as a subject, there was something it's like to experience all the letters jointly; so he had access to his experience of all the letters together, and that access rationally controlled his verbal report of his own P-consciousness. Other results, such as the Lackner and Garrett (1973) dichotic listening experiment, also seem only to exhibit diminished A-consciousness, rather than none at all.

3. See Rosenthal (1990).

4. Being poised is being disposed in a certain way. So, on Block's official account, A-consciousness is a dispositional property (as he concedes [1995]) corresponding to access in the sense of being able to get at something. But the pretheoretic property of consciousness involves the nondispositional property of *actually accessing* one's states. Some dispositional properties coincide with one's consciously accessing one's mental states; for example, conscious states are reportable and introspectible. Still, a state's being conscious is not itself a dispositional property.

5. Rosenthal (1986); (1990); (1993); and elsewhere.

6. What matters is seeming unmediated to the subject. The HOT's causal history figures only to ensure that, not for its own sake, as Block supposes (1995).

7. Block sees it as beyond dispute that dogs, for example, have phenomenal states without thoughts; that's right in this weak sense of "phenomenal state." But dogs do have thoughts, and may well have unsophisticated HOTs. And it is question begging just to assert without evidence that HOTs are not needed for their states to be conscious. In any case, Block concedes that P- and A-consciousness may fall off together in lower species (Block 1995).

8. But *pace* Block (1995), HOTs may well result in repressed states becoming P-conscious; after all, suitable HOTs intuitively seem to help in "getting in touch with one's feelings."

On widening the explanatory gap

A. H. C. van der Heijden,[a] P. T. W. Hudson,[b] and A. G. Kurvink

Department of Experimental and Theoretical Psychology, Leiden University, 2300 RB Leiden, The Netherlands.

[a]heijden@rulfsw.leidenuniv.nl; www.rulfsw.leidenuniv.nl;
[b]hudson@rulfsw.leidenuniv.nl

Abstract: The explanatory gap refers to the lack of concepts for understanding "how it is that . . . a state of consciousness comes about as a result of irritating nervous tissue." By assuming that there are colours in the outside world, Block needlessly widens this gap and Lycan and Kitcher simply fail to see the gap. When such assumptions are abandoned, an unnecessary and incomprehensible constraint disappears. It then becomes clear that the brain can use its own neural language for representing aspects of the outside world. While this may not close the gap, it becomes clearer where we need new concepts.

Block (1995t) acknowledges that he cannot define phenomenal consciousness in any remotely noncircular way: "really all one can do is *point* to the phenomenon. . . . Nonetheless it is important to point properly" (p. 230). Block then points via "synonyms" "what makes a state phenomenally conscious is that there is something 'it is like' (Nagel 1974) to be in that state" (p. 228) and via "examples" "we have P-conscious states when we see, hear, smell, taste, and have pains" (p. 230).

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Block emphasizes that it is phenomenal consciousness that has seemed such a scientific mystery. In his view, "By way of homing in on P(henomenal)-consciousness, it is useful to appeal to what may be a contingent property of it, namely, the famous "explanatory gap" (p. 230–31.) To quote T. H. Huxley (1866), "How it is that anything so remarkable as a state of consciousness comes about as a result of irritating nervous tissue, is just as unaccountable as the appearance of Djin when Aladdin rubbed his lamp." And indeed, a good thing to do is to home in on the mystery via this property. It is important, however, to home in properly.

In Block's view, "it is not easy to see how current approaches to P-consciousness *could* yield an account of it. Indeed, what passes for research programs on consciousness just *is* a combination of cognitive psychology and explorations of neuropsychological syndromes that contain no theoretical perspective on what P-consciousness actually is" (p. 231). In our view, however, the situation is worse. In not taking the natural sciences seriously, philosophers like Block needlessly broaden the explanatory gap. Following Hardin's example in his ("Colour for philosophers") (Hardin 1988), we illustrate this point with examples from the colour-vision domain (audition could be used as well).

In several places in the text, Block expresses, implies, or sympathizes with the point of view that there are colours in the outer world. "Different areas of the visual cortex are differentially sensitive to colour" (p. 231); "Representations of colours . . . of a single object" (p. 231); "Mary (the woman who is raised in a black and white room" (p. 231); "the representational content *that there is a red square in front of me*" (p. 232); "Suppose he gets home by turning right at a red wall. Isn't there something it is like for him to see the red wall – and isn't it different from what it is like for him to see a green wall?" (p. 240); "When the inattentive driver stops at a red light, presumably there is something it is like for him to see the red light – the red light no doubt looks red in the usual way, that is it appears as brightly and vividly to him as red normally does" (p. 241).

If one assumes that there is red, green, black, and white in the outer world, one is indeed confronted with an *awesome chasm*. One must then not only accept and learn to understand that *there is something it is like* to be a functioning visual brain. One must also develop concepts that can make clear how (some attributes of) that *something it is like* can be identical with (some of the) properties of objects in the outer world. The requirement that in *what it is like to be this functioning visual brain* (some of) the properties of the objects in the outer world be literally reproduced is indeed mind-boggling and paralysing. Neither physics nor neurobiology provide the slightest hints about the direction in which these concepts should be sought.

If one accepts that there are no colours in the outer world, the gap becomes appreciably smaller. Of course, one must then still accept and learn to understand that *there is something it is like* to be a functioning visual brain. Now, however, there is no need for concepts that can make clear how (some aspects of) the functioning of the visual brain can make (some attributes of that) *something it is like* identical with (some of the) properties of objects in the outer world. Outside colour can than be described in terms of distributions of electromagnetic energy or photons, as physics will have it, and perceived colour can be characterized in neuroanatomical and neurophysiological terms, as the neurosciences will have it.

There is no reason to assume that the visual brain has ever attempted to reproduce in *what it is like to be a functioning visual brain* (some of) the properties of the objects in the outer world. In other words, there is no reason to assume that there was the colour-colour identity-constraint that is implied by talk about colours in the external world and that needlessly widens the explanatory gap. That there are colours in the external world is a naive idea, unsupported by physics, biology, or psychology. Ultimately it presupposes that the representation (the perceived colour) is represented (as a perceived colour). A perceptual system performs its proper function when it *distinguishes* the relevant things in the outer world. For vision, the information about these relevant things is contained in the structure and composition of the light reflected by the outer world that enters the eyes. For distinguishing the relevant things in the external world, a unique and consistent representation of the corresponding distinctions in the light is all that is required.

So, the visual brain was forced to represent whatever was important for its existence in the external world, by whatever internal means were available or could easily be come by. For this representing, the brain could do nothing but use its own language. The language the brain had available or could easily come by was the language of neurons, connections, and activity patterns over structured groups of neurons. These are the means used by the visual brain to represent the relevant information contained in the electromagnetic reflectance patterns in the outer world. There is no reason to doubt that in the near future neurobiology will be able to tell us exactly how the visual brain performs this job. And, if the intact brain uses that language, and if one happens to be that intact brain, then one perceives a coloured visual world.

Just because Block misses the point about colours, he gets into trouble with two commentators. Lycan (1995), in a footnote, says "Block would do well to note that, in particular, the puzzle of the 'explanatory gap' applies to sensory experience of the fuller, attended sort but not, or not obviously, to qualia strictly so-called" (p. 263). Kitcher (1995) says "At least I don't know how to make sense of the 'what it is like' locution other than in terms of however inchoate a knowledge or belief about a property of the subject's own states. (Were this lacking, what would be the other *relatum* of the explanatory gap with science?)" (p. 259). In other words, both commentators fail to see a gap where Block rightly thinks there is one.

