

The HOROR Theory of Phenomenal Consciousness

Richard Brown

(Forthcoming) *Philosophical Studies*

The final publication is available at Springer via <http://dx.doi.org/10.1007/s11098-014-0388-7>

I. Introduction

'Phenomenal consciousness' has emerged as a widely used technical term to denote a phenomenon that is familiar from our everyday experience. We feel heat, taste chocolate, anxiously think about the future, reflect on a past love, and on and on. Though we are merely at the beginning of our study of the brain we already know a good deal about what is going on in the brain and extended nervous system when we experience these kinds of events but we don't know why the particular activity of the brain, which we theoretically interpret as information-processing, should have these phenomenal aspects. At the common sense level it really does appear as though we could have all of that information-processing and yet not have any phenomenal consciousness. This is the Hard Problem of consciousness as formulated by David Chalmers (Chalmers 1996, 2010) and which goes back at least to Leibniz's famous Mill thought-experiment. We can formulate it in terms of there being something that it is like for the organism (Nagel 1974). Why is there something that it is like for me to have information-processing relating to tissue damage (in the case of pain)? In particular why is it –painful- rather than anything else? Why not tickling? Or tasting chocolate? Or even nothing at all?

Construed this way, addressing the Hard Problem involves giving a theoretical account of something whose existence is known already. It is comparable to taking a particular object firmly in hand, pointing at it, and then asking about its fundamental nature. 'What is this thing right here?' we ask while pointing directly to the thing in question. It is at that point that we begin to offer theoretical accounts of the phenomenon in question but there can be no doubt at the beginning that phenomenal consciousness exists. There is some temptation to follow Descartes and say that it is the thing that I am most certain of in the entire world. I can doubt that I am awake (perhaps this is a dream) but how can I doubt that I am conscious and that I now consciously experience the feel of the keyboard and hear the humming of the air-conditioner? Any theory that denies these seemingly obvious facts is not giving a theory of consciousness that takes seriously the first-person data.

One very popular approach to this problem has been to connect phenomenal consciousness to representations of a certain kind. This strategy is generally known as *representationalism* and there are many varieties currently on the market ranging from reductive physicalist accounts (Tye 2000, Lycan 2001, Dretske 2003, Rosenthal 2005, Gennaro 2012) to non-reductive non-physicalist

accounts (Pautz 2009, Chalmers 2010).¹ Representationalism, in its most general form, holds that phenomenal consciousness consists in, or at least supervenes on, a certain kind of representation and can be neutral with respect to the ontology of those representations. Put another way, representationalism in its most general form holds that when there is something that it is like for one to see blue, hear a trumpet, feel a pain, or think a thought this is because one instantiates a certain kind of representation. Representationalism of some kind has strong intuitive support. Phenomenally conscious experiences have correctness conditions and that suggests that they are representations. Representations have intentional contents which are about, or directed at, their intentional targets and so it is reasonable to explore the connection between intentionality and consciousness (See Lycan 2008 for an overview of arguments for representationalism).

One important division in representational theories of consciousness is that between first-order and higher-order theories. First-order theories claim that the relevant representations, the ones which supervene on or just are what phenomenal consciousness consists in, represent properties in the physical world, or at least represent non-mental properties. This is why they are first-order states. Their representational targets are physical properties in the world. Higher-order theories claim that the relevant representations represent mental properties; in particular they represent my mental life as instantiating mental properties. Mental properties include mental attitude, intentional content, and qualitative content (that is, first order representations of worldly properties like sounds, colors, locations, etc).

Higher-order theories are often interpreted as relying on a special relation between the first-order state that is represented and the higher-order state that does the representing. This interpretation of higher-order theories is explicitly endorsed by many fans and critics of the theory alike (Balog 2000, Lycan 2004, Gennaro 2004, Mandik 2009, Mately 2011, Block 2011a, Kriegel 2011, Kidd 2012,). The interesting thing about this way of thinking about higher-order theories is that it has the tendency to downplay the fact that higher-order theories are representational theories. This is so even though most of the above-cited authors begin their discussion of higher-order theories by emphasizing that it is indeed a kind of representational theory.

