

## How Things Seem to Higher-Order Thought Theorists

JACOB BERGER *Idaho State University*

*Dialogue: Canadian Philosophical Review* 56(3): 503-526.  
(Please quote the published version).

**ABSTRACT:** *According to David Rosenthal's higher-order thought (HOT) theory of consciousness, a mental state is conscious just in case one is aware of being in that state via a suitable HOT. Jesse Mulder (2016) recently objects: though HOT theory holds that conscious states are states that it seems to one that one is in, the view seems unable to explain how HOTs engender such seemings. I clarify here how HOT theory can adequately explain the relevant mental appearances, illustrating the explanatory power of HOT theory.*

**RÉSUMÉ :** *Selon la théorie de la conscience des pensées d'ordre supérieur de David Rosenthal («higher-order thought»), un état mental est conscient seulement dans les cas où l'on est conscient d'être dans cet état par l'intermédiaire d'une pensée d'ordre supérieur appropriée. Jesse Mulder (2016) a récemment émis l'objection suivante : si la théorie des pensées d'ordre supérieur maintient que les états de conscience sont des états pour lesquels il semble à la personne qui s'y trouve qu'elle s'y trouve, le point de vue ne semble pas pouvoir expliquer de quelle manière les pensées d'ordre supérieur engendrent ces impressions. Je précise ici comment la théorie des pensées d'ordre supérieur peut expliquer de manière adéquate les apparences mentales pertinentes, illustrant le pouvoir explicatif de la théorie des pensées d'ordre supérieur.*

**Key words:** HOT theory, consciousness, seemings, awareness, thought

### 1. Introduction

In order to explain any ordinary phenomenon, it is crucial first to get a handle on our commonplace conception of it so that we can begin to theorize about its underlying nature. It has, however, proven notoriously difficult to characterize *consciousness* in a commonsense way. Following Thomas Nagel, many in the contemporary philosophy of mind gloss conscious mental states as those *for which there is something that it is like to be in them*.<sup>1</sup> But as some critics have observed, despite its initial appeal, this description is quite vague and it is unclear whether it amounts to anything more than the circular claim that conscious states are conscious states.<sup>2</sup> At best, it is uncertain how such a description of consciousness provides us with sufficient material with which to theorize about its nature, a lack of clarity which perhaps explains why so many who employ this characterization of consciousness conclude that it cannot be reductively or naturalistically explained.<sup>3</sup>

But most higher-order (HO) theories of consciousness—and here my focus is on David Rosenthal's higher-order thought (HOT) theory—seek to explain consciousness

---

<sup>1</sup> Nagel 1974, p. 436.

<sup>2</sup> See, e.g., Lycan 1996, p. 77.

<sup>3</sup> See, e.g., Nagel 1974; Chalmers 1996; Levine 2001.

by building upon a more informative characterization of it.<sup>4</sup> To understand conscious states, it is reasonable to explore how they differ from their *nonconscious* counterparts. There is, of course, much folk-psychological and experimental evidence that mental states such as perceptual states and thoughts can occur without being conscious.<sup>5</sup> And HOT theory takes as its starting place the following uncontroversial observation about such states: if one is in a mental state, but in no way aware of being in it, then that state is not conscious. The logically equivalent contrapositive of that ordinary claim, Rosenthal's Transitivity Principle (TP),<sup>6</sup> is that a mental state is conscious only if one is somehow aware of oneself as being in it. The TP allows us to hone in on what consciousness is, without smuggling in any questionable theoretical assumptions about it. Building on this foundation, HOT theory hypothesizes that the relevant HO awareness of our mental lives is engendered by having thoughts of a certain kind. That is, HOT theory holds that a mental state is conscious just in case one is aware of oneself as being in that state via a suitable HOT to the effect that one is in that state. Both folk-psychological considerations and a range of experimental findings support HOT theory;<sup>7</sup> it is, I believe, the most promising naturalistic theory of consciousness currently available.<sup>8</sup>

It is thus enshrined in common sense, and the basis for HOT theory, that consciousness is a matter of mental appearance: a matter of which mental states it seems to one that one is in. The fact that consciousness is a matter of mental appearance raises what many regard as a bug—but which is actually a positive feature—of HOT theory: namely, the possibility that a HOT might make it seem that one is in a state that one is not in fact in. It is compatible with HOT theory, for example, that one might consciously be in pain or consciously see a square insofar as one's HOT might render one suitably aware of being in pain or of seeing a square, even if one is not in pain or does not see a square. This possibility of these targetless or mismatching HOTs strikes many as problematic or even incoherent; it forms the basis of perhaps the most widely cited objection to HOT theory.<sup>9</sup> But this issue has been adequately addressed elsewhere by HOT theorists who note that targetless HOTs are perfectly coherent and explanatorily advantageous insofar as many psychological phenomena can be effectively accounted for by the fact that one's conscious states (the states it suitably seems to one that one is in) can diverge from the first-order states that one actually is in.<sup>10</sup>

In the pages of this journal, Jesse Mulder mercifully does not criticize HOT theory along these lines, but rather raises a novel objection to HO views in general, which he

---

<sup>4</sup> For varieties of HO theories, see, e.g., Armstrong 1968; Lycan 1996; Carruthers 2005; Gennaro 2012. For Rosenthal's HOT theory, in particular, see, e.g., Rosenthal 1986; 2002; 2005; 2011.

<sup>5</sup> For review, see, e.g., Berger 2014a.

<sup>6</sup> See, e.g., Rosenthal 2005, p. 4.

<sup>7</sup> See, e.g., Lau and Rosenthal 2011.

<sup>8</sup> HOT theory need not be understood as naturalistic, but it offers a plausible program for naturalistic reduction. HOT theory aims to explain consciousness in terms of ordinary intentional states—thoughts—and it is reasonable to expect that such states can be themselves given a naturalistic treatment (for attempts to naturalize intentional states, see, e.g., Stich and Warfield 1994).

<sup>9</sup> See, e.g., Byrne 1997; Neander 1998; Block 2011a; Edwards and Platchias 2016.

<sup>10</sup> See, e.g., Rosenthal 2005; Weisberg 2011; Rosenthal 2011; Berger 2014b; Brown 2015; Pereplyotchik 2015; Shargel 2016. Rosenthal himself, for example, has proposed that we may regard conscious states in cases of targetless HOTs as merely *notional* states (e.g., 2005, p. 130).

calls the ‘Seeming Problem’.<sup>11</sup> Mulder argues that, even if conscious mental states are states that we seem to be in, HO views—and specifically HOT theory—do not have the resources to explain how HO states can engender the relevant mental appearances. After surveying various features that suitable HOTS are theorized to exhibit, Mulder concludes that they cannot explain how our mental lives appear in the ways constitutive of consciousness.

I explain here why Mulder’s Seeming Problem for HOT theory is just that—a seeming problem. I argue that, as I understand the theory, it already has the resources to explain how HOTS can and often do produce the relevant mental appearances. Addressing this interesting worry not only illustrates the explanatory power of HOT theory, but also clarifies how the theory works. I begin in Section 2 by laying out some basic features of HOT theory and Mulder’s new objection to it. I then clarify the nature of seemings in general in Section 3. In Sections 4 and 5, I explain how HOT theory can reply to Mulder’s concerns. I close in Section 6 by reflecting on the relationship between awareness and seemings as it relates to consciousness.