What Block fails to note is that, like him, these two commentators also think that there are colours in the outside world. Lycan says "One registers such a quale whenever one perceives a coloured object as such," and, again, "Suppose he did in fact stop at a red light. Presumably the light looked red rather than green to him" (p. 263). Kitcher says "To see the rose window, however, we must have some inner state(s) that can serve as surrogate(s) for a blue array, that can carry the amazingly rich information we extract from arrays of colours," and, again, "This account explains . . . why, when people try to describe what it is like to see blue, they are drawn to characterizations that also fit blue itself – namely, cold, similar to 'seeing' purple, and so on; and why the idea that perceiving blue involves a 'bluish' state is so natural" (p. 259). So, according to these commentators, there are coloured objects and there is "blue itself."

Because Block fails to notice their use of colour words, he fails to understand how they compound his error, thus sowing confusion. Kitcher and Lycan compound the problem because they postulate colours in the outside world and, thereby, fail to see any explanatory gap between *colours in the world* and *colours in the system*. For them, if there has to be a gap, it has to be a gap between something like *colours in the system* and *attending to* or *having knowledge about* those colours. Block, not surprisingly, fails to understand where the real problem is. He responds by stating "I find this difference of opinion far more troubling than any other that comes up about consciousness. I really don't know how to explain the vast divergence we see here" (p. 280).

Block teaches us as a general rule that "if you want to get anywhere in theorizing about X you should have a good pretheoretical grip on the difference between X and things that are easily confused with it" (p. 237). This general rule is certainly valid. Especially when you wish to speculate about functions of phenomenal consciousness you should have a good pretheoretical grip on the difference between what shows up in phenomenal consciousness and *things* in the outer world that are easily confused with it. For phenomenal consciousness, Block's colours are certainly such *things*.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Author's Response

Biology versus computation in the study of consciousness

Ned Block

Department of Philosophy, New York University, New York, NY 10003-6688.
nb21@is5.nyu.edu; www.nyu.edu/gsas/dept/philo/faculty/block/

Abstract: The distinction between phenomenal (P) and access (A) consciousness arises from the battle between biological and computational approaches to the mind. If P = A, the computationalists are right; but if not, the biological nature of P yields its scientific nature.

The target article focused on the distinction between P (for phenomenal) consciousness and A (for access) consciousness. P = experience. P-conscious qualities are the qualities of experience such as the phenomenal quality of pain or the sensation of red. A state is A conscious if it is poised for direct control of reasoning, speech, and action. The interest in the A/P distinction arises from the battle between two different conceptions of the mind, the *computational* and the *biological.* The computational approach supposes that all of the mind (including consciousness) can be captured with information processing notions such as computation and function in a system.

According to this view (often called functionalism by philosophers), the level of abstraction for understanding the mind is one that allows multiple realizations: just as one algorithm can be realized electrically or hydraulically, the mind can be realized biologically or electronically. The functionalist thinks that the right level of description for characterizing consciousness is the information processing level, not the level of realization of computation, namely, the biological level. The biological approach makes the opposite bet. If P = A, the functionalist side is right about consciousness. But if consciousness has a biological nature, then the realizations are what count, and we can expect that P and A will diverge.

I hypothesized that cases of P without A exist, but that A without P may not. In all my searching and reviewing suggestions of correspondents, I have seen only one case (in humans) that may well be a case of A without P. Hartmann et al. (1991) describe a case of "inverse Anton's syndrome," an adult whose primary visual cortex had been mostly destroyed, leaving a small island of primary visual cortex. (Thanks to Ralph Adolphs for drawing this to my attention.) This patient cannot discriminate whether the room is dark or illuminated, and he insists that he is blind. If stimuli are presented in the upper right visual field (which projects to the remnant of his primary visual cortex), however, he can recognize faces, facial emotions, and read single words. Yet the patient insists that he does not see anything. When asked how he knows what the word says or whose face it is, he says things like "It clicks" or "I feel it in my mind." There is no sign of hysteria or a psycho-social situation favoring blindness; that is, no reason to believe he is self-deceived. There is damage in the parietal lobes, including the left inferior parietal region. Milner and Goodale (1995) have proposed that phenomenal consciousness requires ventral stream activity plus attention, and that the requisite attention can be blocked by parietal lesions. So perhaps this is a case of visual access without visual phenomenal consciousness. I hope that readers of this journal can comment on whether this is a genuine case of A without P.

R1. Tweaking the definition of "A"

I certainly agree with **Chalmers**'s point that we should tweak the definition of "A" so as to avoid uninteresting cracks between P and A, that is uninteresting cases of P without A or A without P. Of course, it would be easy to redefine "A" in response to each crack between A and P, resulting in an ad hoc gerrymandered notion of A. Since P *has* an information-processing role (I assume), it would be trivial to claim that there are no cracks between P and *that* role. For example, in the target article I gave an example of P which does not result in A because of lack of attention. Assume that Crick and Koch (1995) are right that visual experience is a matter of activity in pyramidal cells of the lower cortical layers of the visual areas in the back of the head. Suppose further, (as Crick and Koch also suggest) that visual information is put in a position in which it can be used for reasoning and control of behavior by being transmitted to the frontal and prefrontal cortex (in the front of the head). So a conscious event in the visual cortex becomes A-conscious by virtue of transmitting information to the frontal and pre-frontal cortex, and those events in the front are later than the P-events in the back, since it takes time for the information to get to the front. If these ideas are right, a crack would appear to open up between P and A because of the myriad ways in which the information in the back might fail to affect the front in the appropriate way. Now many functionalists (especially the variety that hold functionalism as a conceptual truth) would not be bothered by this, for functionalism is prepared to count an event as P-conscious by virtue of its effects at *other times and places.* In fact, functionalism is prepared to include in the defining role of a P-conscious event processes that don't *actually* happen, but *would* happen under certain conditions. But such features, if used to frame a type of information processing, would make it far from a natural kind of information processing.

If the claim that P = A is to be significant, A must be a genuine natural kind that is also a geniune information-processing analog of P. It was in this spirit that I defined A so as to rule out the kind of degraded access involved in blindsight as a case of genuine A. The blindsight patient cannot harness the information from the blind field without being told to guess and being given a set of alternatives. So it is best to think of access-consciousness as involving a form of access that is more full-blooded than what exists in blindsight. To rule out blindsight as a case of A without P, I defined a state as A-conscious if it is poised for *rational* control of reasoning, speech, and action. The word "rational" caused a great deal of misunderstanding and I conceded in the response to the first round of commentary that it was a misleading choice. (I never meant to rule out control that involves poor reasoning – see especially my reply to Kobes in the original response.) "Control" does the job all by itself if understood properly: the information in the blindsight patient's head about what he saw *influences,* but it does not control. In some publications, I have been defining a state as A-conscious if it is poised for *voluntary* or *direct* control. The blindsight patient's guesses are voluntary, but the contents do not control the responses in a voluntary manner. They control via an indirect pathway involving guessing.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Chalmers proposes defining "A" as "direct availability for global control"; he expands on global control saying that he has in mind especially deliberate behaviors. His "deliberate" corresponds to my "voluntary" (and "rational"), and I think both play the same role in eliminating blindsight. No significant difference so far. Also, my *poised* and Chalmers's *directly available* seem to do the same job. As I explained (Block 1995), the reason for "poised" was to rule out cases where access requires processing. For example, we all have a belief about what we had for breakfast this morning, but for many readers, that belief was quiescent until reading this sentence. If we make A a totally dispositional concept, a matter of mere accessibility, then quiescent or inactive beliefs will count as A without P. Chalmers's "directly available" seems designed to do the same job, since as he explains, it is meant to eliminate contents that take some work to retrieve.

