

Benjamin Kozuch

The Received Method for Ruling Out Brain Areas from Being NCC Undermines Itself

Abstract: *Research into the neural correlates of consciousness (NCC) aims to identify not just those brain areas that are NCC, but also those that are not. In the received method for ruling out a brain area from being an NCC, this is accomplished by showing a brain area's content to be consistently absent from subjects' reports about what they are experiencing. This paper points out how this same absence can be used to infer that the brain area's content is cognitively inaccessible, in which case we would expect its content to be absent from subjects' reports whether its content is (phenomenally) conscious or not. If so, such reports cannot count as evidence against that brain area being an NCC, and the received method fails. An alternative method (one suggested in Block, 2007) is considered.*

1. Introduction

Following a protracted absence, the scientific study of experience is again in full bloom. Something currently pursued with enthusiasm is research into the *neural correlates of consciousness* (NCC), those brain areas (or processes)¹ whose activity is closely allied with

Correspondence:

Benjamin Kozuch, University of Alabama, Philosophy Department.

Email: bigben@email.arizona.edu

¹ In this paper, we focus only on the research programme devoted to finding the *brain areas* that are NCC, rather than processes. Henceforth, I drop mention of the latter. For examples of work in the latter research programme, see Crick and Koch (1990), or Hameroff and Penrose (1996).

conscious experience. While there is not yet anything approaching unanimity concerning what the NCC are, the list of candidates at least appears to be shrinking. However, there is a problem that any science of consciousness faces, one that I think proves particularly troubling for NCC research.

A scientific study of consciousness of course requires data concerning consciousness. But what a subject experiences is something discovered only through report. There being this intermediary creates an epistemic space in which a problem appears: any time a subject reports not having experienced a stimulus, it is possible that the stimulus was experienced, it is just that the experience was cognitively inaccessible, and therefore unreportable. To some, this might seem an abuse of the concept of consciousness: the possibility just discussed implies that one could have a conscious mental state and yet not know it, but is this not a contradiction? Indeed, it sounds contradictory. But it need not be.

Here, it helps to appeal to Block's distinction between phenomenal and access consciousness (1995; 2002). *Phenomenal* consciousness refers to the experiential aspect that some mental states have: a mental state is phenomenally conscious if and only if there is something that it is like for one to have that mental state (Nagel, 1974; Block, 1995; Chalmers, 1995).² There is something it is like for one to experience the feel of velour, to have hunger pangs; it is *experiential* mental states like these that phenomenal consciousness picks out. *Access* consciousness refers to mental states having a certain kind of cognitive availability: some mental state is access conscious if and only if its content is poised for use in rational and voluntary control of action — including making reports.³

Appealing to these concepts, we can see more clearly the problem that a scientific study of consciousness might face: in instances where

-
- ² The inclusion of 'for one' in 'something that it is like for one to have that mental state' is sometimes thought to lend support to higher-order theories of consciousness, since it implies that each time a mental state is conscious, it is because some subject is 'aware' of the mental state (Rosenthal, 2002). This 'awareness', however, can be understood in a robust or deflationary way, and only the more robust understanding gives support to higher-order theories (Block, 2007).
- ³ There seems to be some variation in how the concept of 'phenomenal consciousness' is understood. On a weaker understanding, a mental state's being phenomenally conscious implies only that it is *experientially* conscious, in the way described by Nagel and Chalmers. On a stronger understanding, a mental state's being phenomenally conscious implies not only that it is experiential, but that it also fails to be access conscious.

a subject reports not having (phenomenally) experienced a stimulus, the scientist wants to take this to mean that there was in fact no experience of the stimulus; however, there exists the possibility that there was such an experience, it was just not accessed. (Hereafter, I use the terms ‘phenomenal consciousness’, ‘consciousness’, and ‘experience’ equivalently, except where noted. I also use ‘access consciousness’ and ‘access’ equivalently.)

Naturally, this problem arises only if it is possible that experience occurs in absence of access. Block argues strenuously for this thesis (1995; 2005; 2007; 2011; see also Lamme, 2006; 2010), but it has its share of detractors (e.g. Chalmers, 1997; Rosenthal, 2002; Clark, 2009; Cohen and Dennett, 2011).⁴ This issue is probably not resolved soon — perhaps not surprisingly, it turns out difficult to give forceful, non-question-begging arguments for or against access consciousness being necessary for phenomenal consciousness. Given this, it is worthwhile to wonder what consequences the possibility of unaccessed experiences would have for a science of the NCC. That is what this paper does.

As explained below, accomplishing the goals of NCC research requires the ability to identify not just those brain areas whose activity is constitutive of conscious experience, but also those whose activity is not. What I will here argue is that, if it is possible that experience occurs in absence of access,^{5,6} then what seemed our most promising method for ruling out brain areas from being NCC turns out to not work — not *at all*. I will summarize the argument.

In the received method for ruling out brain areas from being NCC, one establishes that a brain area is not an NCC⁷ by showing that its

-
- 4 See fn. 22 for a more comprehensive list of those commentators on either side of this debate.
- 5 The modal claim here is one of *epistemic* possibility. Putting it roughly: some proposition P is epistemically possible if, given one’s present knowledge, one is not justified in thinking that P is false. So the assumption adopted in this paper is that, given our present knowledge, we are not justified in thinking that phenomenal consciousness never occurs in absence of access consciousness.
- 6 The antecedent of this conditional could be understood as saying that it is possible that there are *unaccessed* experiences, or that it is possible that there are *inaccessible* experiences. The assumption adopted below is that both are possible. This issue is discussed further in 2.4.
- 7 The phrase ‘brain area A is not an NCC’ is shorthand for ‘activity in A is never conscious’. Thus when looking for those brain areas that are not NCC, what we are looking for are those brain areas whose activity *never* contributes to conscious experience.

content is consistently absent from subjects' reports about what they are experiencing. But some content's being absent from reports also acts as evidence for its content being inaccessible, since its content being inaccessible would be a good explanation as to why it never shows up in reports. So if a brain area's content is inaccessible, we should expect it to be absent from subjects' reports whether it is phenomenally conscious or not, and its failure to show up in subjects' reports provides *no* evidence for its content being not conscious. In short, the received method for ruling out brain areas from being NCC fails whenever it is used. As I argue below, this negatively affects the prospects for accomplishing some goals of NCC research, including determining the 'neural natural kind' of experience (Block, 2007).

The paper proceeds as follows: in Section 2 I provide background, explaining some of the goals and focus of current NCC research, and describing the received method for ruling out brain areas from being NCC. I also discuss a problem that the use of reports as data faces if it is the case that experience might occur in absence of access. In Section 3 I present the argument described two paragraphs above, concluding that the received method cannot work. In Section 4 I discuss consequences this has for NCC research.

2. The Received Method for Ruling Out Brain Areas from Being NCC

This section performs stage-setting. I start by explaining why accomplishing the goals of NCC research relies on being able to rule out brain areas from being NCC, and then describe the received method for doing this. Next I point out how this method requires using *negative reports*, reports in which subjects indicate some content is not a part of their experience. Then I describe a problem that any use of negative reports potentially faces, one that I later argue is particularly threatening to the received method.