## 2. HOT Theory and the Seeming Problem

At first, the basic motivation for HOT theory, the TP, might seem to explain consciousness circularly: that conscious states are states of which we are conscious or aware. As Rosenthal has made clear, however, expressions such as ‘conscious’ or ‘aware’ are ambiguous in at least three ways.<sup>12</sup> First, we often talk of a creature’s being conscious, as in awake and responsive to stimuli and not asleep or anaesthetized. This phenomenon is what Rosenthal calls ‘creature consciousness’. Second, we often also say that a creature is conscious (or aware) *of* something—as in, I am currently conscious of my computer screen. Rosenthal calls this property of creatures ‘transitive consciousness’. Third, we often speak of a mental state’s being conscious, as opposed to nonconscious as in cases of subliminal perception observed in persons with blindsight or studies of masked priming in healthy individuals.<sup>13</sup> This notion of consciousness, the target of explanation of HOT theory, is what Rosenthal calls ‘(intransitive) state consciousness’. The TP is thus not circular insofar as it seeks to explain a state’s being state conscious in terms of one’s transitive awareness of that state.<sup>14</sup> HOT theory seeks to implement the TP by positing that suitable HOTS are the states in virtue of which people are aware of being in their conscious mental states.

Herein nonetheless lies a problem for HO views such as HOT theory, or so Mulder argues. Mulder asks:

But what exactly does the *transitive* notion involved amount to? HO theorists should be concerned with this question, given that this transitive notion, ‘awareness of,’ is their main explanans for the explanandum of intransitive consciousness (state consciousness). In particular, HO theorists should take

---

<sup>11</sup> Mulder 2016.

<sup>12</sup> See, e.g., Rosenthal 2002, pp. 406-407.

<sup>13</sup> See respectively, e.g., Weiskrantz 1997; Kouider and Dehaene 2007.

<sup>14</sup> For clarity going forward I use ‘conscious’ only in the intransitive (state) sense and ‘aware’ and related expressions only in the transitive sense, unless otherwise specified.

note of platitudes relating to this transitive notion (just like they take note, and make thankful use of, the platitude recorded by the Transitivity Principle).<sup>15</sup>

And, Mulder continues, one such platitude is the following, which he calls the ‘Seeming Platitude’:

Seeming Platitude: What a subject is *aware-of* is how things *seem* to the subject.<sup>16</sup>

The Seeming Platitude links a state of awareness to how things seem to one. Folk psychology holds that if one is aware of X, then it seems to one that X is, in some way, present. If I am currently conscious or aware of my computer screen, then it seems to me that my computer screen is in my presence. As Mulder puts it, “a representational state amounts to awareness-of only if it makes its contents available to the subject, such that things seem to the subject to be the way those contents say they are.”<sup>17</sup> Suitable states of awareness are (or at least entail), to use a somewhat inelegant expression, *seemings*.

HOT theory respects the Seeming Platitude. A suitable HOT that I see a square not only makes me aware of that visual state, but makes it seem to me that I see a square. And Rosenthal himself repeatedly emphasizes that conscious states are states that it suitably seems to one that one is in; he writes:

A state’s being conscious is a matter of mental appearance—of how one’s mental life appears to one. If somebody is in a mental state but doesn’t seem subjectively to be in that state, the state is not conscious.... Higher-order theories exploit this, arguing that a state is conscious only if one is subjectively aware of oneself as being in that state.<sup>18</sup>

Per the TP, for example, the *conscious visual appearance* that there is a square simply is the *suitable mental appearance* that one sees a square.

Mulder argues, however, that HOTs cannot make things appear in this way; he dubs this the ‘Seeming Problem’:

Seeming Problem: HO theories rely on a representational conception of *awareness-of* that does not respect the Seeming Platitude.<sup>19</sup>

In a nutshell, Mulder’s objection is that HOTs cannot implement the TP. HOT theory holds that conscious states are states that it seems to one that one is in via suitable HOTs, but, he argues, HOTs are unable to make our mental lives appear to us in the way that putatively engenders consciousness. And, if that were the case, Mulder’s

---

<sup>15</sup> Mulder 2016, p. 451.

<sup>16</sup> Mulder 2016, p. 451.

<sup>17</sup> Mulder 2016, p. 452.

<sup>18</sup> Rosenthal 2011, p. 431.

<sup>19</sup> Mulder 2016, p. 452.

objection would be a totally fair criticism of the view. Any theory ought to be able to cash its theoretical cheques.

At the outset, I want to mention what I think Mulder is *not* arguing because there is an ambiguity in the Seeming Problem. As we shall see, many theorists assume that seemings are necessarily conscious—that if things seem a way to one, they *ipso facto* consciously seem that way.<sup>20</sup> Thus one might interpret the Seeming Problem as concerning how suitable HOTS could give rise to conscious seemings. That is, a critic of the theory might be sceptical that we should identify the mental appearances engendered by representational HOTS with the appearances of consciousness. After all, it is commonly assumed in the contemporary philosophy of mind that consciousness is an actually existing monadic property of first-order conscious states, determinates of which are often referred to as ‘qualia’ or ‘phenomenal’ characters/qualities.<sup>21</sup> On this kind of view, a visual experience of a square makes it seem that there is a square because it exhibits an appropriate phenomenal property. If that were so, then HOT theory would doubtless seem like a nonstarter. But such views are unmotivated—and to construe the Seeming Problem along these lines amounts to begging the question against HOT theory and its fundamental motivation, the TP.<sup>22</sup> According to the TP, all there is to a state’s being conscious or having a conscious character is for one to suitably seem to be in it.

Moreover, Rosenthal is clear that seemings need not be conscious. Neither the TP nor HOT theory entail that conscious states are states that it *consciously* seems that one is in; rather, they are simply states that we seem to be in. To assume otherwise implausibly complicates consciousness—in an ordinary conscious experience of a square, I need not be consciously aware of myself, much less that I see a square. Indeed, if that were the case, then HOT theory would face a simplistic objection often raised against it: if HOTS are what make mental states conscious, but those HOTS are conscious, then the HOTS themselves would need to be themselves targeted by yet HO HOTS, resulting in a vicious regress. In reply, Rosenthal proposes that suitable HOTS are typically nonconscious states, which are theorized to make states *state* conscious.<sup>23</sup> According to HOT theory, what it is for a conscious state to make things consciously seem a way to one is for one’s HOT to make it nonconsciously seem to one that one is in that state. HOT theory does not deny that HOTS can themselves be conscious too—when one has suitable third-order HOTS about them. But this, according to Rosenthal, is what it is to introspect: to be not just aware, but also consciously aware, of being in some state. The theory is thereby explicitly committed to the existence of nonconscious appearances, which some theorists might find questionable—and I return to this issue to defend this feature of the view in Section 5.

In any case, Mulder grants for the sake of argument that consciousness is a matter of mental appearance and that seemings need not be conscious.<sup>24</sup> His objection, I take it, is that HOTS are not the kinds of states that can make our mental lives appear

---

<sup>20</sup> For review, see, e.g., Tucker 2013.