However there is another way to interpret higher-order theories (Brown 2012a; Lau & Brown forthcoming). One can emphasize the fact that it is a representational view and insist that phenomenal consciousness consists in

¹ It should be noted that Rosenthal uses 'thought' pretty consistently in his writing and has denied that his view is a version of representationalism. However by this he means to deny the view that mental qualities are intentional. He has maintained that mental qualities represent in a distinctive non-intentional manner (see pages 119 and 222 in Rosenthal 2005). And in any case higher-order thoughts are intentional states and so his theory of consciousness (as opposed to mental quality) is straightforwardly representational.

having suitable higher-order representations.² On this alternative way of thinking about the higher-order theory there is no explanatory role for a relation between first and higher-order states. Rather the explanatory power lies in the nature of the higher-order representation in question. I have introduced the acronym 'HOROR' to distinguish this version from the relational version of higher-order theory as defined by the above authors. 'HOROR' stands for 'Higher-Order Representation Of a Representation'. In the case of consciously seeing red, for instance, the higher-order representation will be something to the effect of 'I am having this visual representation of red' (Lau & Rosenthal 2011).

As I see things HOROR theory is the right way to interpret higher-order thought theories like those of David Rosenthal (2005) and the relational reading that is so wide-spread in the literature is mistaken but there are those that disagree (on both accounts).³ I don't think that the debates about the right way to characterize the theory, or whether Rosenthal changed his mind about the relational view (Block 2011b), are very interesting. What is important here is that we have distinct theoretical accounts of phenomenal consciousness on offer. Once one sees that there are multiple theoretical interpretations of the basic phenomenon under investigation we can look at the body of evidence at our disposal and see which of the theories is better supported overall.

In the next section I develop the HOROR theory a bit more and compare it to the same-order and higher-order thought theories. After doing that I move to presenting some considerations in favor of HOROR theory.

II. The HOROR

The central claim of HOROR theory is that phenomenal consciousness consists in instantiating a certain kind of higher-order representation. Suppose that one is having a vivid conscious experience of red. In that case HOROR theory says that the relevant, seemingly non-inferential, higher-order representation will have something like the following as its intentional content: *I, myself, am having this visual representation of red*. This higher-order representation represents that one is oneself having or instantiating a certain first-order mental representation (or perhaps that one is instantiating a certain set of first-order representations). This representation is thought-like in that it is composed of intentional contents and arranged so as to have assertoric force. Because one represents oneself as instantiating various mental properties it will appear to one as though one has

² I follow the tradition in positing that suitable higher-order representations are those that at least appear to the subject as non-inferential

³ For instance Weisberg 2011 and Rosenthal 2011 seem to suggest that they are eliminativists about phenomenal consciousness. If so then they would disagree with the way I have set the problem up but they equate 'phenomenal consciousness' with Block's first-order view and claim there is no such first-order property. I claim that as long as we are careful to use 'phenomenal consciousness' and related 'what it is like' terminology in a theoretical neutral way then everyone in the debate agrees that there is phenomenal consciousness.

these various mental qualities and HOROR theory claims that this is what phenomenal consciousness really is.

These higher-order representations account for the two different characters of phenomenal consciousness (*cf* Kriegel 2012). On the one hand we have the state's subjective character. The subjective character of a phenomenally conscious experience is that part of the experience which makes it the case that it is experienced as being for me. On the HOROR theory this comes from the fact that one represents oneself as being in these various states. On the other hand, the phenomenal character of a conscious experience is that part of the experience which would distinguish a conscious experience of red from a conscious experience of blue, and both of those from the conscious experience of the sound of a trumpet. According to the HOROR theory this will turn out to be the intentional content of the higher-order representation. If one were consciously experiencing a nice vivid red then the relevant higher-order representation in question will have something like *'I am, myself, having this visual representation of red'* as its intentional content. If one were consciously seeing green instead then the content would be something like, *'I am, myself having this visual representation of green'*. According to HOROR theory the difference between a phenomenally conscious experience of green and one of red is in the content of the higher-order representation not in the first-order states that one instantiates.