2.1. Goals of NCC research

When discussing the target of current NCC research, it should first be noted that the term 'neural correlates of consciousness' is a bit of a misnomer. Most NCC researchers seek more than that which merely correlates with consciousness, it being more accurate to say that they

are searching for the *neural basis* of consciousness (Kanwisher, 2001; Crick and Koch, 2003; Block, 2007; Tononi and Koch, 2008).^{8,9} It is also important to note that it is specifically conscious *experience* (AKA ‘phenomenal consciousness’) that most NCC researchers take to be the target of their research (e.g. Baars, 1988; Milner and Goodale, 1995/2006; Damasio, 1999; Crick and Koch, 2003; Tononi and Koch, 2008; Lau, 2010; but see Dehaene and Changeux, 2004).¹⁰ It is conscious experience that gives rise to the vexing problems standing in our way of understanding consciousness (Chalmers, 1995; 1996). Indeed, this is why there is such excitement surrounding NCC research: perhaps such research will shed light where philosophical approaches have thus far been unable to (Baars, 1988; Damasio, 1999; Crick and Koch, 2003; Lau, 2010; but see Chalmers, 1996, Chapter 3).¹¹ Among the various goals of NCC research, there are two that we focus on below.

The first is to determine which brain areas compose the *content* NCC (Frith, Perry and Lumer, 1999; Chalmers, 2000; Rees, Kreiman and Koch, 2002; Crick and Koch, 2003; Hohwy, 2009; Overgaard and Overgaard, 2010), these being those brain areas¹² that directly determine the content of experience.¹³ More precisely, some brain area A is

-
- ⁸ We can understand the neural basis of an experience E to be whatever neural states would have been *minimally sufficient* for the occurrence of E (Chalmers, 2000): for any experience E, there will probably be some neural states that correlate with E, but which are not essential to the occurrence of E; the idea here is that such *mere* correlates can be excised by defining the term ‘neural basis’ in terms of minimal sufficiency.
- ⁹ The particular term ‘neural basis’ is not used by all of the researchers just cited, but it is clear that they have something along these lines in mind.
- ¹⁰ I remind the reader that this paper adopts the weaker understanding of ‘phenomenal consciousness’, it being equivalent to ‘experience’ (see fn. 3). And so none of the researchers just cited should necessarily be taken to countenance the idea of experience without access.
- ¹¹ Crick and Koch, for example, have argued that finding the NCC might make ‘the problem of qualia clearer’, it being ‘fruitless to approach this problem head on’ (Crick and Koch, 2003, p. 119).
- ¹² That there would be *brain areas* that qualify as content NCC assumes that the neural basis of consciousness can — to some significant degree, and in some sensible fashion — be localized to certain brain areas. This is a common assumption in NCC research (see, for example, Milner and Goodale, 1995/2006; Chalmers, 2000; Rees, Kreiman and Koch, 2002; Zeki, 2003), one that this paper joins in.
- ¹³ Content NCC are often contrasted with *background* NCC, those neural systems whose activity correlates with general modes or levels of experience, such as dreaming, being wide awake, or being disorientated. Some have hypothesized the background NCC are in the upper brainstem (e.g. Crick and Koch, 2003; Block, 2009).

a content NCC if and only if content C being represented in A (assuming certain background conditions obtain)¹⁴ is sufficient for C being represented in experience (*cf.* Chalmers, 2000). It could be, for instance, that there is a brain area such that whenever it represents redness, this constitutes an experience as of redness; if so, such an area is a content NCC for colour experience. Much contemporary NCC research is directed at finding the content NCC.

A second goal of NCC research is to discover what we might call *the* NCC, the type of neural activity both necessary and sufficient for experience. Following Block (2007), we can call this the ‘neural natural kind’ of experience. (It is of course possible that there is more than one type of neural activity sufficient for experience, a complication mostly set aside in this paper.)¹⁵ In the case of visual experience, hypotheses concerning the neural natural kind include recurrent loops between the striate and extrastriate visual cortex (Lamme, 2006; Block, 2007; Gennaro, 2012, Chapter 9), synchronized activity between prefrontal and visual areas (Kriegel, 2007), and intermediate-level visual processing that is constituted by gamma vectorwaves (Prinz, 2012).

Important to accomplishing both goals is the ability to show that a brain area is not an NCC. In the case of locating the content NCC, it is important because this goal involves determining not just which brain areas are content NCC, but also those that are not. (That this latter part is a goal of NCC research is made apparent by some examples in 2.2.) The ability to rule out brain areas from being NCC is also important in instances where there are competing hypotheses as to which brain area is the neural basis of some type of conscious content (e.g. experiences as of motion), since it allows us to eliminate some of the hypotheses. In the case of the second goal, that of finding the neural natural kind of experience, being able to rule out brain areas from being NCC is important for showing that some neural property P is *necessary* for experience, since one important kind of confirmation of this hypothesis would come from instances where a brain area both lacks P and is not an NCC.

¹⁴ For example, it might be necessary for the background NCC (see previous footnote) to be operating properly.

¹⁵ A more sophisticated way of understanding this goal of NCC research would recognize this possibility, saying that the goal is to discover the set of neural properties such that (a) at least one of the properties in the set must obtain if there is to be experience, and (b) any of the properties in the set individually obtaining is sufficient for experience.

Plausibly, there are other reasons for which ruling brain areas out from being NCC is important to a science of consciousness.¹⁶ Next we examine what would be the obvious method for accomplishing this.

2.2. *The received method for ruling out brain areas from being NCC*

Finding the content NCC is often thought to be accomplished through a process of *content matching*: one locates the content NCC by finding those brain areas whose content systematically matches the content of experience (Chalmers, 2000; Noë and Thompson, 2004). The motivation is clear enough: say we found a brain area such that, whenever it represented content C (e.g. redness), C was represented in experience; plausibly, this would be evidence for that brain area being the neural basis of experiences that have C as their content. More to present purposes, content matching also provides a way to show that brain areas are *not* NCC. This is accomplished by finding *mismatches* in content between brain areas and experience. This method is also intuitive: say some brain area represented content C (e.g. greenness) though the subject was not having an experience as of C (she was instead, say, experiencing redness, or no colour at all). It would be odd to think that the brain area in question was the neural basis of that experience.¹⁷ Given this, a mismatch in content between experience E and brain area A can count as evidence against A being an NCC. Furthermore, were it found to be the case that the content of A was *consistently* mismatched with the content of experience, one might think that this provides justification for thinking that the brain area is in fact not an NCC. This latter idea — that one shows that a brain area is not an NCC by demonstrating its content to be consistently

¹⁶ Being able to rule out brain areas from being NCC is likely also important in instances where there are multiple candidates for whatever neural property is sufficient for experience, since an instance in which some brain area both has P and is not conscious falsifies the hypothesis that P is sufficient for experience, thereby eliminating it as a candidate. Similarly, instances in which some brain area both lacks P and fails to be conscious might, under proper circumstances, be considered to confirm the hypothesis that P is sufficient for experience.