Chalmers suggests "global control" where I specify the kind of control in much more detail, including specifying that the kind of control of reasoning must involve inferential promiscuity, that is, free use of a representation as a premise in reasoning. I don't see much difference here between Chalmers and me, but it is worth mentioning that the greater specificity does have an advantage. Consider the case I mentioned in the target article of a torture victim who represses the memories of torture. The memories exert a *global* effect on his behavior, causing him to react negatively to places and people that are similar to those involved in the torture; the memories cause slips, affect dreams, and create a global mood. Yet they are not A-conscious. The notion of inferential promiscuity is especially useful in seeing why not.

Now we come to a significant difference. Though our definitions of "A" seem more or less equivalent, there is a crucial difference in interpretation when it comes to thinking about my putative cases of P without A. I gave a number of examples that were designed to exploit the fact that access to a P-content can fail for a variety of reasons, including lack of attention and various forms of blockage. (I mentioned blockage due to repression, information processing limits, fragmentation of the self, and deactivation of centers of reasoning and planning by, for example, anesthesia.) If these cases are genuine cases of P, then they are cases of P without A, because some work would be required to access the blocked representations. Attention would have to be focused or the blockage removed. **Chalmers** does not dispute that any of my cases are cases of P; rather, he tries to avoid such cases by saying "the information was directly available all along; it simply wasn't accessed." But he is trying to have his cake and eat it too, interpreting "directly available" as *poised* for access in order to rule *out* A-consciousness of what I had for breakfast this morning and as *merely potentially available* for access to rule *in* A-consciousness in cases of inattention, repression, limits on information process, fragmentation, and anesthesia. The information about what I had for breakfast was potentially available for access, only not accessed.

Perhaps **Chalmers** will say that accessing the information about what I had for breakfast this morning involves retrieval from memory – which is why it is not access-conscious – whereas the cases of P without A that I mentioned do not. But what about repression? Accessing the repressed images of torture involves retrieval from memory too, yet Chalmers wants to see them as access-conscious. No doubt there is some way of distinguishing between the ways that memory is involved in these two cases. But recall that a candidate definition of "A" must be non-ad-hoc as well as a genuine information processing image of P. To build into our definition of "A" a very fine grained condition distinguishing between two ways of accessing memory looks ad hoc, and it raises the question of why that difference involving memory ought to be included in an information processing image of P.

In sum, there are a variety of ways in which access to representations – both P and non-P representations – can be derailed. Anyone who wants to frame a definition of "A" that cuts between the P and non-P cases to avoid cracks between P and A owes us far more than **Chalmers** has provided. Moreover, P comes in a variety of degrees, of phenomenal flavors, and of representational contents. All would seem to affect the causal properties of P-states. But that raises the issue of whether the role of P has any unity apart from its dependence on the intensity, flavor, and representational properties of the P-states that have that role. Consider the kind *feet*, which, let us suppose, is a natural category. Now consider the causal role of feet, what affects them and how and what they affect. Feet are affected by concrete and high-heeled shoes and in turn affect the air-conditioners in gas pedal plants, the breeding of animals from which shoeleather is taken, and the stock-prices of companies in the foot-jewelry industry. Is the *role* a natural item apart from the feet that mediate the causal relations? I doubt it, and I would guess that the same point applies to the role of P-consciousness.

R2. Does consciousness have a function?

The best explanation for the close correlation between P and A is that P is somehow involved in the machinery of A. By contrast, **Chalmers** favors epiphenomenalism. He objects to my claim that P greases the wheels of A on the ground that there is no conceptual entailment from neural stuff to P-consciousness, so there is no contradiction in the idea of a physical duplicate of me who is a zombie, that is, has no P-consciousness. His argument that P-consciousness must be redundant to the causal mechanisms of A-consciousness is that the zombie has the same physical causal machinery of A-consciousness as I do but has no P. Since the causal machinery works the same way with or without P, P does nothing.

But this argument takes mere logical possibilities much too seriously. Mere logical possibilities do not tell us what the real mechanisms are. Magic is logically possible. The scarecrow of Oz who thinks despite a head of straw is logically possible, but one cannot move from that to any conclusion about the *actual* mechanisms of thinking. My car does not think and has no P-consciousness, but there is a logically possible physical duplicate of it that is a sapient and sentient being whose thinking and P-consciousness plays a role in the operation of the car. In my car, the low-gas light goes on via a simple piece of machinery. That machinery is present in the magic world, but, in addition, there is another mechanism. In the magic world, the fact that the car wants to inform me of the empty tank plays a causal role that is parallel to the physical machinery but nonetheless causally efficacious. Both are causally efficacious; it is a case of overdetermination. The magic world is *merely* logically possible in the sense that there is no contradiction in it.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Sapience and sentience are present in one case, absent in the other. But no conclusion can be drawn about sapience and sentience having no effect.

Moving to a somewhat different topic, I agree with **Chalmers** that one can interpret much of the empirical work on consciousness that I criticized as assuming that P = A (that is, that P is a nondispositional state that provides the basis for the A-disposition). So some of this empirical work can be rescued in a post hoc way by making a distinction that the authors themselves did not see. I acknowledged this in the target article. But not all of this work is equally rescuable in that way. In particular, much of the reasoning I was criticizing has problems of the "trivial or false" variety. Witness Searle's reasoning described in the target article and Crick and Koch's reasoning that V1 is not part of the neural correlate of consciousness because V1 does not project to frontal cortex and projection to frontal cortex is required for direct control of behavior. This is trivial of A and false (or at least unjustified) for P. (See Block 1996b.)

One final point: **Chalmers** notes that model 3 is implausible, apparently assuming that I thought otherwise. I indicated that model 3 might be difficult to refute empirically, not because I thought the model might actually be right, but rather because of the usual problems with refuting epiphenomenalism. Refutations of the view always end up being more methodological than experimental. So called "simplicity" has to figure very strongly in refutation of such ideas.

R3. Consciousness and the self

Many of the commentators in round one felt that neither A nor P corresponds very well to the intuitive notion of consciousness. The problem was that neither P nor A required that one's *self* have access to one's own conscious states. A-consciousness is a purely information-theoretic idea that does not explicitly involve the self. One can speak of one state controlling another without explicitly putting any self in the picture. Although I mentioned various connections with the self in talking about P, none loomed large. Both **Browne** and **Rosenthal** criticize me on this basis, as did many in round one. Church (1995), Harman (1995), Lycan (1995), Kitcher (1995), and Levine (1995) criticized my view explicitly on this ground, but many of the critics in round one were obliquely critical about this. For example, Baars (1995) expanded my P-consciousness as "personal consciousness" and Armstrong (1995) suggested that "A" would be better replace by "I" for introspection.