<sup>21</sup> See, e.g., Nagel 1974; Chalmers 1996; Levine 2001.

<sup>22</sup> For discussion, see, e.g., Rosenthal 2011; Berger 2014b.

<sup>23</sup> For discussion of this objection and reply, see, e.g., Rosenthal 2005, p. 9.

<sup>24</sup> See respectively, e.g., Mulder 2016, pp. 455-456 and p. 452.

in the ways constitutive of consciousness—and going forward I interpret the Seeming Problem in this manner.

Mulder observes that there are broadly three ways to reply to this problem. First, one could pursue what he calls the ‘Ostrich Strategy’, arguing that any kind of representation makes things seem a way to one. Second, one could endorse what he calls the ‘Repeat Strategy’, on which a HOT is necessary not only to make one aware of oneself as being in a state, but also to make a distinct HOT a state that engenders mental appearances. Third, one could endorse what he calls the ‘Alternative Strategy’, on which the fact that HOTs can make things seem a way is explained not simply by the fact that HOTs are HO, but by appeal to other properties of the HOTs themselves.

The Repeat Strategy plainly faces a vicious regress, as Mulder indicates. According to that proposal, HOT theory might explain how HOTs engender conscious seemings by repeating the appeal to HO states of awareness. If I have a HOT that I see a square, then that HOT makes me aware of myself as seeing a square. On the Repeat proposal, my HOT makes it seem to me that I see a square only if I have yet a third-order HOT to the effect that I have that HOT. But this view is unworkable for the same reasons that HOT theory is often mistakenly thought to face a regress. And, as noted, this is not how Rosenthal develops the theory.

Unlike the Repeat Strategy, the Ostrich Strategy allows that HOTs can make things seem ways to subjects. Mulder correctly observes, however, that this strategy implausibly trivializes the notion of ‘seeming’. I agree with Mulder that many information-carrying states of a person—and arguably many kinds of *mental* representations—do not make anything seem a way to one. Neither a state of my retina nor my hope that others will endorse HOT theory make, in any straightforward way, anything seem any way to me, even if my retinal state carries information about edges and my hope represents that state of affairs.

But it is clear that Rosenthal does not endorse this approach either. Rosenthal’s methodology for explaining consciousness in general is to characterize it in an ordinary way—in terms of our awareness of states—and then to theorize about what kind of state could implement that awareness. To do so, he mobilizes the resources of folk psychology. And, in many places, Rosenthal has emphasized that not any kind of mental representation can make one aware of things. For example, he writes:

There are two ways we are ordinarily conscious of things: by sensing them and by thinking about them as being present. I may, for example, be conscious of somebody by seeing that person or simply by thinking about the person as being nearby.<sup>25</sup>

Folk psychology holds, for example, that I can be aware of the colour red by visually sensing it or by having a thought about it. This is not to say that just any kind of thought will make one aware of something, though. Rosenthal notes that common sense has it that, for example, “[h]aving a thought about Napoleon doesn’t make one conscious of him unless one thinks about him as being present.”<sup>26</sup> So Rosenthal proposes restrictions on what kinds of mental representations can make one aware of things in

---

<sup>25</sup> Rosenthal 2005, p. 182.

<sup>26</sup> Rosenthal 2005, p. 182, fn. 16.

the way necessary to generate consciousness. That is to say, HOT theory pursues a variety of the Alternative Strategy.

The fact that there are two commonsense ways by which we can become aware of things—by thought or by sensation—explains why HO theories fall into two general categories: HOT theories, on which the relevant awareness is generated via thoughts, and higher-order-perception (HOP) theories, on which that awareness is generated via quasi-perceptual/sensory states.<sup>27</sup> And Rosenthal has persuasively argued that the HOP theories are problematic—for example, that perceiving always involves sensory qualities, but it is unclear what the sensory qualities of our states of HOP might be.<sup>28</sup> Rosenthal thus identifies HOTs as the best candidate for generating consciousness.

The idea that thoughts of a certain kind could not only make one aware of things, but also make things seem a way to one, is therefore taken to be engrained in common sense. If that assumption were false, then HOT theory could not get off the ground. But, even if that assumption is true, there remains the question of whether we can successfully theorize about which thoughts could make our mental lives appear in the way that is distinctive of consciousness. Mulder alleges that this project cannot succeed. If Mulder were right and it turned out that thoughts cannot engender seemings, or the right sort of seemings, then folk psychology would arguably lack the resources to explain consciousness. In that case, we would need to move beyond standard HOT and HOP theories; perhaps we would need to posit a new psychological mechanism that implements the TP.<sup>29</sup> I argue shortly, however, that HOT theory already enjoys the material needed to explain how HOTs can and often do engender the relevant seemings. Before evaluating Mulder's arguments, however, it will first be helpful to clarify what seemings in general might be.

### 3. On Seemings and HO Awareness

Part of the problem with assessing Mulder's objection is that he himself is not particularly explicit about how he conceives of seemings in general. I suspect that this is in part because he aims to remain neutral regarding the relationship between seemings and states of awareness; he writes:

Notice, by the way, that I do not commit myself to a precise statement as to the relation between seeming and awareness-of. Perhaps they are just one and the same phenomenon; perhaps seeming is a feature of all and only states that provide awareness of their contents (or *vice versa*); perhaps there can even be seemings which do not amount to any awareness-of (although I don't think that this last option makes sense, pre-theoretically). My argument only requires the claim that the awareness-of on which the HO approach rests implies seeming....<sup>30</sup>

---

<sup>27</sup> For the latter kind of view, see, e.g., Lycan 1996.

<sup>28</sup> See, e.g., Rosenthal 2005, pp. 182-183.

<sup>29</sup> Lau and Rosenthal (2011, p. 366) discuss, for example, what they call the 'higher-order statistical inference view' wherein higher-order decision processes assess statistical features of first-order representations. Such decision processes are not included in folk psychology; they are instead posits of theoretical psychology.

<sup>30</sup> Mulder 2016, pp. 452-453.

Pace Mulder, folk psychology does appear to include cases of appearances that do not engender awareness of things. To adapt Rosenthal's example above, even if it may seem to me that Napoleon was ambitious, it is far from obvious that I am thereby aware of Napoleon in any ordinary sense of 'aware'. But the possibility of these sorts of seemings is not a concern for Mulder, as his worry regards whether certain states of awareness—HOTs—can engender seemings. More to the point, then, I am doubtful that folk psychology includes any cases of states of awareness that do not engender seemings. For present purposes, however, I argue here only that the sort of awareness that HOTs engender entail seemings. As I understand HOT theory, it holds that suitable HOTs simply are seemings, though it is compatible with there being other sorts of states (such as sensory states) that are or entail seemings too.