¹⁷ Justification for this is found in what is known as the ‘isomorphism constraint’, which says that, as a matter of nomological necessity, if some brain area A forms the neural basis of some experience E, then A has the same representational content as E (Noë and Thompson, 2004).

mismatched with the content of experience — I will call the *received method* for ruling out brain areas from being NCC.¹⁸

One brain area looking like it might succumb to the received method is V1 (also known as the primary visual cortex), which is the low-level visual area where a majority of visual information first enters the cortex. There are numerous experiments seeming to demonstrate content mismatches between V1 and visual experience. Probably the best-known of these would be a series of single-cell studies involving binocular rivalry (Logothetis and Schall, 1989; Leopold and Logothetis, 1996; Sheinberg and Logothetis, 1997), in which relatively stable activity was observed in V1, in contrast to the alternating images experienced by subjects during binocular rivalry. In another study, very fine gratings, ones subjects report to look like a uniform field, nonetheless produced after-effects in V1 that corresponded to the grating's orientation (He and MacLeod, 2001). And, in another experiment, V1 was shown to represent rapidly flickering colours that subjects experienced as only one colour (Gur and Snodderly, 1997). Of course, these examples, on their own, only show that activity in V1 is sometimes not constitutive of visual experience. But the idea here would be that, as content mismatches between V1 and conscious experience continue to amass, it gradually increases the degree to which we are justified in thinking that V1 is not an NCC (Chalmers, 2000; Rees, Kreiman and Koch, 2002; Tononi and Koch, 2008). For similar reasons, some commentators have been tempted to say that the so-called *dorsal visual processing stream* is not an NCC (Milner and Goodale, 1995/2006; Jacob and Jeannerod, 2003): The dorsal stream is thought to provide those representations that guide visuomotor action, and yet there are numerous studies in which visuomotor actions appear unaffected by consciously experienced visual illusions (e.g. Bridgeman *et al.*, 1979; Aglioti, Goodale and Desouza, 1995; Ellis, Flanagan and Lederman, 1999; Kroliczak *et al.*, 2006).

The received method looks important to NCC research, since it appears able to provide strong evidence against a brain area being an NCC. As seen above, this is something important to accomplishing the goals of NCC research.

¹⁸ I note that the received method is not *explicitly* used in NCC research. Nonetheless, the reasoning embodied in the received method is frequently at work in NCC research, if only implicitly (see, for example, Chalmers, 2000; Prinz, 2000; 2012; Rees, Kreiman and Koch, 2002; Brogaard, 2011).

2.3. *Negative reports and the received method*

Something needed for any scientific study of consciousness is data regarding consciousness itself. But consciousness is usually considered private: while each person knows consciousness intimately, one can be acquainted only with one's own experience (Nagel, 1974). That experience would be a first-person phenomenon poses a problem for a *scientific* study of consciousness, since it is often thought that only data that are *public* are admissible in science (Hempel, 1952; Feigl, 1953; Sellars, 1956; Popper, 1959; Railton, 1985; Baars, 2003). In response to this, there has been broad agreement that it will be *reports* that constitute the publicly available evidence concerning consciousness (see, for example, Marcel, 1988; Chalmers, 1998; Baars, 2003; Jack and Roepstorff, 2003; Rees, 2007).¹⁹

Below we are concerned with *negative reports*, reports in which a subject indicates some content is not a part of her experience. When a subject reports experiencing one continuous colour, for instance, though V1 represents the stimulus as flickering, this counts as a negative report about the content of V1. Negative reports are what make content mismatches possible: it is only through the subject giving a negative report about the content of a brain area that we could discover that the brain area's content is mismatched with experience. Thus the received method, with its reliance on content mismatches, must make use of negative reports.

2.4. *A problem for negative reports*

Recall Block's distinction between phenomenal and access consciousness: phenomenal consciousness picks out mental states that are experiential, access consciousness picks out mental states the contents of which can be used in rational, voluntary control of action, including reports. The phenomenal/access consciousness distinction highlights a potential problem for negative reports, which is that any time a subject gives a negative report about some content, the content might actually be phenomenally conscious, but fails to be reported because it is not access conscious.

¹⁹ The idea of a 'report' should not be equated with *verbal* report. As the term is commonly used in psychology, one can report being conscious of the stimulus by pushing a button, raising a hand, giving a nod, and so on.

We consider an example. Damage to the right parietal cortex can result in *hemispatial neglect*, a failure to acknowledge objects in the left visual field (Driver and Vuilleumier, 2001).²⁰ A neglect patient might not notice that the left half of a depicted horse is actually a cow (Peru *et al.*, 1996), or that the left half of a sketched house is in flames (Marshall and Halligan, 1988), though indirect measures indicate that such information continues to be represented (the neglect patient might say that an unaltered horse picture looks ‘more real’, or pick the house that is not on fire as the one she would prefer to live in). Because of parietal areas’ known role in attention, neglect is often thought a disorder of attention. But some commentators go beyond this, taking it also to be a disorder of consciousness, claiming that in hemispatial neglect there is no experience of left-located items (Driver and Vuilleumier, 2001; Koch and Tsuchiya, 2007; Prinz, 2007; Vosgerau and Newen, 2008).²¹ For others, this is hasty: perhaps when the neglect patient fails to acknowledge an object there is an experience of the object, but an inability to access the experience prevents its report (Lamme, 2006; Block, 2007; Jacob and De Vignemont, 2010). A lack of access might be expected given the impairments of attention that neglect involves.

Hemispatial neglect makes vivid what might be a general problem for negative reports: in any case, where a subject reports some content C is not a part of her experience, the report underdetermines whether C is experiential or not, since the negative report might result, not from C failing to be experiential, but rather from its failing to be accessed.

Of course, this is not a problem if experience never occurred in the absence of access. In the argument presented in the next section, I take the possibility of experience without access as an assumption, a starting point for an investigation. I do not think that this is helping myself to too much, since the idea that experience occurs without

²⁰ Hemispatial neglect can also manifest as a failure to acknowledge the left part of individual objects, the coming two examples being instances of this.

²¹ More precisely, we can say that these commentators mean to deny that any of the visual states representing left-located items are experientially (i.e. phenomenally) conscious. Here, it is hard to choose terminology that makes everyone happy, as some higher-order theorists (e.g. Rosenthal, 1997) would say that the hemispatial neglect patient is conscious of the left-located items (‘transitively’ conscious), in virtue of her visual system bearing representations of left-located items, but that these representations *themselves* fail to be (‘intransitively’) conscious. The issue we are concerned with here is whether these representations are experientially conscious.

access is currently the subject of vigorous debate,²² one probably not resolved soon. Indeed, there are a number of experiments plausibly interpreted as instances of experience without access. For example, in some studies of hemispatial neglect,²³ neglected faces have been shown to bring about significant activation in the *fusiform face area* (Driver and Vuilleumier, 2001; Rees *et al.*, 2000; 2002), a high-level visual area sometimes thought to be the neural basis of conscious face perception (Tong *et al.*, 1998).²⁴ In addition, there are a collection of studies based around partial report protocols (Sperling, 1960; Landman, Spekreijse and Lamme, 2003; Lamme, 2010) that are sometimes interpreted as instances in which subjects consciously perceive more than they are able to access and therefore report (Block, 2007; 2011; Lamme, 2010), in part because the abilities subjects display in these experiments are thought to rely on items being represented in recurrent loops between visual areas (Lamme, 2004), a type of neural activity closely associated with conscious perception (Pascual-Leone and Walsh, 2001; Silvanto *et al.*, 2005; Camprodon *et al.*, 2010). Naturally, whether one takes these studies (see also Vandenbroucke *et al.*, 2013) to demonstrate experience without access will largely be a function of one's pre-existing theoretical predilections. But the point here is not that the explanation according to which these are instances of experience without access is our *only* plausible explanation of these experiments, but rather that it is *one* of them.