Officially, **Rosenthal**'s conclusion is that P entails A (unless one or the other is phony), so it is useless to look for cases of P without A. I say that this is his *official* conclusion because *actually* he thinks cases of P without A are completely obvious and uncontroversial. Rosenthal has adopted the misleading strategy of *redefining* both "P" and "A" so that the *P-redefined* entails the *A-redefined*, even though in his view, as in mine, *P does not entail A.* What is misleading about this procedure is that the redefinitions are not made explicit. I confess that my first thought on reading Rosenthal's reply was that for the reason just mentioned, the disagreement between us was completely verbal. But on reflection, I see that the redefinitions he offers are natural expressions of the clash of our points of view about the importance of self-consciousness, and this clash is an important one to get clear about.

Let me explain. **Rosenthal** and I mean the same thing by "state with phenomenal content." Phenomenal contents are specific types or categories of experience such as the experience of the sensation of red or the feeling of pain. In my terminology, a state with phenomenal content is just a P-conscious state; I do not distinguish between the two. But Rosenthal rejects the equivalence, *state with phenomenal content = P-conscious state.* His argument starts with the claim that we can have *unconscious* states that belong in one of these P-content categories, for example, an unconscious pain or an unconscious sensation of red in subliminal perception. Such unconscious pains and sensations, he notes, are not A-conscious. By my definitions (which I will be using here), they are cases of P without A, and so we see that Rosenthal accepts P without A in *my* senses of these terms as uncontroversial. Indeed, Rosenthal is much more liberal about P without A than I am. I think that there is only a very remote *possibility* that subliminal perception is P and thus only a remote possibility that subliminal perception involves P-content without A. Suppose the letter "Q" is flashed too briefly for the subject to report on it, but long enough to influence later choices. Rosenthal seems to assume that such "perceptions" are states with phenomenal content. (I expect this is because he has a very thin notion of phenomenal content. But let us put this issue to one side, accepting with Rosenthal that there are uncontroversial cases of P, in my terms, without A.)

Here is where the basic difference in perspective comes in. **Rosenthal** holds that these P-without-A states (in my senses of the terms) are not *conscious* states *at all*, for there is "nothing it's like *for a subject* to be in that state." In other words, P without A, if it exists, is not *real* consciousness because it need not involve access to the self, to the subject him or herself. So he rejects my notion of P (because it is not what he thinks of as *consciousness*). He holds that the cases of P without A are not real cases of P-*consciousness* without A. Since he thinks access to the self is required for genuine consciousness, he redefines "P" as what we might call "{P + self-access}."

Okay, so that is part of the story. But we still haven't seen how a redefined P will necessarily involve A. Does {P + self-access} entail A? No, because as **Rosenthal** notes, A is a purely information-processing notion that *also* involves no connection to the self. So Rosenthal changes my A too to what we might call "{A + self-access}," my A plus the added condition that the self has access to the state. He regards this as the pre-theoretic intuitive sense of A: "Much in Block's discussion relies on this pretheoretic notion of A-consciousness, rather than the official connection with inference and the control of speech and action." So Rosenthal's claim that P entails A amounts to the claim that the redefined P, {P + self-access} entails the redefined A, {A + self-access}. There is more to the story here about why {P + self-access} entails {A + self-access}, but I will not go into it. My point is only that the claim that redefined-P entails redefined-A does not challenge my claim that P without A is at least conceptually possible. Those cases of P that *don't involve self-access* may be the very cases that do not involve A. For example, in the target article I mentioned a case in which one is having an intense conversation oblivious to a loud noise, even though one has raised the volume of one's voice to compensate for it. Once one notices the noise, one might realize that one was *hearing it all along.* This is plausibly a case of P without A, and one that does not involve self-access.

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Rosenthal notes that A is dispositional. Being poised for direct mental and behavioral control is a disposition. But consciousness is not dispositional, he says. For one's sensation of red to be conscious in his preferred sense, one must oneself be conscious of it, and that is not dispositional. I can agree that that nondispositionality of the P-conscious sensation of red shows that P-consciousness is not A-consciousness. But this does not show that we should reject A in favor of higher order thought. The sensation of red is also P-conscious. And P consciousness, like higher order thought, is not dispositional.

The key problem with higher order thought as the main or only nondispositional notion of consciousness is that it is too intellectual. Consider the dog in pain mentioned by both Kitcher (1995) and me (Block 1995). Surely a dog with a pain that *hurts* (and therefore exerts direct control of behavior) is in a conscious state in a reasonable, intuitive sense of the term even if the dog has no higher order thought about the pain! And the dog may have a conscious pain in this sense even if it does not have a sufficient grip on the concept of pain or the concept of the self to *have* the thought that "I, myself have a pain."

The verbal aspect of **Rosenthal**'s point can also be seen by noting that Rosenthal has no complaint against the naturalness or importance of what I call P-consciousness. For him, it is the category *state with phenomenal content*. And he makes no criticism of my notion of A except that, leaving out the self, it is not the "pretheoretic notion of A-consciousness." The criticism is that neither P nor A deserve to be called categories of "*consciousness*." So the verbal aspect is that the word "consciousness" should not be applied to them. But there is also an implicit substantive and nonverbal complaint, namely, that I have left out *the main thing* in a notion of consciousness. What is this main thing that I left out? For Rosenthal it is the higher order thought, one state being about another. I agree that higher order thought is important, but I have scanted it because both P and A are more primitive and fundamental. It is P that engenders the famous explanatory gap. We have a promising research program into the nature of thought. There is no reason to suppose that higher order thought will not yield to it. But there is something else that I might be said to have left out. Armstrong, Baars, Church, Harman, Kitcher, Levine, and Lycan, all first round commentators, mention some sort of connection with the self. I will try to come to grips with this issue, starting with Browne's argument.

Browne regards the relations of access to the self as the heart of the intuitive conception of consciousness. He says that reducing this intuitive conception to A-consciousness will simply leave out the intuitive idea of access to the self. Recall that a representation is A-conscious to the extent that it is poised for direct mental and behavioral control. The informational relations involved in direct control of reasoning and action (e.g., informational promiscuity) make no mention of the self and do not in any explicit way clarify the intuitive notion of self-access. So, according to Browne, reducing the intuitive idea to A is not initially promising. The other alternative mentioned by Browne is the idea of reducing the intuitive idea of self-consciousness to P (or perhaps adopting a version of P that includes it). His objection to this idea is that it is unexplanatory. P does not help to explain anything about access to the self.

R4. Deflationism about the self

My disagreement with **Browne** (and many of the other commentators) hinges on my deflationism about the self. (See White 1991, for a worked out picture along these lines.) This is not tantamount to being an eliminativist like Hume or Dennett. My view is that the upshot of work in cognitive psychology and cognitive neuroscience is that we (ourselves) are loose federations of centers of control and integration, and for this reason, the intuitive idea of the self as a monolithic integrated entity is an illusion. The conflict between the intuitive conception and the emerging scientific picture was first captured in a convincing manner in Nagel's 1971 paper, "Brain bisection and the unity of consciousness." Nagel argued that the fragmentation observed in split brain patients exists to some degree in normal people, and this challenges our intuitive concept of the self. This sort of idea has been widened and elaborated for many years now by many psychologists and neuropsychologists, for example by Gazzaniga and his colleagues (see also, Dennett 1991). Gazzaniga (1985) tries to explain many ubiquitous cognitive phenomena in terms of the relations among "sub-selves," especially the efforts of some sub-selves to rationalize the behavior of other sub-selves. The most impressive evidence involves cases where knowledge is accessible via one part of the body, but not another. Goodale and Milner (1992) note a double dissociation: some patients cannot describe the orientation of a slot but act appropriately towards it, others show the reverse. Marcel (1993) notes a situation in which blindsight patients can access information better if responding by button-presses than verbally, and better still by eye blinks. Such phenomena are observed not only in brain damaged patients, but also in normals.