But what is a seeming in general? It is well known that 'seems' and related appearance expressions are ambiguous in various ways—and it is clear that at a minimum HOT theory is operating with the intentional or representational use of 'seems', on which it seems to one that X is present just in case one represents X's presence in some appropriate way.<sup>31</sup>

Although common sense holds that thoughts of a certain kind can make things appear ways, some theorists have recently urged that seemings are *sui-generis* experiences, quite unlike sensory states or the other ordinary propositional attitudes found in folk psychology.<sup>32</sup> Chris Tucker describes this kind of view thus:

But what is a seeming? A seeming that P is neither a belief that P nor an inclination to believe P... it is a certain kind of experience with propositional content. What distinguishes seemings from other experiences is their peculiar phenomenal character. Huemer... refers to this character as "forcefulness," but I prefer the name "assertiveness".... The phenomenology of a seeming makes it feel as though the seeming is "recommending" its propositional content as true or "assuring" us of the content's truth.<sup>33</sup>

If such a view were correct, then HOT theory would indeed fall prey to the Seeming Problem: HOTs, which are just ordinary propositional thoughts, would not be able to make things seem a way to one. One might instead recommend a *higher-order seeming* theory, wherein the states of awareness that engender consciousness are HO *sui-generis* seemings of the sort posited by Tucker. But such a view would face the regress problem that is often thought to undermine Rosenthal's HOT theory: if seemings are necessarily conscious, but made conscious by seemings, then the view is unworkable.<sup>34</sup>

Such a *sui-generis* account of seemings is in any case far from uncontroversial. Jake Quilty-Dunn has recently argued, quite persuasively to my mind, against this view and in favour of the idea that seemings often are ordinary occurrent beliefs—that is,

---

<sup>31</sup> For a classic discussion of the ambiguity of appearance vocabulary, see, e.g., Chisholm 1957, chapter 4; on the representational sense of 'seems', see, e.g., Shoemaker 1982, p. 365.

<sup>32</sup> See, e.g., Tucker 2013.

<sup>33</sup> Tucker 2010, p. 530.

<sup>34</sup> One might attempt to save the higher-order-seeming theory by arguing that seemings are *sui-generis* propositional attitudes that need not be conscious. But this revision of HO theories is unnecessary because Rosenthal's version of HOT theory is unproblematic.

thoughts.<sup>35</sup> The main evidence for the view that seemings are not thoughts is the observation that we can and often do have seemings that contradict our beliefs. Tucker reports that the view that seemings can be beliefs “is not widely endorsed these days because it faces the Problem of Known Illusions.... It might seem to me that the half-submerged stick is bent, even though I don’t believe that it is.”<sup>36</sup> But, as Quilty-Dunn argues, it is quite plausible that we often have many contradictory beliefs/thoughts, as in cases of implicit racist biases that contradict one’s explicit beliefs in racial equality.<sup>37</sup> In the submerged-stick example, perhaps it seems to one that the stick is bent insofar as one has the perceptual thought—that is, a thought formed quickly or automatically on the basis of sensation—that the stick is bent, though one more reflectively believes that it is not bent. Since the latter state is held more strongly than the former, the former is, to borrow a phrase of David Armstrong’s, “held in check” by the latter insofar as it may rarely manifest in behaviour.<sup>38</sup> Thus it remains open that seemings are simply thoughts of a certain kind. And, even if Tucker is correct that seemings have a distinctive assertoric character when conscious, it is plausible that this is often simply the character of ordinary assertoric thoughts, which can occur without being conscious too.

Moreover, even if *some* seemings are *sui-generis* experiences, it remains open that some seemings are ordinary assertoric thoughts as well. But I need not evaluate this alternative conception of seemings in detail here because to assume that no thoughts are seemings would be question begging in this context—and in any case it is not the approach that Mulder pursues.

Mentioning *sui-generis* accounts of seemings is nonetheless instructive because it reveals a reasonably neutral way to understand seemings in general. To accommodate what I take to be the central insight of the *sui-generis* view, it appears that folk psychology conceives of seemings as nonfactive mental states that have an assertoric character: they present things as being certain ways to subjects. Characterized this way, it is at least open at this stage that HOTs can make things seem a way. Mulder denies this. In the next section, I evaluate his arguments, arguing that he provides no compelling reason to deny that HOT theory can pursue an effective version of the Alternative Strategy.

#### **4. The Alternative Strategy Revisited**

Even though Rosenthal maintains that there are restrictions on the nature of suitable HOTs—that is, the kinds of HOTs hypothesized to engender consciousness—Mulder argues that none of the features of HOTs could explain how they can make things seem a way to us. As Mulder correctly observes, HOTs are typically “claimed to be occurrent, non-inferential, assertoric and conceptual, and they are thought to involve a first-person reference that is linked to a mental state at the present moment.”<sup>39</sup> But he maintains that these features will not do the trick. He argues, for example, that there are many

---

<sup>35</sup> Quilty-Dunn 2015. One reason that one might doubt that beliefs are seemings is that seemings are occurrent states, whereas beliefs are often assumed to be dispositional. But, as I note shortly, HOTs are theorized to be occurrent too.

<sup>36</sup> Tucker 2013, p. 4.

<sup>37</sup> Quilty-Dunn 2015, p. 567. For other recent defences of the idea that we can have contradictory beliefs, and especially in the case of implicit bias, see, e.g., Egan 2008; Mandelbaum 2016.

<sup>38</sup> Armstrong 1968, p. 154.

<sup>39</sup> Mulder 2016, p. 457.

cases of occurrent states that do not manifest such appearances, as in states of the early visual system of which one is unaware.<sup>40</sup>

Mulder is alive to the fact that these conditions on suitable HOTs are theorized not to be individually sufficient, but to be jointly necessary and sufficient for engendering consciousness.<sup>41</sup> But he nonetheless seems not to appreciate the full nature of these states. What Mulder specifically leaves out is consideration of the relevant contents: suitable HOTs must not only make “first-person reference that is linked to a mental state at the present moment,” as he claims, but must occurrently and conceptually assert that one is *currently in a state*: they have contents such as <I see that there is a square> or <I think that *p*>.

Why? The reasons are straightforward. After all, not any mental representation of a mental state will make one aware of it—much less make that state *state* conscious. My having a hope that you have a future mental state is a HOT insofar as it is a thought about a mental state, but it clearly will not make your future state *state* conscious and arguably does not even render me aware of that state. It is therefore a requirement on the contents of suitable HOTs that they not only represent a state, but also represent the state as somehow present. We recall that folk psychology holds that the only kind of thoughts that can make one aware of things are thoughts of things as present. But one cannot simply have a thought about a state as present, lest my thoughts about your current mental states be able to make those states conscious. So another necessary condition on suitable HOTs is that they must represent oneself as being in mental state. But even this condition alone is insufficient—for I can imagine that I am in pain without rendering myself aware of that pain, much less rendering the pain conscious.<sup>42</sup> Thus in order to be not only necessary but also sufficient for consciousness, suitable HOTs must *mentally assert that one is currently in a state*.

These restrictions on the nature of suitable HOTs explain why many purported objections to HOT theory are not really problems for the view. They explain, for example, why what is often called the ‘problem of the rock’—the worry that if HOT theory were true, then one could have a HOT about a rock, but it is absurd to think that a rock could be thereby conscious—is simply a misunderstanding of the theory.<sup>43</sup> One cannot have the assertoric thought that one is currently in a rock.

It is for these reasons as well that a suitable HOT that one is in a state makes it seem to one that one is in that state. Just as common sense has it that the thought that one’s computer is present not only makes one aware of one’s computer but also makes it seem that one’s computer is present by asserting to the subject that fact, the HOT that I see a square makes me aware of myself as seeing a square *and* makes it seem to me that I see a square.<sup>44</sup> That HOTs can make things seem a way is just a special case of how some thoughts in general can generate appearances—they just concern features of one’s mind, rather than features of the world. Moreover, suitable HOTs have the assertoric force distinctive of seemings in general because they are thoughts that one is

---

<sup>40</sup> Mulder 2016, pp. 458–459.