On HOROR theory phenomenal consciousness is not a property that the first-order representation has.⁴ I admit that this is counter-intuitive. We might have thought that phenomenal consciousness should be a property that my first-order sensing of tissue damage should have, but on this view it is not. Because of this some may think that HOROR theory is too revisionary in its construal of what phenomenal consciousness is. If phenomenal consciousness is a higher-order representation then that is as good as saying that the thing we thought we were talking about doesn't exist. I think this is a mistake. This is only a cost to the theory if it is at odds with what we discover scientifically. In the course of our investigations into the nature of physical reality we have made many surprising discoveries, including the nature of combustion, the relationship between electricity and magnetism, relativity theory, and quantum mechanics, just to name a few and I am sure more are in store. I do not find HOROR theory any more surprising than the fact that there is no absolute simultaneity. If the view is supported by philosophical reasoning as well as empirical considerations it should be taken seriously in spite of what our intuitions tell us. It is a reasonable hypothesis about the fundamental nature of the thing that we are the most certain

⁴ I can be neutral with respect to the nature of first-order representations and I take no stance on whether they are intentional, or accounted for by some kind of Quality Space Theory. However I am committed to the claim that first-order states have intentional and qualitative contents and that they account for the behavior of the creature for the most part. When I consciously experience red, in the typical case, I will have a first-order state that has qualitative red as its content and a higher-order representation of myself as instantiating that first-order state. The first-order state accounts for my discrimination behavior and the higher-order state accounts for what it is like for me.

of in all of the world. In this sense I could not be further from eliminativism. I believe that consciousness is a real feature of our world and that one theory about the nature of that stuff is that it amounts to a certain kind of higher-order representation.

From what has been said so far one may wonder whether this is a same-order or a higher-order view. Same-order theories of consciousness are self-representational theories. That is, they appeal to a distinctive kind of relationship between the higher-order and first-order contents of a particular conscious state. They must both be parts of a single state. HOROR theory will then look superficially like these kinds of theories. The higher-order representation consists of a thought-like state that represents the subject as instantiating various mental properties and so the theory does claim that phenomenal consciousness consists in a single representational state. However the state does not represent itself. It represents the subject, not that very state. We can put this by saying that the relevant higher-order representation is phenomenally conscious but not state conscious.

State consciousness is distinct from phenomenal consciousness. Phenomenal consciousness has to do with whether there is something that it is like for one to be in various states. HOROR theory posits that phenomenal consciousness consists in one representing oneself as being in various first-order states. That is what phenomenal consciousness is on this view. So it cannot occur unconsciously in that sense but it can occur when it is not state-conscious. State consciousness has to do with whether a given mental state is the representational target of a suitable higher-order representation. A mental state is state-conscious just in case it is targeted by a higher-order representation. HOROR theory claims that the relevant higher-order representation is phenomenally conscious but not because it is the target of a further higher-order representation. Thus it is a possibility on HOROR theory that one might have a phenomenally conscious state, by having a suitable higher-order representation, yet fail to have a conscious state in the state-conscious sense. This might happen if one were to not actually be in the first-order state that one represents oneself as being in.⁵ However, typically it will be the case that one is in states with both properties. For instance in the typical case of consciously seeing red one will have a first-order state which represents the physical property –physical red–, and a higher-order representation that one is in fact in that first-order state. In that case the first-order state is state-conscious and the higher-order representation is phenomenally conscious.

⁵ This is an example of the so-called empty higher-order representations, which I will discuss in more detail in the next section. I will note here, though, that I am neutral with respect to the question of whether it is the notional state that has the property of being conscious. This may not be as strange as it sounds; however, one potential problem is that of explaining why it isn't always the notional state that has this property.

Some may think that backing off of the relational view means that one has given up the basic principle of the higher-order theory known as the Transitivity Principle (Rosenthal & Weisberg 2008). This principle states that a conscious state is one that I am conscious of myself as being in, or equivalently, that a state that I am in no way aware of myself as being in doesn't count as a conscious state. The transitivity principle is usually invoked to give an explanation of state-consciousness and as we have seen I do accept that it gives us an explanation of that property. What about in the case of phenomenal consciousness? Here things are bit trickier.