So there seems to be reason to think that the possibility of experience without access is — at very least — not to be brusquely dismissed. But this leaves open the possibility that content must at least be *accessible* if it is to be experiential. Consider, however, that accessibility is a *dispositional* property: some content is accessible if and only if it has the *potential* to be accessed (e.g. if it would be accessed, if mechanisms of access attempted to recover it). But the property of being conscious appears *occurrent*: some content is

²² Commentators giving arguments suggesting that phenomenal experience and access experience cannot come apart include Dennett (1995); Chalmers (1996; 1997); Kriegel (2006); Clark (2007; 2009); Rosenthal (2002; 2007); Brown (2011); Cohen and Dennett (2011); those giving arguments suggesting that they can come apart include Wolfe (1999); Dehaene and Naccache (2001); Lamme (2006); Dretske (2006; 2007); Wallhagen (2007); Block (1995; 2001; 2007; 2011).

²³ More specifically, the studies involve a type of neglect known as *visual extinction*, a disorder in which subjects will ignore a left-located item if a right-located item is presented simultaneously.

²⁴ This last example comes from Block (2007).

conscious if and only if it *actually* manifests the property of being experiential. So, if accessibility is necessary for experience, then it looks as if some content's having a *dispositional* property (accessibility) can be used to explain its having an *occurrent* property (being experiential). But this is just the opposite of how one would think that dispositional and occurrent properties are related: intuitively, the dispositional property of fragility that a glass possesses is explained by whatever occurrent (probably microphysical) properties it has, not the other way around. Indeed, it is sometimes doubted whether dispositional properties are causally efficacious at all (Prior, Pargetter and Jackson, 1982).²⁵

Additionally, the idea that accessibility is necessary for experience seems committed to strange predictions. Consider a case in which there is some accessible content C in brain area A, which is conscious from times t-1 to t-2 (somewhere around a few seconds), but never accessed. Next consider the same scenario, but this time, at some point between t-1 and t-2, the connections between A and whatever brain areas might have accessed its content are severed, leaving A's other connections untouched. Now, if accessibility is necessary for experience, then at the very instant that the connections are severed, C becomes unconscious. And this is the case even though the connections between brain area A and whatever brain areas might have accessed content C were inactive at the time of severance; and even though activity in A continued the same before and after the severance. Certainly, one could not (without begging the question) rule this scenario out, that C's losing the dispositional property of accessibility would immediately result in C losing the occurrent property of being conscious; but this scenario is at least less intuitive than one in which C continues to be conscious.

What we saw earlier is that it is hard to dismiss the possibility of experience without access. This raised the question as to whether *accessibility* might be required for experience, but now we have seen

²⁵ It is true that some first-order (Dretske, 1995; Tye, 1995) and higher-order (Carruthers, 2000) representational theories of consciousness appeal to a dispositional property in order to say what makes a mental state conscious. But this is usually thought to be a *cost* of holding one of these theories (Rowlands, 2001; Jehle and Kriegel, 2006). We should expect the same in the present case.

reason to think that it is not. And so it looks as if it is possible that experience occurs without accessibility.²⁶

Be that as it may, I remind the reader that this possibility is something that I intend to take as an assumption. My present goal has just been to provide some motivation for adopting this assumption. Shortly, we look at what consequences the possibility of experience without accessibility would have for NCC research.

I will restate this section's main points. One is that accomplishing the goals of NCC research requires (or at least would benefit from) the ability to rule out brain areas from being NCC. Another is that our most promising method for doing so (i.e. the received method) makes use of negative reports. The final point is that, if experience happens in absence of access, then a negative report has the potential to mislead, since there is the possibility that some given negative report results, not from a lack of experience, but rather access.

3. The Failure of the Received Method

The received method for ruling out brain areas from being NCC, the reader will remember, recommends finding numerous instances in which the contents of experience are mismatched with the contents of a brain area, using this to infer that the brain area is not an NCC. Let us formalize the reasoning upon which the received method is based. In using the received method, one advances an argument like this:

²⁶ At this point, one might ask for an account of what an inaccessible experience is. If one has an idea as to what an experiential mental state is, it is not difficult to picture an inaccessible experience. When one has an experiential mental state, there is something it is like for that subject to have the mental state. Now picture the same experiential state occurring, but under conditions in which its content is not able to be enlisted in higher-level cognition, voluntary control of action, or reports; i.e. the kinds of use to which content that is access conscious can be put. An inaccessible experience is just this.

Then again, perhaps not everyone will find this so clear-cut. We could imagine someone (e.g. a higher-order theorist) claiming that, when she reflects upon conscious experience, it is obvious that it essentially involves some kind of access — namely, to whatever worldly properties the experience represents — and that this calls into question whether the idea of an inaccessible experience is even intelligible. Unfortunately, it is unclear whether there is any non-question-begging line of argument in favour of either the idea that access is essential to experience, or that it is not. While this issue remains a difficulty for the position for which I argue below, space dictates that it be left aside, and I take it as an assumption that the idea of an inaccessible experience is coherent. I thank an anonymous referee for pushing me on this point.

(Received Method Argument, First Pass)

1. If some brain area A is consistently involved in content mismatches, then the content of A is not conscious
2. Brain area A is consistently involved in content mismatches
- C. The content of brain area A is not conscious

So far, so good. However, whether there is a content mismatch in any instance is not something to which the scientist has direct access, this being discovered only through negative reports. We revise the argument to reflect this:

(Received Method Argument, Final Pass)

1. If subjects consistently give negative reports about content in some brain area A, then the content of A is not conscious
2. Subjects consistently give negative reports about content in brain area A
- C. The content of brain area A is not conscious

Formulated this way, the received method faces a formidable challenge. More specifically, I think that any time Premise 2 is shown to be true about some brain area A, it leads to Premise 1 being false about A. This is because Premise 2 being true also sets in motion the reasoning of this argument:

(Master Argument)

1. If subjects consistently give negative reports about content in some brain area A, then the content of A is inaccessible
2. If the content of brain area A is inaccessible, then negative reports about content in A provide no evidence for its content being not conscious
- C. If subjects consistently give negative reports about content in some brain area A, then negative reports about content in A provide no evidence for its content being not conscious

Justification for each premise of the Master Argument is provided shortly. For the time being, I would have the reader notice the upshot, which is that any time one manages to carry out the recommendation of the received method — that is, any time one is able to show that the content of some brain area A is consistently the subject of negative reports (i.e. show that Premise 2 of the Received Method Argument is true about that brain area) — one thereby also triggers the reasoning of the Master Argument. But the conclusion of the Master Argument implies that Premise 1 of the Received Method Argument is false

about brain area A, in which case the Received Method Argument no longer goes through. So, if the Master Argument is sound, then the received method fails in any instance it is used.