So I take it that there is a good scientific basis for what might be called deflationism about the self; regarding the self as a loose federation. This fact is what underlies my disagreement with **Browne** and others. To begin, my notion of A-consciousness *does* involve the self, the only self that *really exists*. The self-consciousness that they hanker after is a mirage. For a representation to be informationally promiscuous, to directly control behavior and speech, *is* for it to be *self*-conscious, given what the self *really* is. The definition of access-consciousness is implicitly relativized to a system. For a representation to dominate activity within that system is for it to be *as* self-conscious as it *can* be. Browne's dissatisfaction with A because it leaves out the self depends on ignoring the relevant science. I said in the target article that one should take intuitions about consciousness very seriously. But these intuitions can only be taken seriously insofar as they do not conflict with scientific fact, and one of the few facts in this area is that the intuitive notion of the self is in large part illusory. So the dissatisfaction with A that many of the critics have expressed, that it does not involve any connection with the self, is a mistake. A does involve self-consciousness in the only sense in which self-consciousness is real.

R5. Is A a kind of consciousness at all?

Bringsjord points out that I waffle on whether A is a kind of consciousness that can exist without P. I expressed some sympathy (but did not actually endorse) Searle's (1992) and

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

Burge's (1996) claim that a zombie which has no P consciousness has no consciousness of any sort, even if it has the information processing aspects of P. I think they are on to something important about the ordinary notion of consciousness: P is the core and A is conceived of by many of us as a kind of consciousness only against a background of P (as I noted in the replies in the first round). But I have two reasons for seeing A as an independent kind of consciousness. First, I think we all use both A and P to some extent in thinking about consciousness. I refer the reader to the discussion of Searle in the target article. Searle officially denies that A is a kind of consciousness, but I have caught him using "consciousness" in the sense of A. There is also the curious fact I noted in the target article that many people appear to have a concept of consciousness in which A appears to be the core and P is a subsidiary sensory aspect of A that is not even necessary for consciousness. It would be surprising if there were no echo of this in those of us who officially see consciousness as P.

A second reason is that I am less concerned with our ordinary use of "conscious" than with the important scientific issue of the relation between P and its information processing image, namely, A. Are they phenomena of different sorts? Can one can exist without the other? The ordinary concepts of consciousness are vague and even if Searle and Burge are right, not too much violence is done to the ordinary concept by treating A without P as a form of consciousness. [As Block (1996b) shows, both A and P are present in the pages of *Nature.*]

R6. How representational is P?

Güzeldere & Aydede find my views on the representational properties of P incoherent. Let me summarize the relevant claims so you can judge for yourselves:

1. Some P-contents are not at all representational, or at least, there is nothing about P that *requires* that P-contents be representational. In the target article, I gave the example of orgasm, but I am not totally sure about it. What I am sure about is that what *matters* about the phenomenal content of orgasm is nothing representational.
2. So P-content is not representational *per se.* (This is just another way of saying that there is nothing about P that requires it to be representational.)
3. Some specific P contents *are* representational *per se;* that is, some specific P-contents have an essential representational aspect. The example I used was the image or visual experience of circles (as opposed to squares). I noted that it is a feature of these P contents that the squares are packable but the circles are not.
4. Some other specific P contents are representational, but not *per se.* According to me, the inverted spectrum thought experiment shows that the P-content that represents red might have represented green.

I think the appearance of incoherence that Güzeldere & Aydede are worried about comes from the ease of confusing the claim that P is not essentially representational with the claim that some specific P-contents are essentially representational. Art is not essentially representational but some items of art are.

Gamble raises the interesting issue of how P could be representational at all. She says P is an intrinsic property whereas representation is relational. But why can't an intrinsic property represent via a relation? Consider the red color of a section of a map. Suppose the redness is an intrinsic property. Still, it can be used to represent altitude. Gamble says that cognitive science must treat P as a representation if it hopes to study it. I don't see why cognitive science can't study the *function* of something that is not representational. No doubt this depends on how one chooses to define "cognitive science." But using "cognitive science" so that Gamble is right, still some other field could study P, call it cognitive biology.

R7. Is there a fallacy?

Güzeldere & Aydede say that Schacter's notion of consciousness is more like A than P. But their quotations do not seem to me to support this view. Güzeldere & Aydede quote Schacter as speaking of "access to consciousness." Is this supposed to be *access to access-consciousness?* Charity requires rejecting this reading. My view is that consciousness is a mongrel concept containing elements of both P and A. Schacter (and Crick & Koch 1995b) are closer to P than A. But the important point is that by using a single notion of consciousness (that includes elements of both P and A), they end up with a dilemma: triviality or falsehood. This also applies to Crick and Koch (1995a). If they mean A, it is trivial that V1 is not conscious; but if they mean P it is perhaps false. Consider Searle (1992): the epileptics are missing "consciousness" and therefore flexibility. If it is P that is meant, the premise is very likely false. If A is meant, the reasoning is trivial. It is trivial that missing A leads to lack of flexibility because A *includes* flexibility in the relevant sense.

Searle does not make the P/A distinction, but if we make it, we can reinterpret him as saying that P is missing in the epileptics, and that explains the missing A. But even this much charity will not save his argument, since it is very implausible that they are missing P. **Bringsjord** tries to make it plausible that this happens all the time, for example, when we drive "automatically." But this is a very implausible view of automatic driving. Here is an experiment we can all perform. Next time you are going out on a long drive, get your companion to note when you seem to have spaced out and to ask you the following question: "What did you just see?" I will tell you my result: I remember the last car I passed, the last curve in the road and the like. I have been told that pilot work using the Nissan driving simulator (at the Nissan laboratory in Cambridge MA) yields the same result: a moving window of memory of about 30–45 seconds. (Unfortunately, I have been unable to confirm this report.) Bringsjord seems to assume that because there is no long term memory of P, there is no P.

What about Searle's contradiction? **Bringsjord** gives a stunning application of the principle of charity in explaining away Searle's contradiction. I submit that my diagnosis (switching between using "consciousness" to mean A and P) was far more plausible.

R8. Representation and function

Gilman and I are to some extent at cross purposes, as I can explain by distinguishing between representationism and functionalism. Functionalism is the view that the nature of

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

experience can be completely captured by the role of experiences in the mental economy, how they affect other mental states and behavior, and how they are themselves affected by stimulation. Suppose that when I both touch and look at the corner of a cube, I have experiences in the two modalities with the same representational content but different phenomenal feels. One phenomenal feel in sight, another in touch, but no representational difference. This need not disturb a *functionalist*, since there are such large and obvious functional differences between sight and touch. A functionalist has the resources to explain the phenomenal difference. But a *representationist*, by contrast cannot accept experiences that have the same representational content but different phenomenal content, for representationism goes beyond functionalism in trying to cash out all phenomenal character in terms of the ways the world is represented to be. Similarly, a functionalist need not be troubled if the experience of orgasm has no representational content at all, for its functional role (e.g., its motivational role) can serve to distinguish that experience from other experiences.