In one sense HOROR theory does endorse the transitivity principle. This is because on that view phenomenal consciousness consists in *implementing* the transitivity principle. That is to say that it consists in *representing* oneself as being in various mental states. I interpret the transitivity principle as claiming that a phenomenally conscious mental state consists in representing oneself as instantiating first-order representation. In this sense HOROR theory claims that it is the mechanism of state-consciousness, which really is what phenomenal consciousness turns out to be. However, there is a sense in which HOROR theory does back off of the transitivity principle with respect to phenomenal consciousness. This is because HOROR theory does allow that there are phenomenally conscious states that the subject is in no way aware of being in. This is just the case of having a suitable higher-order representation that one is not aware of oneself as having, which happens most of the time. When I consciously see red, for example, this is because I am instantiating a suitable higher-order representation but I am not aware of myself as being in that higher-order state. In order for that to happen I would have to instantiate a further state representing me as being in the previous state, which is not what happens in ordinary perceiving. I see this as an advantage of the theory rather than a drawback, and I will discuss it more fully in the next section.

Some may object to the claim that the higher-order state is itself phenomenally conscious on the grounds that this is a property of persons not states (Berger 2013). Even the 'what it is like for one' talk seems to suggest that the thing that has the property is the 'one' not the state in virtue of which that creature has the property. There is a way of taking this point that does not do violence to the argument of this paper. The issue is which kind of state is responsible for the person having this property? Here this critic will agree that it is the higher-order state that makes this the case.

But there is also a way of reading this claim that just makes the point that I want to make in the first place. A state is phenomenally conscious when there is something that it is like for one to be in that state. Thus we can define a state being phenomenally conscious in terms of its making the creature in question instantiate this property. If one is thinking about phenomenal consciousness in this way then we will say that the higher-order state is phenomenally conscious since it is the state that there is something that it is like for you to have. So while

it is true that the person has the property of being phenomenally conscious the nature of that property, which the creature has, is instantiating the right kind of higher-order representation.

Before leaving this section I want to make clear that HOROR theory, as I see it, is in principle neutral with respect to the metaphysics of consciousness. This is exactly what we should expect given that it is, in its most general form, simply a version of representationalism about consciousness. The central claim of HOROR theory is that phenomenal consciousness could be a kind of higher-order representation. If one thought that one could account for the relevant representations and intentional contents in purely physical terms then one could construct a reductive version of HOROR theory. On the other hand if one thought one could not account for the representations or their representational contents in physicalist terms one could have a non-reductive version of HOROR theory. This non-reductivism could stem from the fact that these representations or their contents are themselves not physical. Or it could stem from the fact that these representations or their contents emerge from physical processes, or just are certain physical processes. If it turned out that a functional duplicate failed to instantiate these kinds of representations in the actual world then we would know that these representations depend on some physical or biological feature of the brain or its physical composition. At that point whether you called the resulting view representationalism or not seems like mostly a verbal issue.

Up until this point I have merely been trying to clarify the view. I haven't yet given any reasons to think that HOROR theory is right. One may grant that it is a theoretical possibility that one could occupy but why believe it? In the next section I will argue that we do have some good reasons to take HOROR theory seriously.

III. The Case for HOROR Theory

When I have a phenomenally conscious pain or consciously see blue it is not merely that there is something that it is like to have these experiences but also that there is something that it is like *for me*. Sometimes this is brought out by appeal to the phrase itself 'what it is like for me to see blue' has two parts. One is the blueness of the experience. That is what distinguishes it from what it is like for me to see red, or yellow, or hear a bell, etc. The other is that what it is like *is like something for me*; I represent the experience as belonging to me in some way.