Now we go carefully through each premise of the Master Argument.

3.1. If subjects consistently give negative reports about content in some brain area A, then the content of A is inaccessible

Let us consider some hypothetical brain area A, whose content is consistently subject to negative reports. At this point, the received method says that these reports show that the content of A is not conscious, but let us, for now, put aside the question of whether these reports really do this. Instead, let us assess what the consistent negative reports say about the *accessibility* of content in A.

It seems that such reports are evidence for content in brain area A being inaccessible to the mechanisms of report. On the one hand, consider that the inaccessibility of the content of A would act as a good explanation for why subjects consistently give negative reports about the content of A; this is to say, it is plausible to think that subjects make such reports *because* the content of A is inaccessible to the mechanisms of report (*cf.* Brogaard, 2011). On the other hand, were the content of A accessible, it would be hard to explain why its content never shows up in reports. Given these considerations, the first premise of the Master Argument seems well supported: if subjects consistently give negative reports about content in some brain area A, then the content of A is inaccessible.

3.2. If the content of brain area A is inaccessible, then negative reports about content in A provide no evidence for its content being not conscious

In approaching the premise entitling this section, let us revisit earlier thoughts. Remember the reason for which the scientist resorts to reports when studying experience, this being the privacy of experience. This was to be overcome by having subjects report upon what they are and are not experiencing. Closer to present purposes, the idea would be that a negative report about content C could be used as evidence for C being not conscious. This is analogous to how a path through a cloud chamber is taken as evidence for an electron having the trajectory of that path, or how Koplik spots are taken as evidence for a patient having measles. But, as I explain now, when it comes to a

brain area whose content is always inaccessible (what I will refer to as an ‘inaccessible brain area’), negative reports do not provide *any* evidence for its content being not conscious (*cf.* Block, 2007).

According to how evidence is often understood (see, for example, Hacking, 1975), an observation of property P is evidence for some other property P* also obtaining if and only if P is a *reliable indicator* of P*, where reliable indication is grounded in covariation between the two properties. Koplik spots are evidence for measles, for instance, because they accompany measles (and no other disease). It is natural to extend this reasoning, saying that the *strength* of evidential support that an observation of P provides for P* is a function of *how* reliably P covaries with P*. If it were the case, for example, that P obtains only if P* also obtains, then one could be *sure*, on basis of having observed P, that P* also obtains. Likewise, if P* had *no* covariational relationship with P — more specifically, if P would obtain whether or not P* obtains, then an observation of P plausibly provides *no* evidence in favour of P* obtaining: if measles neither cause nor prevent Koplik spots, then an observation of Koplik spots would say nothing about whether or not the patient has measles.

This brings us to the issue with which we are presently concerned, which is whether a negative report about content C is evidence against C being (phenomenally) conscious, if C is in an inaccessible brain area. If what was just discussed is correct, such reports provide no evidence for this at all: since C is an inaccessible brain area, C’s being conscious could not make it any more likely to be reported as such. Given that negative reports about C have *no* covariational relationship with C being conscious, they are a (maximally) unreliable indicator of whether C is conscious or not. We must conclude that, if some content C is in an inaccessible brain area, a negative report about C offers no evidence — not even weak evidence — for C being not conscious.

3.3. *The failure of the received method*

The idea behind the received method is that if one finds that subjects consistently give negative reports about the content of a brain area, then one can infer from these reports that that brain area is not an NCC. However, what we just saw is that if one successfully establishes that a brain area’s content is consistently the subject of negative reports, one also establishes that its content is inaccessible, since such reports are best explained by such content being inaccessible. But if its content is inaccessible, then negative reports are an entirely unreliable

indicator of whether its content is conscious or not, and therefore provide *no* evidence for its content being not conscious. In sum, in the very process of following the recommendations of the received method, we lose the evidential value of those reports that we had hoped to use to show that a brain area is not an NCC. The received method undermines itself, failing in each instance it is used.

Let us now revisit brain area V1 and the dorsal stream. In 2.2, we saw what appeared to be a notable body of evidence against V1 being an NCC: V1 represented a stimulus as being composed of fine gratings, though subjects reported it as appearing uniform; V1 represented flickering colours, though subjects reported the stimulus to look of only one colour; and so on. We also saw similar evidence in the case of the dorsal visual processing stream, in the form of numerous instances in which dorsally guided visuomotor actions appeared unaffected by consciously experienced visual illusions.

However, if the content of V1 and the dorsal stream were inaccessible, this would prevent these negative reports from counting as evidence against them being NCC, for reasons just discussed. Indeed, there is neuroanatomical data suggesting that these brain areas *are* inaccessible, as both V1 and the dorsal stream have sparse or no direct connections to those areas in the prefrontal cortex thought necessary for making reports (Felleman and Van Essen, 1991). But, as seen above, there is perhaps a stronger reason for thinking this: if the dorsal stream and V1 were inaccessible brain areas, this provides a good explanation as to why subjects consistently give negative reports about the content of these brain areas. If so, then the negative reports cannot be used as evidence for these brain areas being not NCC, and the received method fails in the case of V1 and the dorsal stream.

Earlier in the paper, we observed that it seemed as if the received method provided some notable, if not conclusive, evidence against V1 and the dorsal stream being NCC. Surprisingly, it looks to provide none at all.

3.4. A conscious spinal cord? (An objection)

One might respond to the argument just given by rejecting the kind of reasoning it employs, alleging that it leads to absurd conclusions. Consider that, since neural activity in the spinal cord is inaccessible, we would expect negative reports about its activity whether it is conscious or not. But does this not mean that, if we accept the reasoning of the last section, we cannot reject the idea that the spinal cord might

be conscious? And yet we *know* that the spinal cord is not conscious, indicating something is wrong with the reasoning used — or so concludes this objection.

However, what I specifically argued for above is that *negative reports* are unable to provide evidence for an inaccessible brain area being not conscious. If we transpose the reasoning in question so that it concerns the spinal cord, we get the conclusion that negative reports cannot provide evidence against the spinal cord being conscious. So if the objection in question is to go through, it must be the case that negative reports are the only means we have for ruling out the spinal cord from being conscious.

Probably, there are other means available. It has sometimes been argued that an absence of evidence for experience serves as evidence for an absence of experience (Huxley, 1874/2002).²⁷ In the case of the spinal cord, we are hard-pressed to find anything looking like evidence for its being conscious. So, if the principle just mentioned is correct, then we *do* have evidence for the spinal cord being not conscious. But when it comes to inaccessible brain areas, there does not seem to be the same kind of lack of evidence as in the case of the spinal cord.