<sup>41</sup> See, e.g., Mulder 2016, p. 459.

<sup>42</sup> I thank an anonymous referee for this journal for this clarification and example.

<sup>43</sup> See Goldman 1993, p. 366; cf. Gennaro 2012, Section 4.3.

<sup>44</sup> Though Mulder speculates that HOT theorists might consider self-representation as a potential additional condition on seemings (2016, pp. 459–460), this is not to say that all seemings must be self-representational, as my example illustrates.

currently in a mental state. Put another way, they make present to one's mind that one's mind is some way.

Mulder, however, purports to have counterexamples—examples of allegedly suitable HOTS that do not manifest the relevant mental appearances. For example, he writes, “Obviously, one may again point to Freud-inspired counterexamples: ‘suppressed’ thoughts to the effect that one is in certain mental states (anger, say) fit the bill but are supposed to be exactly *not* seeming-providing.”<sup>45</sup> And Mulder does seem correct that simply suppressing the thought that one is angry need not put one in a conscious state of anger.

But it is not entirely clear how to interpret Mulder's case. If, for example, prior to suppressing the HOT that one is angry, one formed it on the basis of testimony from a psychotherapist, then the HOT will fail to meet an additional condition on suitable HOTS. As Rosenthal often observes, suitable HOTS must not only assert that one is currently in a state, but also make one aware of that state in a way that does not depend on one's awareness of any inferences or observations. After all, simply being told that one is in a state need not make that state conscious. This requirement is in place because folk psychology has it that our awareness of our conscious states always at least feels direct. But this way of understanding Mulder's Freudian case would violate this requirement, whether or not the HOT is suppressed.<sup>46</sup>

Perhaps Mulder assumes that suppressing the HOT would render one completely unaware of any mediation—perhaps suppressing it includes suppressing one's memory of being told about the anger. In that way, Mulder's argument echoes Georges Rey's objection wherein one has a suitable HOT that one is in a state, but then represses that conscious state.<sup>47</sup> Rey's example is arguably stronger insofar as it stipulates that one has a suitable HOT, but then purports to imagine that one fails to be in the target conscious state. Mulder's case involves a suppressed HOT, which arguably does not meet the relevant conditions. In any case, as Rosenthal argues in reply to Rey's objection, cases of repression may not involve a loss of consciousness of the repressed state; instead, they may simply involve a lack of *attention* to those conscious states.<sup>48</sup> It thus remains open that, even if one suppresses a HOT, one remains aware not only of that HOT, but also of the reasons why one has it (e.g., one would likely remember, though perhaps not attend to, the fact that one's therapist made one aware of the anger).

Mulder nonetheless offers a thought experiment to put pressure on the assumption that suitable HOTS must not arise on the basis of inference or observation.

---

<sup>45</sup> Mulder 2016, p. 459.

<sup>46</sup> Since HOT theory does not hold that every HOT engenders consciousness, it is open that, just as some ordinary thoughts of things as present can make things seem a way, Mulder's suppressed HOT could make it seem to one that one is angry, without making one consciously angry. That is to say, it is unclear whether Mulder aims to offer a counterexample to the claim that HOTS of any kind can make things seem a way or to the claim that certain HOTS can give rise to consciousness. But, even if Mulder's worry is the latter problem, it is far from clear that Mulder's example challenges either claim.

<sup>47</sup> Rey 2000, p. 227. I thank an anonymous referee for this journal for drawing my attention to the parallel between these arguments. While Rey discusses *repression* and Mulder discusses *suppression*, there is some psychoanalytic debate about whether these are distinct processes—and the expressions are often used interchangeably (see, e.g., Akhtar 2009, p. 278). For ease, I assume here that they are the same process.

<sup>48</sup> Rosenthal 2000, p. 234.

He imagines that, after being told by my therapist that I am angry, thereby forming the HOT that I am angry on the basis of an inference of which I am aware, a clever neuroscientist somehow erases from my mind any trace of the inference so that it now seems that I am aware of myself as being angry not on the basis of inference. Here, Mulder observes, my “resulting HO state would be subjectively indistinguishable from a proper one,”<sup>49</sup> and so it is unclear why an actual lack of mediation would be necessary for HOTs to engender consciousness.

To clarify, however, Rosenthal’s view is not that suitable HOTs cannot actually arise on the basis of inference or observation, but that they must not *seem* to arise on that basis: suitable HOTs must be, as Rosenthal puts it, ‘subjectively unmediated’, even if they are in fact mediated.<sup>50</sup> Whether a HOT engenders consciousness thus does not depend on extrinsic factors of which the subject may be unaware including, for example, whether one’s HOT is in fact mediated. Consciousness depends only on factors of which the subject is aware: whether or not it seems to one that one’s HOT depends on inference or observation.

Mulder acknowledges that HOT theory might hold this, but, he argues, this condition is circular; he writes that this kind of “move relies on how things seem to the subject—and it is clear that the requirement of *seeming* non-inferentiality is of no help for explaining seeming.”<sup>51</sup> But this is not so. To say that a HOT arises in a subjectively unmediated way is just to say that one has the thought, whether or not one mentally represents (consciously or otherwise) any of the facts that lead to why one has that thought. So, for example, my psychoanalyst-inspired HOT that I am angry is not subjectively unmediated insofar as I have a mental representation of the cause of that thought—namely, a thought about my therapist’s testimony. But, if I have no such representations, either as a matter of the normal psychological course of things or because of neuroscientific intervention, then my HOT would be subjectively unmediated.

Mulder might urge that in his neuroscientific-intervention example one would still not be consciously angry. If so, his argument would be similar to Ned Block’s recent objection to the requirement that suitable HOTs must arise in seemingly non-inferential ways. Block imagines a situation in which one simply forgets the source of one’s HOT, but such a HOT does not plausibly make its target conscious.<sup>52</sup> But Rosenthal has replied that, without adding details to the case, merely forgetting the source of a HOT would not be sufficient to render it seemingly non-inferential. As Rosenthal puts it, it “would still seem that one’s awareness rested on some such mediating factors, despite one’s inability to recall which.”<sup>53</sup> This is particularly the case in Block’s example, wherein it is hard to imagine what it would be like to form a HOT on the basis of inference, forget those inferences, and then it not seem in any way that one’s awareness is based on *something*.

But, if Mulder’s imagined neuroscientific intervention really makes the HOT subjectively unmediated, then there is no non-question-begging reason to suppose that

---

<sup>49</sup> Mulder 2016, p. 458.

<sup>50</sup> See, e.g., Rosenthal 2005, pp. 183-184.

<sup>51</sup> Mulder 2016, p. 458.

<sup>52</sup> Block 2011b, p. 446.

<sup>53</sup> Rosenthal 2012, p. 1429.

it would not engender consciousness. Such a case would be, as Mulder himself acknowledges, subjectively indistinguishable from any normal instance of being consciously angry insofar as it one would be aware of being angry in a way that feels direct. Whatever residual doubt one may have about the state's being conscious arguably stems from one's failing to truly imagine what it would be like for *all* subjective mediation to be removed.