The notion that there is some kind of awareness associated with phenomenal consciousness, or there being something that it is like for one, is widely accepted (Rosenthal 2005, Block 2007, Brown 2012b, Kriegel 2012, Sebastian forthcoming).⁶ One thing to note right off that but is that HOROR theory has a

⁶ Even those dualists who are attracted to panquality-ism, like David Chalmers (forthcoming), will agree to this. He may think that there can be phenomenal redness as a fundamental part of our

ready explanation of why this should be so. A phenomenally conscious experience, on that theory, consists in representing oneself as instantiating or having some first-order mental representation, so of course it is like something *for you* to be in that higher-order state. Phenomenal consciousness is a matter of mental appearances and this state represents you yourself as having a mental life characterized by various mental properties. Since this state attributes these properties to you and represents them as being states that you are at that moment in, it will appear to you that you yourself are having the relevant first-order representations. First-order views do not have a satisfactory account of the kind of awareness that seems to be at least partially constitutive of phenomenal consciousness.

Block has, for instance, appealed to a deflationary notion of awareness (2007). But this does not seem to be able to do justice to the notion of awareness that is at play here (Brown 2012c). For instance on a deflationary notion of self-representation a state represents itself merely in virtue of one being in the state in question. The main difficulty for this view lies in explaining the notion of self-representation that is in play here. It cannot be anything like the self-representation that a higher-order theory would posit. So, it cannot be an intentional or conceptual representation of the state or its content that is doing the work. The problem is further compounded when we take into account Block's allowance of unconsciousness perception. On his view a first-order representation of red can occur consciously or unconsciously. When it is conscious in the phenomenal sense there will be something that it is like for one to have the experience. On Block's view that will mean that the first-order state has come to represent itself. But what has changed? On Block's account it may have something to do with the firing strength of the relevant neural areas. But how can this plausibly be interpreted as the state representing itself? Further how can this account for the kind of awareness that we discover phenomenologically? It is not impossible that a first-order account can be given but as of now there is none in the offing and so to ignore HOROR theory and its advantage here seems to be evidence of a theoretical bias.

What we have seen so far is that the HOROR theory has an advantage explaining why phenomenally conscious experiences are like something for the subject of those experiences. By itself this is enough to merit taking the theory seriously but when coupled with the absence of a first-order alternative account that is at least as good as the one provided by HOROR theory we can see that the HOROR theory clearly has the advantage here.

Additional evidence for this position comes from what we might call the argument from concept acquisition. Rosenthal has made this argument in various places (Rosenthal 2005, Rosenthal 2012). Anyone who has had experience with wine

ontology but that will not count as a conscious experience in the sense we are interested in since it is not experienced by anyone (cf Chalmers 2013 for an endorsement of this line of thought in response to Hellie).

will know that acquiring a new word will sometimes allow one to make finer-grained distinctions in the experience that one has. One interpretation of what is going on here is that learning the new word results in one's having a new concept and the application of this concept allows one to represent one's mental life in a more fine-grained way. This results in more phenomenal properties in one's experience. It carves it up in a more fine-grained way phenomenologically. On HOROR theory that amounts to the claim that one represents one's mental life as instantiating different mental qualities.

A note of caution should be sounded here. I am not trying to claim that the argument from concept acquisition shows that higher-order theories are true. But it does at least suggest that concept application is importantly involved in generating phenomenally conscious experience. I am willing to grant that what the argument suggests is that concept application can make a difference in phenomenal experience and that we need a further argument to show that the way in which it makes this difference is by being applied in a suitable higher-order representation. It could be the case that it does this by changing the first-order states in some way. Thus though this argument is suggestive our ultimate decision about which of these models is right will depend on empirical evidence and so will be considered in the next section. Even so, though, the argument from concept acquisition is still important in that it shows that there is no conceptual absurdity in the claim made by HOROR theory.

Another thing to notice is the ways in which the HOROR theory allows us to capture the same kinds of intuitions that are usually thought to favor the first-order view. If one looks at the literature from the last 20 years or so and every time one sees the word 'qualia', or 'phenomenal property', or what have you, you substitute in 'higher-order representation of oneself as instantiating first-order mental representations' one will end up with something that is true on the HOROR theory. I will not go through all of the individual intuitions but will mention just one representative one, which I mentioned earlier already. On the HOROR theory it will turn out to be true that a phenomenally conscious state cannot occur unconsciously in the phenomenal sense. First-order states can occur so, and in fact always do, but if phenomenal consciousness consists in the appropriate kind of higher-order representations then when those representations occur they must be conscious. This captures the first-order theorists' intuition that consciousness is an intrinsic non-relational property. Phenomenal consciousness consists in me representing my mental life as instantiating various mental properties and this is why there is something that it is like for me, so in addition to capturing the intuition we are able to explain why it is that way.