As **Gilman** notes, I believe in a "nonrepresentational, nonfunctional notion of phenomenal consciousness." Although phenomenal consciousness represents and functions, it cannot be completely accounted for in these terms. However, I did *not* try to argue for the *nonfunctional* part in the target article. The strategy of the target article, was to try to put some of the controversies aside to discuss a distinction (between P and A) that was to some extent at least visible even if my position in those controversies is mistaken. However, I *did* argue that P-content goes beyond the representational. I did not give my strongest argument for that conclusion (namely, the Inverted Earth argument, presented in 1990; 1994; 1996) but I did make some brief remarks in that direction, discussing the impoverished representational content of orgasm (as compared with its truly impressive phenomenal character). And I also had a discussion of sensations with the same representational content in different modalities. My purpose was to head off an identification of P with A, one that surfaced in the commentaries of Armstrong, Dennett, Farah, and Tye, in the first round.

Here's why **Gilman** and I are largely at cross purposes. He argues against my point about the experience of orgasm partly by appealing to its functional properties. He says "Phenomenal contents may vary in a more fine-grained way than natural language labels for those contents, but is such variation obviously nonrepresentational *and nonfunctional*?" He summarizes my remarks about the experience of orgasm as suggesting that "there is so much to the experience of orgasm that one couldn't possibly exhaust 'all that' with a representational *or functional* account." And he notes that there is no in-principle problem to "representational *or functional* accounts of the evaluative part of an experience." (Emphasis added in all these quotations.) Sure, the evaluative function of the experience of orgasm is entirely immune from my point that this experience is *representationally* impoverished; however, I wasn't trying to argue against functionalism, but only against the stronger view: representationism.

We are not entirely at cross purposes, however. **Gilman** does also defend the representationist point of view. For example, he notes correctly that we cannot expect all of representational content to be expressible in natural language; for example, recognitional dispositions often constitute a kind of content that is not expressible in English. But are we to take seriously the idea that the phenomenal character of orgasm is exhausted by a kind of *recognition*? On the face of it, having the orgasm-experience and recognizing it are very different. Perhaps recognizing the experience *changes* the experience somewhat. But surely recognition does not wholly create the experience. (What about the first time?) And there is no plausibility in the idea that an orgasm experience requires any sort of categorization. Couldn't an animal, or even a person, have something like that experience without the recognition?

R9. P = A?

I argued that just as the concept of water differs from the concept of H_2O, so the concept of P and A differ. The real question, I suggested, was whether as a matter of empirical fact, just as water = H_2O, so P = A. (Since A is dispositional whereas P is not, what this comes to is that all and only P-states have the A role.)

Pöppel presents evidence that 30–40 Hz oscillations (each one lasting roughly 30 msec) are the basis of consciousness. For example, if a type of anesthesia is used that suppresses these oscillations, subjects feel that no time has elapsed when they wake up. ("When does the operation start?") Types of anesthesia that do not suppress the oscillations promote implicit recall of tapes played under anesthesia. (Patients exposed to a recording of a Robinson Crusoe story are much more likely to associate Crusoe with "Friday" after the operation; see Schwender et al. 1994). Pöppel mentions another interesting temporal matter: evidence for mechanisms of presemantic automatic sensory integration that take 2–3 seconds. Access to P must take place within such a 2–3 second window. So is the idea this? There are two basic mechanisms of consciousness; the 30–40 Hz oscillations underlie P, and the 2–3 second integration mechanism underlies A. I take it that with mechanisms that differ in their time scale in this way, we could have P without A. For a P event might occur and fade out before the integration required for A can take place.

Noë denies that the concepts of P and A differ. He argues that perception intrinsically involves both P and A. Even if he is right about this, it falls short of the conclusion. Perception could essentially involve two nonidentical things. Moreover, I mentioned a number of nonperceptual cases. Recall the Freudian example of the repressed image of the red room in which the patient was tortured. I argued that the repressed image could be P without being A. (The case is hypothetical, but recall that we are talking about the *conceptual* possibility that P and A come apart.) But Noe sometimes appears to use "perception" to *mean* experience, namely, P. On this interpretation, there is no doubt that experience is intrinsically P. The only issue, then, is whether experience is intrinsically A, the issue of the next to last paragraph of Noë's comment.

Noë gives two reasons why P contents must be A, but neither applies to nonperceptual cases like the Freudian case. The first is that experience by its nature has a rational import. Surely the repressed image potentially has a rational bearing, but one cannot use it unless it becomes A-conscious. The second is that he doubts that one would credit someone with P unless one were willing to credit the person with A too. But one might have all sorts of indirect

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

evidence of the P content of the red image, including the person's own testimony *after* the psychotherapy is successful and the image becomes A conscious. The patient might tell us that once he recovered access to the image, he realized that he had always had the image, but the pain associated with it kept him from acknowledging it even to the point of realizing that he had it or realizing that it showed that he had been tortured. Even if one insists on the latter sort of evidence, there could be a period during which the image was P without being A. (Some models of memory, e.g., the Headed Records view of Morton, 1991, have room for such phenomena.)

Mangan agrees that there is a conceptual possibility of P diverging from A, but he is certain that in fact P = A. He seems to think that I argue as follows: a difference in concepts, therefore difference in fact. But that is not my argument. I say that we do not know whether P = A. There is certainly reason to take apparent cases of P without A (and one apparent case of A without P) seriously. Mangan says that research on P is doing well on the assumption that P = A. But is it really doing well when we have no idea how anything physical could have P, when we have proposals that the field seriously considers drawing on quantum mechanics, whose rationale seems to be that both quantum mechanics and consciousness are mysterious? Mangan mentions my analogy: perhaps P is like the liquid in a hydraulic computer, and A is like the computation. P is the hardware implementation of A. Mangan wonders whether P can "completely" implement A. But if the analogy is correct, then we have to wonder whether there are other implementations of A, just as a given computation may be realized electrically instead of mechanically. There can be hydraulic fluid without the hydraulic computer and an electronic version of the computer without any fluid. How does Mangan rule out the analogous possibilities in the case of P and A?

Bogen wonders whether the right hemisphere might have A without P. He is sure it has A, and if his theory of P in terms of the ILN is right, it has P too. Perhaps some reader can shed more light on the issue. On dreaming, Bogen agrees with Revonsuo (first round) that dreams may be P without A. In dreaming, one's representations *are* poised to control behavior, but behavioral systems are paralyzed, so there is no behavior. Dream contents are A; so they do not provide a case of P without A.

R10. The explanatory gap

Van der Heijden et al. think that the explanatory gap is made to seem wider than it is by assuming that, for example, roses are red and violets are blue. If you suppose that a rose is red, then, according to them, you have to suppose that red is "literally reproduced" in P-consciousness. And if red is "literally reproduced" in P-consciousness, it is no surprise that it seems almost impossible to explain P-consciousness in neural terms. They suggest that we give up the "color-color identity constraint" that insists that we have red both in the world and in the mind. Here is where they go wildly, unbelievably wrong. They say that we should give up the idea that a rose or anything else is ever red. The only redness, they say, is mental redness. But why not hold instead that roses are red, giving up the idea that red is "literally reproduced" in P-consciousness? Why not reject the "color-color identity constraint" by rejecting colors in the mind? Why not construe talk of red in the mind as a misleading way of expressing the fact that P-conscious states *represent* the world as red? And a representation of red need not itself be red (like the occurrences of the word "red" here). This idea is spelled out further in Block (1983) and Tye (1995, Ch. 4).