We recall, however, that many theorists assume that all seemings are conscious. And it is surely correct that in Mulder's case it would not *consciously* seem to one that one is angry. On Rosenthal's view, it would only seem that way were one to introspect, thereby rendering one's second-order HOT conscious. I therefore suspect that some of what may be driving some of the scepticism that HOTs can make things seem a way is tacit resistance to the idea that *nonconscious* states can engender appearances. It is, however, a central plank of HOT theory that nonconscious thoughts of a suitable kind can not only make one aware of things, but also make things seem a way to one. While this idea may strike some as outlandish, it is not so hard to believe once properly understood.

## 5. Nonconscious Seemings

I hazard that many theorists assume that seemings must be conscious because paradigmatic or obvious cases of seemings are conscious. After all, all of the cases of seemings of which we are aware from the first-person perspective are conscious. But similar considerations once motivated thinkers to assume that other sorts of mental activity—and perhaps all mentality—must be conscious.<sup>54</sup> Once we allow that other kinds of mental states such as perceptual states—and, in particular, intentional states—can occur nonconsciously, it becomes increasingly unclear why seemings cannot occur without being conscious too.

Much experimental evidence supports the view that seemings can occur nonconsciously. Consider, for example, a standard study involving masked priming. In such experiments, participants are presented with stimuli, which are then quickly followed by what are known as 'masks' that prevent the stimuli from being consciously perceived (evidenced by the facts that participants report not seeing the stimuli and are at chance in forced-choice tasks for selecting them from among options). But it is natural to claim that such stimuli are subliminally perceived insofar as they *prime* downstream behaviour in complex ways. For example, participants in such studies might be faster at identifying a square as a square if first primed by a square than if primed by a circle. It is natural to claim that a participant sees the square and is aware of the square. And, if that is the case, then, perhaps a bit more controversially, it is also reasonable to say that it seems to the person that there is a square. One might doubt that the state makes things seem a way to the subject insofar as the subject is unaware that the state makes things seem that way to her. But it would be mysterious to explain why the participant primed by a square would act in square-type ways if the state did not assert, and thus make present to her at least in some way, that there is a square.<sup>55</sup>

---

<sup>54</sup> For discussion, see, e.g., Berger 2014a.

<sup>55</sup> Proponents of the view that seemings are *sui-generis* experiences naturally interpret such evidence otherwise. Tucker, for example, explains blindsight phenomena—wherein people with severely damaged visual cortices report not seeing anything within certain visual regions but are nonetheless remarkably

Of course, the participant does not *consciously* see the square. It does not consciously seem to her that there is a square. The state of seeming/awareness is itself nonconscious insofar as it makes things seem a way to her, but it does not seem to her that she is in it. This way of describing things is only paradoxical if one fails to distinguish the (perceptual) state of *transitive* consciousness in virtue of which it seems to the person that there is a square from the fact that this state is (intransitively) *state* nonconscious. State conscious states, by contrast, are states that it seems to one that one is in; in consciousness, things *seem to seem* a way to one. Rosenthal explicitly defends this distinction; he writes:

[W]hatever initial air of oddity there is to this idea, there is good reason to sustain the distinction between how things seem and how they seem to seem. The content of one's sensory states defines how things seem to one, even when the sensory states aren't conscious.<sup>56</sup>

HOT theory capitalizes on this insight. Just as nonconscious perceptual or sensory states can make things seem a way to one, without seeming to seem a way to one, so too can HOTs nonconsciously make things seem a way to one. And, again, according to the TP, all there is for it to consciously seem that there is a square is for it to seem to one that there seems to be a square. It only consciously seems to one that there seems to be a square when one is aware not of one's first-order states, but of one's HO states of awareness. But that is introspection, not ordinary consciousness. So long as we keep these folk-psychological constructs separate, the idea that nonconscious HOTs could make one's mind seem a certain way to one (without seeming to seem that way) is perfectly intelligible.

One might nonetheless reject the idea that nonconscious states can make things seem a way because one doubts that nonconscious states are genuinely mental—that describing nonconscious states as 'thoughts' or 'perceptual states' is a kind of loose talk.<sup>57</sup> On this view, such states are, to use Daniel Dennett's expression, *subpersonal* insofar as they are attributable only to individuals' subsystems and not to whole persons.<sup>58</sup> And, if nonconscious states are not mental, then they arguably cannot be seemings. But this view is questionable insofar as the relevant nonconscious states—

---

good at distinguishing stimuli in those fields—as involving a loss of visual imagery in those blind areas. But, he argues, the fact that individuals with blindsight are able to make what they regard as guesses about those stimuli reveal that the mechanisms for producing conscious, though arguably weak, seemings are relatively intact (Tucker 2010, pp. 530-531). Tucker's explanation is somewhat mysterious, however: if participants' visual imagery is missing, then what explains how the seeming mechanisms function normally? It is plausible that visual states are what typically cause, or simply are, the relevant seemings. This is why the HOT-theoretic explanation of blindsight is preferable: a person with blindsight might nonconsciously see a square, and on that basis it nonconsciously seems to her that there is a square; it is just that it does not seem via a suitable HOT that it seems to her that there is a square. The case is much clearer in studies of masked priming wherein participants do not even have the weak impression that they can guess what masked stimuli are and are at chance for selecting them in forced-choice tasks.

<sup>56</sup> Rosenthal 2005, p. 243.

<sup>57</sup> For a debate about the existence of nonconscious mental states (in particular, nonconscious perceptual states), see, e.g., Phillips and Block 2016.

<sup>58</sup> Dennett 1969, p. 93.

including HOTs—play virtually the same functional roles as their conscious counterparts except that we are not subjectively aware of them. Moreover, to deny there are genuinely nonconscious mental states is to deny a central motivation for HOT theory and so does not form the kind of independent objection to the theory that Mulder aims to raise.

Similar considerations undermine Mulder’s concern that, if HOT theory were true, the relevant sorts of seemings might occur during sleep. Mulder writes:

If awareness-of (‘transitive consciousness’) is consistent with being asleep, as Rosenthal asserts, it would follow that, while immersed in a deep, dreamless sleep, things seem to one to be various ways. Although one might use the term ‘seeming’ in that way, that is not the way it is used in the Seeming Platitude—after all, that platitude allows for a distinction between representations that do and those that do not enjoy such seeming, whereas Rosenthal’s notion of awareness-of does not. Thus read, Rosenthal is, in effect, pursuing the Ostrich Strategy, which leads to an objectionable trivialization of seeming.<sup>59</sup>

But Rosenthal does not adopt the Ostrich Strategy. While there are constraints on which representations constitute states of transitive awareness (and hence seemings), they do not rule out the possibilities that those representations can occur without being state conscious or in creatures that are not creature conscious. Mulder’s remarks are curious since he is sensitive to the tripartite distinction between state, transitive, and creature consciousness. On the assumption that these are distinct uses of ‘consciousness’, it is open that they could come apart. And some theorists, including Rosenthal, have argued that there may be cases of state conscious states in the absence of creature consciousness, as in instances of lucid or vivid dreaming.<sup>60</sup> Mulder does, however, assume that for things to seem a certain way, the individual must be creature conscious. After introducing the Seeming Platitude, he writes that the fact that things seem a way to a subject “involves, minimally, that the subject is awake, or at least not in a deep coma or immersed in a dreamless sleep (in other words, that the subject is ‘creature conscious,’ as it is often put).”<sup>61</sup> But even this minimal assumption is questionable. Presuming transitive consciousness is distinct from creature consciousness, there is no independent reason to add this stipulation. It is arguable that we can become aware of things—it can seem to us that there are things—even while asleep or anaesthetized.<sup>62</sup>

These considerations do, however, illustrate an ambiguity in Mulder’s Seeming Platitude, which may help clarify a bit why his Seeming Problem may seem initially compelling. There is a way in which Mulder is correct that suitable HOTs do not make things seem a certain way, depending on how we understand ‘things’. When considering in general states that make things seem a way—such as first-order perceptual states—such states typically make one aware of *things in the external world*.