Consider that, though we lack anything like *direct* evidence for there sometimes being experience in inaccessible brain areas, there is of course no lack of evidence for there sometimes being experience in brain areas, full stop. This is especially true when it comes to the kind of brain area being considered in this paper, *cerebral* brain areas. The cerebral cortex is the wrinkly outermost layer of the brain, a phylogenetically recent addition where higher functions like memory, perception, and reasoning are in large part carried out. The cerebral cortex is also often thought to be where consciousness resides (at least in part). The upshot is, even if a brain area is inaccessible, if it is a cerebral brain area (like V1 or the dorsal stream), there is still *some* evidence for activity in it sometimes producing experience, because the inaccessible brain area is of the *genus* cerebral brain area, some of whose members we already know to sometimes produce experience.²⁸

²⁷ Quoth Huxley, ‘in the matter of experience... we may hold by the rule, *de non apparentibus et de non existentibus eadem est ratio*’ (1874/2002, p. 25); or ‘the reason is the same respecting things which do not appear, and those which do not exist’.

²⁸ We might not yet know *which* brain areas produce experience, but we do know that *some* brain areas must (on the assumption that the brain is responsible for conscious experience).

In sum, while there appears to be an absence of evidence for experience in the spinal cord, the same is not true in the case of inaccessible brain areas. Given this, there is no worry that the reasoning used to argue against the received method leads to our being unable to dismiss the possibility of a conscious spinal cord.

4. Consequences and Conclusions

Let us take stock. Earlier in the paper (2.1) we saw reasons why being able to show that a brain area is not an NCC is important to a science of consciousness. One reason comes from how content NCC research aims to identify not just those brain areas that are content NCC, but also those that are not. The ability to rule out brain areas from being NCC is also useful in cases where there is more than one hypothesis concerning what the neural basis of some type of conscious content is (e.g. motion experience), since it allows us to rule out some of the hypotheses. Finally, being able to rule out brain areas from being NCC is helpful toward the end of finding the neural natural kind of experience, since it can provide us with instances where a brain area both lacks neural property P and is not an NCC, a confirmation of the hypothesis that P is necessary for experience.

This raised the question as to how we might rule out brain areas from being NCC. The received method looked like the natural way to do so, but turned out not to work: to show that a brain area is not an NCC, the received method recommends that we accumulate instances where content in a brain area is the subject of negative reports; however, doing this turns out to be good only for showing that a brain area's content is inaccessible. Given this, what a science of consciousness needs is an alternative method for showing brain areas to be not NCC; ideally, one that can show this about *inaccessible* brain areas.

We have at least one at our disposal (*cf.* Block, 2007). The method starts by finding whatever neural property P is the neural natural kind of experience (see 2.1), then taking an inaccessible brain area to be an NCC if and only if it has P. And so it seems that this *indirect method* (as I will call it) provides an alternative way to rule out brain areas

from being NCC.^{29,30} Perhaps the failure of the received method does not leave a science of consciousness without recourse.³¹

The indirect method, however, has sizeable downsides. For one thing, if the indirect method is all we have at our disposal, it will probably be some time before we can actually show that some brain area is not an NCC: doing so requires knowing the neural natural kind of experience, and plausibly this is something that we will not know for some time. Another drawback of having only the indirect method is that it might prevent our ever having high confidence in an hypothesis concerning the neural natural kind of experience, *qua* the kind of neural activity both necessary and sufficient for experience. To support the hypothesis that some neural property P is necessary for experience, what we want are a number of instances in which some content's being conscious comes and goes with P. Generally speaking, the more brain areas in which this relationship is observed, the more confident we can be. But if all we have is the indirect method, then the inductive base from which the hypothesis that P is the neural natural kind of experience is generalized can consist only of those brain areas whose content is accessible. Conceivably, these could be few in number. Notably, they will probably not include V1 and the dorsal stream, brain areas that the received method promised to deliver as part of the evidence to be used in determining the neural natural kind of experience.

-
- ²⁹ The trick here is to measure neural activity in those brain areas whose content we know to be (a) accessible (because subjects sometimes give 'positive reports' about its content), and (b) not always conscious (because subjects sometimes give *negative* reports about its content), the idea being to look for some neural property P that comes and goes with the positive reports. Here, I leave aside the question of whether it will turn out to be difficult to implement the indirect method; but if it does, the failure of the received method has worse consequences for NCC research than the ones I describe just below.
- ³⁰ Of course, this method relies on being able to disentangle the neural natural kind of access from the neural natural kind of experience, but there is reason to be optimistic that this can eventually be done. Consider Shea's method (2012; see also Block, 2005; 2007): the idea here is to start by looking to see what natural kinds (property-clusters that generate reliable inductions) are associated with subjects' reports of being conscious of a stimulus. If experience happens without access, we should eventually discern two (rather than one) natural kinds. The one that is more closely associated with reports is to be considered the neural natural kind of accessibility, with the remaining one being the neural natural kind of experience.
- ³¹ I am grateful to an anonymous referee at *Journal of Consciousness Studies* for pointing out the importance and potential viability of this alternative method.

In sum, while the failure of the received method does not look fatal to a science of consciousness, it does mean that a science of consciousness might progress more slowly, and be able to draw less certain conclusions, than we might otherwise have hoped. What would be most helpful at this point is the development of new methods for ruling out brain areas from being NCC. Of course, none of this is a problem if it could definitively be shown that experience cannot happen in absence of access. Thus far, it is unclear how good the prospects for this are.³²

References

- Aglioti, S., Goodale, D. & Desouza, J. (1995) Size contrast illusions deceive the eye but not the hand, *Current Biology*, **5**, pp. 679–685.
- Baars, B. (1988) *A Cognitive Theory of Consciousness*, New York: Cambridge University Press.
- Baars, B. (2003) Introduction: Treating consciousness as a variable: The fading taboo, in Baars, B., Banks, W. & Newman, J. (eds.) *Essential Sources in the Scientific Study of Consciousness*, Cambridge, MA: MIT Press.
- Block, N. (1995) On a confusion about a function of consciousness, *Behavioral and Brain Sciences*, **18**, pp. 227–287.
- Block, N. (2001) Paradox and cross purposes in recent work on consciousness, *Cognition*, **79** (1), pp. 197–219.
- Block, N. (2002) Concepts of consciousness, in Chalmers, D. (ed.) *Philosophy of Mind: Contemporary and Classical Readings*, Oxford: Oxford University Press.
- Block, N. (2005) Two neural correlates of consciousness, *Trends in Cognitive Sciences*, **9** (2), pp. 46–52.
- Block, N. (2007) Consciousness, accessibility, and the mesh between psychology and neuroscience, *Behavioral and Brain Sciences*, **30**, pp. 481–548.
- Block, N. (2009) Comparing the major theories of consciousness, in Gazzaniga, M. (ed.) *The Cognitive Neurosciences IV*, Cambridge, MA: MIT Press.
- Block, N. (2011) Perceptual consciousness overflows cognitive access, *Trends in Cognitive Sciences*, **15** (12), pp. 567–575.
- Bridgeman, B., et al. (1979) Relation between cognitive and motor-oriented systems of visual position perception, *Journal of Experimental Psychology: Human Perception and Performance*, **5**, pp. 692–700.
- Brogaard, B. (2011) Conscious vision for action versus unconscious vision for action?, *Cognitive Science*, **35** (6), pp. 1076–1104.
- Brown, R. (2011) The myth of phenomenological overflow, *Consciousness and Cognition*, **21** (2), pp. 599–604.