As was noted in the introduction HOROR theory is distinct from those theories that posit a relation between the first-order state and the higher-order state as playing an important role in the explanation of phenomenal consciousness. These kinds of theories face problems from empty higher-order thoughts (Block 2011a). In those kinds of cases one has a higher-order representation without

the first-order state occurring. According to HOROR theory there will be no difference in the phenomenology of the experience, though there will be a difference in the mental functioning and so in the creature's behavior. Since there is no relation between first and higher-order states empty higher-order thoughts are no problem for the HOROR theory. They are a problem for specific versions of the relational higher-order theory. In particular they are a problem for those versions of the theory according to which a first-order state or representation is a necessary condition for there being a phenomenally conscious state. But in the empty higher-order cases there is no first-order state. So the theory has a necessary condition which is not met and which is in conflict with the sufficient condition. HOROR theory has no such problems.

What the argument from empty higher-order representations shows is that there is a problem for any higher-order theory of consciousness that denies that the higher-order representation at least partially determines what it is like for one to have the experience. This includes theories like those of Kriegel (2011) and Gennaro (2012). However there is an alternative way of developing the relational account (Lau & Brown forthcoming). On this alternative account the higher-order representation is responsible for making it the case that one is having a phenomenally conscious experience but the first-order state is responsible for the exact way the experience is like for you to have. So according to this 'joint-determination' view if one had an empty higher-order representation there would still be some phenomenal consciousness, though it would be partial. The experience would be something to the effect that one is seeing something or other but one would be lacking the specific phenomenal properties contributed by the first-order states. This kind of view counts as higher-order because it claims that phenomenal consciousness consists in being aware of oneself as being in the first-order state. It interprets this as having a referential relation between the first-order state and the higher-order state.

This kind of view avoids the problem of empty representations but it does so at the expense of giving up the explanatory power of the higher-order strategy. If one has a first-order representation of physical red and then one comes to have a higher-order representation to the effect that one is in THAT STATE, where THAT STATE is the address of, or some kind of pointer to, the first-order state, the question becomes one of how the first-order state contributes its property to the content of the higher-order state, and thus to what it is like for the creature. If one wants to say that the phenomenal redness of my conscious experiences comes from the first-order state then we can ask why it is the case that it is not phenomenally red even when it is not represented by the higher-order state? In addition we can ask how that state contributes the phenomenal redness to the mental appearances. How can thinking 'I am in that state' (pointing at some state) make one aware of oneself as being in the state in respect of its mental properties? This view might make more sense as a non-reductive view. If one is pointing at a primitive non-reducible phenomenal redness then it is easy to see why it is like seeing red for one.

To conclude this section I think that the joint-determination view is an interesting and important alternative way of thinking about higher-order theories. I think HOROR theory has the explanatory advantage but only further empirical and philosophical work will be able to settle this issue. In addition to this we can conclude that the relational higher-order views that deny that phenomenal consciousness is (at least) partially constituted by the higher-order state as well as first-order views more generally have difficulties.

IV. Conclusion

My aim in this paper has been to explore a higher-order representational approach to phenomenal consciousness and to give some reasons why we might take such a view seriously. This approach offers a way of combining the merits of several distinct strands of representational theorizing.⁷

⁷ A previous version of the paper was delivered at the CUNY Graduate Center Colloquium Series September 5th 2012 and portions were also delivered at the Association for the Scientific Study of Consciousness meeting July 13th 2013. I am grateful to participants for very helpful discussion. I would also like to especially thank Jake Berger, Ned Block, David Chalmers, Hakwan Lau, Pete Mandik, and David Rosenthal for helpful discussion and comments on previous drafts.