---

<sup>59</sup> Mulder 2016, p. 461.

<sup>60</sup> See, e.g., Rosenthal 2002, p. 417, fn. 3; Windt et al. 2016.

<sup>61</sup> Mulder 2016, p. 452.

<sup>62</sup> For experimental evidence that people can learn new information by perceiving stimuli during sleep, see, e.g., Windt et al. 2016.

A visual perception of a square makes it seem that there is a square in the external world. In the case of consciousness, by contrast, what are the things that seem a way to the subject? If we accept the TP, the answer is: one's mind.

Mulder's Seeming Problem may therefore seem compelling at first insofar as it trades on a confusion of the ways things *in the external world* seem (which are a function of one's first-order states) with how things *in one's internal mental world* seem (which, according to HOT theory, are a function of one's HOTs).<sup>63</sup> In ordinary consciousness, suitable HOTs, strictly speaking, do not make one aware of features of the external world. They make one aware of one's mental states. In that way, Mulder's criticism may seem persuasive. But since suitable HOTs represent first-order states as rendering one aware of the external world, they nonetheless do make one aware, albeit obliquely, of the external world. That is, they make the external world seem to seem a certain way.

A question nonetheless remains: supposing, as Tucker and others do, that conscious seemings have a distinctive assertoric feel, what secures the assertoric character of the conscious seemings that HOTs engender? Even if HOTs are themselves assertoric states, one might worry that nothing could explain the assertoric nature of conscious states since HOTs are themselves seldom conscious. But the answer is straightforward: suitable HOTs present oneself as being in states that make things seem a way. The assertoric nature of the conscious seeming is inherited from the fact that the states that HOTs assert that one is in are themselves assertoric. And this includes HOTs when they are conscious in introspection. Since HOT theory conceives of the character of consciousness as a matter of what states one is suitably aware of being in, being suitably aware of having certain thoughts, including HOTs, makes things consciously seem a way.

## **6. Seemings, Awareness, and Consciousness**

As noted in the Introduction, much of the reason that debates in consciousness remain so vexed is that there is little consensus about how to understand consciousness in a commonsense way. The TP seeks to provide that gloss, but it, and thus consciousness, is often misunderstood. Indeed, I suspect that much of what motivates Mulder's concerns about HOT theory stem from the shift in discussions of the TP and HOT theory between uses of 'awareness' and related expressions, which are factive, and uses of 'seeming' and related expressions, which are nonfactive. Ordinary usage has it that one cannot be aware of *X* unless *X* actually exists, though it can seem to one that *X* is present even if *X* is not. And the TP is often framed in terms of both. Sometimes it is said that a conscious state is a state that one is aware of being in; sometimes it is said that a conscious state is a state that it seems to one that one is in. Mulder's Seeming Platitude seeks to bridge these two readings, but the Seeming Problem holds that HOTs are unable to respect that platitude.

As I have illustrated here, however, there are increasingly accurate ways of framing the TP. Though the TP is often glossed as the claim that one is in a conscious mental state only if one is aware of that state, my being aware of a mental state of yours plainly has nothing to do with your state's being state conscious. So a more accurate,

---

<sup>63</sup> Cf. Carruthers' (2005, p. 42) distinction between what *the world is like* from a creature's perspective and what a creature's *perception of the world is like* from its perspective.

though somewhat more tedious, way of putting the TP is that a mental state is conscious only if one is aware of *oneself* as currently being in that state. Some of these ways of characterizing the TP are fine in a pinch, but the fact that it is often put loosely can lead to confusion.<sup>64</sup>

In this case, I think the factivity of ‘awareness’ may mislead some into worrying that HOTs cannot engender both awareness of our mental lives and appearances of those mental lives. But, as some HOT theorists have urged,<sup>65</sup> there is always a way to shift between the above two ways of putting the TP because if it seems to one that one is in a state in a way that engenders consciousness, then what one is really aware of is not that state, but of oneself as being in that state. And, since being in a state that engenders awareness of any kind requires a thinker, one’s states of awareness are always factive in respect of what they render one aware of—namely, oneself. This shows why it is fine to move from the reading of the TP as the claim that conscious states are states that it seems to one that one is in to the claim that they are states that one is aware of oneself as being in.

But the fact that we use ‘awareness’ in this way of glossing the TP, which is factive, does not preclude us from moving in the other direction too. I hope to have shown that the claim that conscious states are states that we are suitably aware of ourselves as being in simply entails the claim that they are states that it seems to us that we are in. Moreover, even though ‘awareness’ is factive in respect of the objects of one’s awareness, it is not in respect of those objects’ properties: I cannot be aware of my computer if my computer is not there, but I can be aware of my computer as being on when in fact it is off. That is, even if I am correctly described as being aware of my computer, it can seem to me that it is on when it is off.<sup>66</sup> Likewise, though a suitable HOT cannot make me aware of myself as being in a state if I do not exist, if I have a HOT at all then I do exist, and that HOT can make it seem to me that I am in a state that I am not in.

If the forgoing is right, then how things actually seem to me (how things seem to me, given my first-order states) might come apart from how things consciously seem to me (what mental states it seems to me that I am in due to my HOTs). And that strikes many as absurd. Perhaps this worry is what at bottom motivates Mulder. But notice that this is just the targetless/mismatching-HOT objection to HOT theory, which Mulder claims he is unconcerned with and which I have noted is not really a problem. If we are moved by the TP, then conscious states are states that it seems to us that we are in. But that does not require that we actually be in those states. It only requires that we be in suitable HOTs to engender those mental appearances.

## 7. Conclusions

In the end, even if Mulder were correct that HOTs as they are currently understood do not explain the relevant mental appearances, we should understand his criticism not as

---

<sup>64</sup> For more on such confusions, see, e.g., Berger 2014b, Section 4.

<sup>65</sup> See, e.g., Weisberg 2011, pp. 418-419; Berger 2014b, pp. 838-839.