³² I would like to thank Brian Fiala, Terry Horgan, Uriah Kriegel, Shaun Nichols, and two anonymous referees from the *Journal of Consciousness Studies* for discussion and comments on drafts of this paper.

- Camprodon, J.A., Zohary, E., Brodbeck, V. & Pascual-Leone, A. (2010) Two phases of V1 activity for visual recognition of natural images, *Journal of Cognitive Neuroscience*, **22** (6), pp. 1262–1269.
- Carruthers, P. (2000) *Phenomenal Consciousness: A Naturalistic Theory*, Cambridge: Cambridge University Press.
- Chalmers, D. (1995) Facing up to the problem of consciousness, *Journal of Consciousness Studies*, **2** (3), pp. 200–219.
- Chalmers, D. (1996) *The Conscious Mind*, New York: Oxford University Press.
- Chalmers, D. (1997) Availability: The cognitive basis of experience, *Behavioral and Brain Sciences*, **20**, pp. 148–149.
- Chalmers, D. (1998) On the search for the neural correlate of consciousness, in Hameroff, S., et al. (eds.) *Toward a Science of Consciousness*, Cambridge, MA: MIT Press.
- Chalmers, D. (2000) What is a neural correlate of consciousness?, in Metzinger, T. (ed.) *Neural Correlates of Consciousness: Empirical and Conceptual Questions*, Cambridge, MA: MIT Press.
- Clark, A. (2007) What reaching teaches: Consciousness, control, and the inner zombie, *British Journal for the Philosophy of Science*, **58**, pp. 563–594.
- Clark, A. (2009) Perception, action, and experience: Unraveling the golden braid, *Neuropsychologia*, **47**, pp. 1460–1468.
- Cohen, M. & Dennett, D. (2011) Consciousness cannot be separated from function, *Trends in Cognitive Sciences*, **15** (8), pp. 358–364.
- Crick, F. & Koch, C. (1990) Towards a neurobiological theory of consciousness, in *Seminars in the Neurosciences*, vol. 2, pp. 263–275, Philadelphia, PA: Saunders Scientific Publications.
- Crick, F. & Koch, C. (2003) A framework for consciousness, *Nature Neuroscience*, **6** (2), pp. 119–125.
- Damasio, A. (1999) *The Feeling of What Happens*, New York: Harcourt Inc.
- Dehaene, S. & Changeux, J. (2004) Neural mechanisms for access to consciousness, in Gazzaniga, M. (ed.) *The Cognitive Neurosciences III*, Cambridge, MA: MIT Press.
- Dehaene, S. & Naccache, L. (2001) Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework, *Cognition*, **79** (1), pp. 1–37.
- Dennett, D. (1995) The path not taken, *Behavioral and Brain Sciences*, **18**, pp. 252–253.
- Dretske, F. (1995) *Naturalizing the Mind*, Cambridge, MA: MIT Press.
- Dretske, F. (2006) Perception without awareness, in Gendler, T. & Hawthorne, J. (eds.) *Perceptual Experience*, pp. 147–180, Oxford: Clarendon Press.
- Dretske, F. (2007) What change blindness teaches us about consciousness, *Philosophical Perspectives: Philosophy of Mind*, **21**, pp. 215–230.
- Driver, J. & Vuilleumier, P. (2001) Perceptual awareness and its loss in unilateral neglect and extinction, *Cognition*, **79**, pp. 39–88.
- Ellis, R., Flanagan, J. & Lederman, S. (1999) The influence of visual illusion on grasp position, *Experimental Brain Research*, **125**, pp. 109–114.
- Feigl, H. (1953) The scientific outlook: Naturalism and humanism, in Feigl, H. & Brodbeck, M. (eds.) *Readings in the Philosophy of Science*, New York: Appleton-Century-Croft.
- Felleman, D. & Van Essen, D. (1991) Distributed hierarchical processing in the primate cerebral cortex, *Cerebral Cortex*, **1**, pp. 1–47.

- Frith, C., Perry, R. & Lumer, E. (1999) The neural correlates of conscious experience: An experimental framework, *Trends in Cognitive Sciences*, **3**, pp. 105–114.
- Gennaro, R. (2012) *The Consciousness Paradox*, Cambridge, MA: MIT Press.
- Gur, M. & Snodderly, D. (1997) A dissociation between brain activity and perception: Chromatically active cortical neurons signal chromatic activity that is not perceived, *Vision Research*, **37**, pp. 377–382.
- Hacking, I. (1975) *The Emergence of Probability*, Cambridge: Cambridge University Press.
- Hameroff, S.R. & Penrose, R. (1996) Conscious events as orchestrated space-time selections, *Journal of Consciousness Studies*, **3** (1), pp. 36–53.
- He, S. & MacLeod, D. (2001) Orientation-selective adaptation and tilt after-effect from invisible patterns, *Nature*, **411**, pp. 473–476.
- Hempel, C. (1952) *Fundamentals of Concept Formation in Empirical Science*, Chicago, IL: University of Chicago Press.
- Hohwy, J. (2009) The neural correlates of consciousness: New experimental approaches needed?, *Consciousness and Cognition*, **18** (2), pp. 428–438.
- Huxley, A. (1874/2002) On the hypothesis that animals are automata, and its history, in Chalmers, D. (ed.) *Philosophy of Mind: Contemporary and Classical Readings*, Oxford: Oxford University Press.
- Jack, A. & Roepstorff, A. (2003) Why trust the subject?, *Journal of Consciousness Studies*, **10** (9–10), pp. v–xx.
- Jacob, P. & De Vignemont, F. (2010) Spatial coordinates and phenomenology in the two-visual systems model, in Gangopadhyay, N. (ed.) *Perception, Action, and Consciousness: Sensorimotor Dynamics and Two-Visual Systems*, Oxford: Oxford University Press.
- Jacob, P. & Jeannerod, M. (2003) *Ways of Seeing: The Scope and Limits of Visual Cognition*, Oxford: Oxford University Press.
- Jehle, D. & Kriegel, U. (2006) An argument against dispositionalist HOT theory, *Philosophical Psychology*, **19** (4), pp. 463–476.
- Kanwisher, N. (2001) Neural events and perceptual awareness, *Cognition*, **79**, pp. 89–113.
- Koch, C. & Tsuchiya, N. (2007) Attention and consciousness: Two different processes, *Trends in Cognitive Science*, **11** (1), pp. 16–22.
- Kriegel, U. (2006) Consciousness: Phenomenal consciousness, access consciousness, and scientific practice, in Thagard, P. (ed.) *Handbook of Philosophy of Psychology and Cognitive Science*, Amsterdam: North-Holland.
- Kriegel, U. (2007) A cross-order integration hypothesis for the neural correlate of consciousness, *Consciousness and Cognition*, **16** (4), pp. 897–912.
- Kroliczak, G., et al. (2006) Dissociation of perception and action unmasked by the hollow-face illusion, *Brain Research*, **23**, pp. 9–16.
- Lamme, V.A. (2004) Separate neural definitions of visual consciousness and visual attention: A case for phenomenal awareness, *Neural Networks*, **17** (5), pp. 861–872.
- Lamme, V.A. (2006) Towards a true neural stance on consciousness, *Trends in Cognitive Sciences*, **10** (11), pp. 494–501.
- Lamme, V.A. (2010) How neuroscience will change our view on consciousness, *Cognitive Neuroscience*, **1** (3), pp. 204–220.