References

Armstrong, D. (1995) Perception-consciousness and action consciousness? *Behavioral and Brain Sciences* 18:247–48. [DG, DGa, DMR, rNB]

Aydede, M. (1995) An analysis of pleasure vis-à-vis pain. Unpublished manuscript, University of Chicago. [GG]

Baars, B. (1995) Evidence that phenomenal consciousness is the same as access consciousness. *Behavioral and Brain Sciences* 18:249. [rNB]

(1988) *A cognitive theory of consciousness*. Cambridge University Press. [DJC]

Bachmann, T. (1995) More empirical cases to break the accord of phenomenal and access-consciousness. *Behavioral and Brain Sciences.* 18:249–51.[DG]

Berlucchi, G., Aglioti, S., Marzi, C. A. & Tassinari, G. (1995) Corpus callosum and simple visuomotor integration. *Neuropsychologia* 33:923–36. [JEB]

Block, N. (1980) Troubles with functionalism. In: *Readings in philosophy of psychology*, vol. 1. Harvard University Press. [SB]

(1983) Mental pictures and cognitive science. *The Philosophical Review XCII* 4:499–541. [rNB]

(1990) Inverted earth. In: *Philosophical perspectives 4: Action theory and philosophy of mind*, ed. J. Tomberlin. Atascadero. [rNB]

(1994) Consciousness. In: *A companion to philosophy of mind*, ed. S. Guttenplan. Blackwell. [rNB]

(1995r) How many concepts of consciousness? *Behavioral and Brain Sciences* 18:272–84. [DG, DMR, rNB]

(1995t) On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18:227–87. [DB, JEB, SB, DG, AVDH, EP]

(1996a) Mental paint and mental latex. In: *Philosophical issues*, ed. E. Villanueva. Atascadero: Ridgeview. [rNB]

(1996b) How to find the neural correlate of consciousness. *Trends in Neuroscience* 19:2. [NB]

Bogen, J. E. (1993) The callosal syndromes. In: *Clinical neuropsychology*, 3rd ed., ed. K. M. Heilman & E. Valenstein. Oxford University Press. [JEB]

(1995a) On the neurophysiology of consciousness: 1. Overview. *Consciousness and Cognition* 4:52–62. [JEB]

(1995b) On the neurophysiology of consciousness: 2. Constraining the semantic problem. *Consciousness and Cognition* 4:137–58. [JEB]

Bogen, J. E. & Vogel, P. J. (1962) Cerebral commissurotomy in man: Preliminary case report. *Bulletin of the Los Angeles Neurological Society* 27:169–73. [JEB]

Bringsjord, S. (1992) *What robots can and can't be*. Kluwer. [SB]

(1995) Pourquoi Hendrik est – Is une menace pour la littérature généréé par ordinateur? (trans. M. Lenoble). In: *Littérature et informatique la littérature généréé par orinateur*, ed. Alain Vuillerain. Artois Presses Universite. (English version forthcoming, MIT Press). [SB]

Burge, T. (1996) Two kinds of consciousness. In: *Consciousness: Philosophical and scientific debates*, ed. N. Block, O. Flanagan & G. Güzeldere. MIT Press. [rNB]

Chalmers, D. J. (1996) *The conscious mind: In search of a fundamental theory*. Oxford University Press. [DJC, rNB]

Church, J. (1995) Fallacies or analyses? *Behavioral and Brain Sciences* 18:251–52. [GG, rNB]

Cooper, L. A. & Shepard, R. N. (1973) Chronometric studies of the rotation of mental images. In: *Visual information processing*, ed. W. G. Chase. Academic Press. [SB]

Corballis, M. C. (1995) Visual integration in the split brain. *Neuropsychologia* 33:937–59. [JEB]

Crick, F. & Koch, C. (1990) Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2:263–75. [DJC]

(1995a) Are we aware of neural activity in primary visual cortex? *Nature* 375:121–23. [rNB]

(1995b) Why neuroscience may be able to explain consciousness (sidebar). *Scientific American* 12(95):92. [rNB]

Delacour, J. (1995) An introduction to the biology of consciousness. *Neuropsychologia* 33:1061–74. [JEB]

Dennett, D. (1991) *Consciousness explained*. Little Brown. [GG]

(1995) The path not taken. *Behavioral and Brain Sciences* 18:252–53. [AN]

Dretske, F. (1969) *Seeing and knowing*. University of Chicago Press. [GG]