<sup>66</sup> Some disagree. Edwards and Platchais (2016) recently maintain that awareness is factive both in respect of the object and the object’s properties. But this view is motivated by questionable epistemic concerns and strains common sense. Would we really want to deny that I am aware of my computer if I misperceive a single feature of it in the slightest degree?

a refutation of HOT theory, but as a request for more information about how it should work. There is, of course, an ordinary psychological notion of how things (including our mental lives) seem to us. And we can theorize what kind of mental states could explain such seemings. Rosenthal argues that the appearances of consciousness are explained by HOTs of a certain kind. If that were not quite right, or were not the full story, then what we would need is more theorizing—more explanation of how HOTs or some other kind of mental activity explain our mental appearances. I do not believe that we need more theorizing of this sort. But the fact that Rosenthal’s project provides a way to move forward in thinking about consciousness is perhaps its greatest strength, even if its details were to require more work.

**Acknowledgements:** I thank Richard Brown, Zoe Jenkin, Myrto Mylopoulos, David Pereplyotchik, Jake Quilty-Dunn, Evan Rodriguez, Jim Skidmore, Pamela Wahl, Russell Wahl, and Josh Weisberg for their comments on drafts of this piece and/or conversations about this material. Special thanks to David Rosenthal for many helpful discussions of HOT theory over the years.

## References

Akhtar, Salman

2009 *Comprehensive Dictionary of Psychoanalysis*. London: Karnac Books.

Armstrong, David M.

1968 *A Materialist Theory of the Mind*. New York: Routledge and Kegan Paul.

Berger, Jacob

2014a “Mental States, Conscious and Nonconscious,” *Philosophy Compass* 9 (6): 392-401.

Berger, Jacob

2014b “Consciousness is Not a Property of States: A Reply to Wilberg,” *Philosophical Psychology* 27 (6): 829-842.

Block, Ned

2011a “The Higher-Order Approach to Consciousness is Defunct,” *Analysis* 71 (3): 419-431.

Block, Ned

2011b “Response to Rosenthal and Weisberg,” *Analysis* 71 (3): 443-448.

Brown, Richard

2015 “The HOROR Theory of Phenomenal Consciousness,” *Philosophical Studies* 172 (7): 1783-1794.

Byrne, Alex

1997 “Some Like it HOT: Consciousness and Higher-Order Thoughts,” *Philosophical Studies* 86 (2): 103-129.

Carruthers, Peter

2005 *Consciousness: Essays from a Higher-Order Perspective*. Oxford: Oxford University Press.

Chalmers, David J.

1996 *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.

Chisholm, Roderick

- 1957 *Perceiving: A Philosophical Study*. Ithaca, New York: Cornell University Press.
- Dennett, Daniel C.  
1969 *Content and Consciousness*. New York: Routledge and Kegan Paul.
- Edwards, James, and Dimitris Platchias  
2016 "Epistemic Warrants and Higher-Order Theories of Conscious Perception," *Pacific Philosophical Quarterly* DOI: 10.1111/papq.12161.
- Egan, Andy  
2008 "Seeing and Believing: Perception, Belief Formation, and the Divided Mind," *Philosophical Studies* 140 (1): 47-63.
- Gennaro, Rocco  
2012 *The Consciousness Paradox: Consciousness, Concepts, and Higher-Order Thoughts*. Cambridge, MA: MIT Press.
- Goldman, Alvin  
1993 "Consciousness, Folk Psychology, and Cognitive Science," *Consciousness and Cognition* 2 (4): 364-382.
- Kouider, Sid, and Stanislas Dehaene  
2007 "Levels of Processing During Non-Conscious Perception: A Critical Review of Visual Masking," *Philosophical Transactions of the Royal Society B* 362 (1481): 857-875.
- Lau, Hakwan, and David M. Rosenthal  
2011 "Empirical Support for Higher-Order Theories of Conscious Awareness," *Trends in Cognitive Sciences* 15 (8): 365-373.
- Levine, Joseph  
2001 *Purple Haze: The Puzzle of Consciousness*. Oxford: Oxford University Press.
- Lycan, William G.  
1996 *Consciousness and Experience*. Cambridge, MA: MIT Press.
- Mandelbaum, Eric  
2016 "Attitude, Inference, Association: On the Propositional Structure of Implicit Bias," *Noûs* 50 (3): 629-658.
- Mulder, Jesse M.  
2016 "A Seeming Problem for Higher-Order Theories of Consciousness," *Dialogue: Canadian Philosophical Review* 55 (3): 449-465.
- Nagel, Thomas  
1974 "What Is It Like to Be a Bat?," *Philosophical Review* 83 (4): 435-450.
- Neander, Karen  
1998 "The Division of Phenomenal Labor: A Problem for Representational Theories of Consciousness," *Philosophical Perspectives* 12: 411-434.
- Pereplyotchik, David  
2015 "Some HOT Family Disputes: A Critical Review of *The Consciousness Paradox* by Rocco Gennaro," *Philosophical Psychology* 28 (3): 434-448.
- Phillips, Ian, and Ned Block  
2016 "Debate on Unconscious Perception," in Bence Nanay (ed.), *Current Controversies in Philosophy of Perception*, New York: Routledge, pp. 165-192.

- Quilty-Dunn, Jake  
2015 "Believing in Perceiving: Known Illusions and the Classical Dual-Component Theory," *Pacific Philosophical Quarterly* 96 (4): 550-575.
- Rey, Georges  
2000 "Role, Not Content: Comments on David Rosenthal's 'Consciousness, Content, and Metacognitive Judgments,'" *Consciousness and Cognition* 9 (2): 224-230.
- Rosenthal, David M.  
1986 "Two Concepts of Consciousness," *Philosophical Studies* 49 (3): 329-359.
- Rosenthal, David M.  
2000 "Metacognition and Higher-Order Thoughts," *Consciousness and Cognition* 9 (2): 231-242.
- Rosenthal, David M.  
2002 "Explaining Consciousness," in David Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*, Oxford: Oxford University Press, pp. 406-421.
- Rosenthal, David M.  
2005 *Consciousness and Mind*. Oxford: Clarendon Press.
- Rosenthal, David M.  
2011 "Exaggerated Reports: Reply to Block," *Analysis* 71 (3): 431-437.
- Rosenthal, David M.  
2012 "Higher-Order Awareness, Misrepresentation and Function," *Philosophical Transactions of the Royal Society B* 368 (1594): 1424-1438.
- Shargel, Daniel  
2016 "The Insignificance of Empty Higher-Order Thoughts," *Journal of Cognition and Neuroethics* 4 (1): 113-127.
- Shoemaker, Sydney  
1982 "The Inverted Spectrum," *Journal of Philosophy* 79 (7): 357-381.
- Stich, Stephen P., and Ted A. Warfield (eds.)  
1994 *Mental Representation: A Reader*. Oxford: Basil Blackwell.
- Tucker, Chris  
2010 "Why Open-Minded People Should Endorse Dogmatism," *Philosophical Perspectives* 24 (1): 529-545.
- Tucker, Chris (ed.)  
2013 *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. New York: Oxford University Press.
- Weisberg, Josh  
2011 "Misrepresenting Consciousness," *Philosophical Studies* 154 (3): 409-433.
- Weiskrantz, Lawrence  
1997 *Consciousness Lost and Found: A Neuropsychological Exploration*. Oxford: Oxford University Press.
- Windt, Jennifer M., Tore Nielsen, and Evan Thompson  
2016 "Does Consciousness Disappear in Dreamless Sleep?," *Trends in Cognitive Sciences* 20 (12): 871-882.