- Landman, R., Spekreijse, H. & Lamme, V.A. (2003) Large capacity storage of integrated objects before change blindness, *Vision Research*, **43** (2), pp. 149–164.
- Lau, H. (2010) Theoretical motivations for investigating the neural correlates of consciousness, *Wiley Interdisciplinary Reviews: Cognitive Science*, **2** (1), pp. 1–7.
- Leopold, D. & Logothetis, N. (1996) Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry, *Nature*, **379**, pp. 549–553.
- Logothetis, N. & Schall, J. (1989) Neuronal correlates of subjective visual perception, *Science*, **245**, pp. 761–763.
- Marcel, A. (1988) Phenomenal experience and functionalism, in Marcel, A. & Bisiach, E. (eds.) *Consciousness in Contemporary Science*, Oxford: Oxford University Press.
- Marshall, J. & Halligan, P. (1988) Blindsight and insight in visuospatial neglect, *Nature*, **336** (22/29), pp. 766–767.
- Milner, A. & Goodale, M. (1995/2006) *The Visual Brain in Action*, Oxford: Oxford University Press.
- Nagel, T. (1974) What is it like to be a bat?, *Philosophical Review*, **83**, pp. 435–450.
- Noë, A. & Thompson, E. (2004) Are there neural correlates of consciousness?, *Journal of Consciousness Studies*, **11** (1), pp. 3–28.
- Overgaard, M. & Overgaard, R. (2010) Neural correlates of contents and levels of consciousness, *Frontiers in Psychology*, **1**.
- Pascual-Leone, A. & Walsh, V. (2001) Fast backprojections from the motion to the primary visual area necessary for visual awareness, *Science*, **292** (5516), pp. 510–512.
- Peru, A., Moro, V., Avesani, R. & Aglioti, S. (1996) Overt and covert processing of left-side information in unilateral neglect investigated with chimeric drawings, *Journal of Clinical and Experimental Neuropsychology*, **18** (5), pp. 621–630.
- Popper, K. (1959) *The Logic of Scientific Discovery*, New York: Basic Books.
- Prinz, J. (2000) A neurofunctional theory of visual consciousness, *Consciousness and Cognition*, **9**, pp. 243–259.
- Prinz, J. (2007) Accessed, accessible, and inaccessible: Where to draw the phenomenal line, *Behavioral and Brain Sciences*, **30**, pp. 521–522.
- Prinz, J. (2012) *The Conscious Brain*, New York: Oxford University Press.
- Prior, E., Pargetter, R. & Jackson, F. (1982) Three theses about dispositions, *American Philosophical Quarterly*, **19** (3), pp. 251–257.
- Railton, P. (1984) Marx and the objectivity of science, in *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pp. 813–826, Baltimore, MD: Philosophy of Science Association.
- Rees, G. (2007) Neural correlates of the contents of visual awareness in humans, *Philosophical Transactions of the Royal Society B: Biological Sciences*, **362** (1481), pp. 877–886.
- Rees, G., et al. (2000) Unconscious activation of visual cortex in the damaged right-hemisphere of a parietal patient with extinction, *Brain*, **123** (8), pp. 1624–1633.
- Rees, G., et al. (2002) Neural correlates of conscious and unconscious vision in parietal extinction, *Neurocase*, **8** (5), pp. 387–393.

- Rees, G., Kreiman, G. & Koch, C. (2002) Neural correlates of consciousness in humans, *Nature Reviews: Neuroscience*, **3**, pp. 261–270.
- Rosenthal, D. (1997) A theory of consciousness, in Block, N., Flanagan, O. & Güzeldere, G. (eds.) *The Nature of Consciousness*, Cambridge, MA: MIT Press.
- Rosenthal, D. (2002) How many kinds of consciousness, *Consciousness and Cognition*, **11**, pp. 653–665.
- Rosenthal, D. (2007) Phenomenological overflow and cognitive access, *Behavioral and Brain Sciences*, **30**, pp. 522–523.
- Rowlands, M. (2001) Consciousness and higher-order thoughts, *Mind & Language*, **16** (3), pp. 290–310.
- Sellars, W. (1956) Empiricism and the philosophy of mind, *Minnesota Studies in the Philosophy of Science*, **1** (19), pp. 253–329.
- Shea, N. (2012) Methodological encounters with the phenomenal kind, *Philosophy and Phenomenological Research*, **84** (2), pp. 307–344.
- Sheinberg, D.L. & Logothetis, N.K. (1997) The role of temporal cortical areas in perceptual organization, *Proceedings of the National Academy of Sciences USA*, **94**, pp. 3408–3413.
- Silvanto, J., et al. (2005) Striate cortex (V1) activity gates awareness of motion, *Nature Neuroscience*, **8** (2), pp. 143–144.
- Sperling, G. (1960) The information available in brief visual presentations, *Psychological Monographs: General and Applied*, **74** (11), p. 1.
- Tong, F., et al. (1998) Binocular rivalry and visual awareness in human extrastriate cortex, *Neuron*, **21**, pp. 753–759.
- Tononi, G. & Koch, C. (2008) The neural correlates of consciousness, *Annals of the New York Academy of Sciences*, **1124** (1), pp. 239–261.
- Tye, M. (1995) *10 Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, Cambridge, MA: MIT Press.
- Vandenbroucke, A., et al. (2014) Seeing without knowing: Neural signatures of perceptual inference in the absence of report, *Journal of Cognitive Neuroscience*, **26** (5), pp. 955–969.
- Vosgerau, G. & Newen, A. (2008) Orthogonality of phenomenality and content, *American Philosophical Quarterly*, **45**, pp. 329–348.
- Wallhagen, M. (2007) Consciousness and action: Does cognitive science support (mild) epiphenomenalism?, *British Journal for the Philosophy of Science*, **58**, pp. 539–561.
- Wolfe, J. (1999) Inattentional amnesia, in Coltheart, V. (ed.) *Fleeting Memories*, pp. 71–94, Cambridge, MA: MIT Press.
- Zeki, S. (2003) The disunity of consciousness, *Trends in Cognitive Sciences*, **7** (5), pp. 214–218.

Paper received August 2014; revised April 2015.