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press

(1981) *Knowledge and the flow of information.* MIT Press. [GG]
(1995) *Naturalizing the mind.* MIT Press. [GG]
Farah, M. J. (1994) Visual perception and visual awareness after brain damage: A tutorial overview. In: *Consciousness and unconscious information processing: Attention and performance,* ed. C. Umilta & M. Moscovitch. MIT Press. [DJC]
(1995) Is consciousness of perception really separable from perception? *Behavioral and Brain Sciences* 18:254–55. [DG]
Gazzaniga, M. (1985) *The social brain.* Basic Books. [rNB]
Gazzaniga, M., Fendrich, R. & Wessinger, C. (1994) Blindsight reconsidered. *Current Directions in Psychological Science* 3(3):93–6. [DG]
Gazzaniga, M. & LeDoux, J. E. (1978) *The integrated mind.* Plenum. [rNB]
Gerstner, G. E. & Fazio, V. A. (1995) Evidence of a universal perceptual unit in mammals. *Ethology* 101:89–100 [EP]
Goodale, M. & Milner, A. D. (1992) Separate visual pathways for perception and action. *Trends in Neuroscience* 15:20–25. [rNB]
Graham, G. (1995) Guilty consciousness. *Behavioral and Brain Sciences* 18(2):255–56. [SB]
Güzeldere, G. (1997) *The many faces of consciousness.* Ph.D. dissertation, Stanford University. [GG]
Hardin, C. L. (1988) *Colour for philosophers.* Hacket Publishing Company. [AVDH]
Harman, G. (1990) The intrinsic quality of experience. In: *Philosophical perspectives, vol. 4,* ed. J. Tomberlin. Ridgeview. [GG]
(1995) Phenomenal fallacies and conflations. *Behavioral and Brain Sciences* 18:256–57. [rNB, SB, DG, GG]
Hartmann, J. A. et. al. (1991) Denial of visual perception. *Brain and Cognition* 16:29–40. [rNB]
Humphrey, N. (1992) *A history of the mind.* Simon & Schuster. [AVDH]
Huxley, T. H. (1866) Lessons in elementary psychology. Quoted in: *A history of the mind,* by N. Humphrey, 1992. Simon & Schuster. [AVDH]
Katz, L. D. (1995) On distinguishing phenomenal consciousness from the representational functions of mind. *Behavioral and Brain Sciences* 18:258–59.[GG]
Kinsbourne, M. (1995) The intralaminar thalamic nuclei: Subjectivity pumps or attention-action coordinators? *Consciousness and Cognition* 4:167–71. [JEB]
Kitcher, P. (1995) Triangulating phenomenal consciousness. *Behavioral and Brain Sciences* 18:259–60. [AVDH, DMR, rNB]
Kobes, B. W. (1995) Access and what it is like. *Behavioral and Brain Sciences* 18:260. [GG]
Kowal, S., O'Connell, D. C. & Sabin, E. J. (1975) Development of temporal patterning and vocal hesitations in spontaneous narratives. *Journal of Psycholinguistic Research* 4:195–207. [EP]
Lackner, J. & Garrett, M. (1973) Resolving ambiguity: Effects of biasing context in the unattended ear. *Cognition* 1:359–37. [DMR]
Lambek, J. (1961) How to program an infinite abacus. *Canadian Mathematical Bulletin* 4:295–302. [SB]
Levine, J. (1995) Phenomenal access: A moving target. *Behavioral and Brain Sciences* 18:261. [DG, GG, rNB]
Libet, B. (1993) The neural time factor in conscious and unconscious events. In: *Experimental and theoretical studies of consciousness* (Ciba Foundation Symposium 174).Wiley. [DJC]
Lloyd, D. (1995) Access denied. *Behavioral and Brain Sciences* 18(2):261–62. [SB]
Lycan, W. G. (1995) We've only just begun. *Behavioral and Brain Sciences* 18:262–63. [DG, DG, DGa, GG, AVDH, DMR, rNB]
Madler, C. & Pöppel, E. (1987) Auditory evoked potentials indicate the loss of neuronal oscillations during general anaesthesia. *Naturwissenschaften* 75:42–43. [EP]
Mangan, B. (1993a) Dennett, consciousness, and the sorrows of functionalism. *Consciousness and Cognition* 2(1):1–17. [BM]
(1993b) Taking phenomenology seriously: The "fringe" and its implications for cognitive research. *Consciousness and Cognition* 2(2):89–108. [BM]
Marcel, A. J. (1993) Slippage in the unity of consciousness. In: *Experimental and theoretical studies of consciousness.* Wiley. [rNB]
Mates, J., Müller, U., Radil, T. & Pöppel, E. (1994) Temporal integration in sensorimotor synchronization. *Journal of Cognitive Neuroscience* 6:332–40. [EP]
McDowell, J. (1994) The content of perceptual experience. *Philosophical Quarterly* 44:190–205. [AN]
Milner, A. D. & Goodale, M. A. (1995) *The visual brain in action.* Oxford University Press. [rNB]
Morton, J. (1991) Cognitive pathologies of memory: A headed records analysis. In: *Memories, thoughts, and emotions: Essays in honor of George Mandler,* ed. W. Kessen, A. Ortony & F. Craik. Erlbaum. [rNB]
Nagel, T. (1971) Brain bisection and the unity of consciousness. *Synthese* 22:396–413. [AVDH, rNB]
(1974) What is it like to be a bat? *Philosophical Review* 83:435–50. [AVDH]
Natsoulas, T. (1995) How access-consciousness might be a kind of consciousness. *Behavioral and Brain Sciences* 18(2):264–65. [SB]
Pani, J. R. (1982) A functionalist approach to mental imagery. 23rd Annual Psychonomic Society Meeting. [SB]
Páre, D. & Llinás, R. (1995) Consciousness and preconscious processes as seen from the standpoint of sleep-waking cycle neurophysiology. *Neuropsychologia* 33:1155–68. [JEB]
Peterson, L. B. & Peterson, M. J. (1959) Short-term retention of individual items. *Journal of Experimental Psychology* 58:193–98. [EP]
Pöppel, E. (1970) Excitability cycles in central intermittency. *Psychologische Forschung* 34:1–9. [EP]
(1988) *Mindworks: Time and conscious experiences.* Harcourt Brace Jovanovich. (Orig. 1985: *Grenzen des Bewusstseins,* dva, Stuttgart). [EP]
(1989) Taxonomy of the subjective: An evolutionary perspective. In: *Neuropsychology of visual perception,* ed. J. Brown. Erlbaum. [EP]
(1994) Temporal mechanisms in perception. *International Review of Neurobiology* 37:185–202. [EP]
Pöppel, P., Held, R. & Frost, D. (1973) Residual visual function after brain wounds involving the central visual pathways in man. *Nature* 243:295–96. [EP]
Revonsuo, A. (1995) Conscious and nonconscious control of action. *Behavioral and Brain Sciences* 18(2):265–66. [JEB, SB]
Rey, G. (1995) Block's philosophical anosognosia. *Behavioral and Brain Sciences* 18:266–67. [GG]
Rosenthal, D. M. (1986) Two concepts of consciousness. *Philosophical Studies* 49:329–59. [DMR]
(1990) Why are verbally expressed thoughts conscious? Report no.32, Center for Interdisciplinary Research. University of Bielefeld. [DMR]
(1993) Thinking that one thinks. In: *Consciousness,* ed. M. Davies & G. Humphreys. Blackwell. [DMR]
Schleidt, M., Eibl-Eibesfeldt, I. & Pöppel, E. (1987) A universal constant in temporal segmentation of human short-term behavior. *Naturwissenschaften* 74:289–90. [EP]
Schwender, D., Madler, C., Klasing, S., Peter, K. & Pöppel, E. (1994) Anesthetic control of 40 Hz brain activity and implicit memory. *Consciousness and Cognition* 3:129–47. [EP]
Searle, J. (1983) *Intentionality.* Cambridge University Press. [SB]
(1992) *The rediscovery of the mind.* MIT Press. [NB]
Shallice, T. (1972) Dual functions of consciousness. *Psychological Review* 79:383–93. [DJC]
(1988a) *From neuropsychology to mental structure.* Cambridge University Press. [DJC]
(1988b) Information-processing models of consciousness: Possibilities and problems. In: *Consciousness in contemporary science,* ed. A. Marcel & E. Bisiach. Oxford University Press. [DJC]
Shepard, R. N. (1995) What is an agent that it experiences P-consciousness? And what is P-consciousness that it moves an agent? *Behavioral and Brain Sciences* 18:267–68. [GG]
Sperling, G. (1960) The information available in brief visual presentations. *Psychological Monographs* 74. [DJC, DMR]
Sperry, R. W. (1974) Lateral specialization in the surgically separated hemispheres. In: *Neuroscience 3rd study program,* ed. F. O. Schmitt & F. G. Worden. MIT Press. [JEB]
Stich, S. P. (1978) Beliefs and subdoxastic states. *Philosophy of Science* 45:499–518. [DB]
Tye, M. (1995a) Blindsight, orgasm and representational overlap. *Behavioral and Brain Sciences* 18:268–69. [DG]
(1995b) *Ten problems of consciousness.* MIT Press. [GG, rNB]
Vollrath, M., Kazenwadel, J. & Krüger, H.- P. (1992) A universal constant in temporal segmentation of human speech. *Naturwissenschaften* 79:479–80. [EP]
White, S. (1991) *The unity of the self.* MIT Press. [rNB]
Zaidel, E. (1978) Concepts of cerebral dominance in the split brain. *Cerebral correlates of conscious experience,* ed. Buser & Rougeul-Buser. Elsevier/North-Holland Biomedical Press. [JEB]
Zaidel, E., Zaidel, D. W. & Bogen, J. E. (1996) Disconnection syndrome. In: *The Blackwell dictionary of neuropsychology,* ed. J. G. Beaumont, R. Keneally & M. Rogers. Blackwell. [JEB]
(1995) *Behavioral and Brain Sciences.* Editorial commentary BBS 18:272. [SB]

https://doi.org/10.1017/S0140525X97240057 Published online by Cambridge University Press