Work Cited

Balog, Katalin (2000). Phenomenal Judgment and the HOT theory: Comments on David Rosenthal's "Consciousness, Content, and Metacognitive Judgments". *Consciousness and Cognition* 9 (2):215-219.

Berger, Jacob (2013). Consciousness is not a property of states: A reply to Wilberg. *Philosophical Psychology*.

Block, Ned (2007). Consciousness, Accessibility, and the Mesh between Psychology and Neuroscience. *Behavioral and Brain Sciences* 30:481--548.

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 (2012a). Review of 'The Consciousness Paradox: Consciousness, Concepts, and Higher-Order Thoughts' by Rocco J. Gennaro. *Notre Dame Philosophical Reviews*.

Brown, Richard (2012b). The Brain and its States. In Shimon Edelman, Tomer Fekete & Neta Zach (eds.), *Being in Time: Dynamical Models of Phenomenal Experience*. John Benjamins.

Brown, Richard (2012c). The Myth of Phenomenological Overflow. *Consciousness and Cognition* 21(2): 599-604

Carruthers, Peter (2008). Higher-order theories of consciousness. *Stanford Encyclopedia of Philosophy*.

Chalmers, David (1996) *The Conscious Mind: In Search of a Fundamental Theory* Oxford University Press

Chalmers, David (2010) *The Character of Consciousness* Oxford University Press

Chalmers, David (2013) "The Contents of Consciousness: Reply to Hellie, Peacocke, and Siegel". *Analysis* 73:345-68

Chalmers, David (forthcoming) Panpsychism and Panprotopsychism. Available here: <http://consc.net/papers/panpsychism.pdf>

Dretske, Fred (2003). Experience as representation. *Philosophical Issues* 13 (1):67-82.

Gennaro, Rocco J. (2004). Higher-order theories of consciousness: An overview. In Rocco J. Gennaro (ed.), *Higher-Order Theories of Consciousness*. John Benjamin.

Gennaro, Rocco J. (2012) *The Conscious Paradox: Consciousness, Concepts, and Higher-Order Thoughts* The MIT Press

Kidd, Chad (2011). Phenomenal consciousness with infallible self-representation. *Philosophical Studies* 152:361-383.

Kriegel, Uriah (2011). Self-Representationalism and the Explanatory Gap. In J. Liu & J. Perry (eds.), *Consciousness and the Self: New Essays*. Cambridge University Press.

Lau, Hakwan & Brown, Richard (forthcoming). The Emperor's New Phenomenology? The Empirical Case for Conscious Experience without First-Order Representations. In Adam Pautz & Daniel Stoljar (eds.), *Themes from Block*. MIT.

Lau, Hakwan and Rosenthal, David (2011) Empirical Support for Higher-Order Theories of Consciousness

Lycan, William G. (2001). A simple argument for a higher-order representation theory of consciousness. *Analysis* 61 (269):3-4.

Lycan, William G. (2001). The case for phenomenal externalism. *Philosophical Perspectives* 15 (s15):17-35.

Lycan, William G. (2004). The superiority of HOT to HOP. In Rocco J. Gennaro (ed.), *Higher-Order Theories of Consciousness: An Anthology*. John Benjamins.

Lycan, William (2008) "Representational Theories of Consciousness", *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/consciousness-representational/>>.

Mandik, Pete (2009). Beware of the unicorn: Consciousness as being represented and other things that don't exist. *Journal of Consciousness Studies* 16 (1):5-36.

Matey, Jennifer (2011). Reduction and the determination of phenomenal character. *Philosophical Psychology* 24 (3):291-316.

Nagel, Tom (1974) What is it like to be a bat? *Philosophical Review* LXXXIII, 4: 435-50

- Pautz, Adam (2009). What are the contents of experiences? *Philosophical Quarterly* 59 (236):483-507.
- Prinz, Jesse (2012) *The Conscious Brain: How Attention Engenders Consciousness* Oxford University Press
- Rosenthal, David M. (2002). How many kinds of consciousness? *Consciousness and Cognition* 11 (4):653-665.
- Rosenthal (2005) *Consciousness and Mind*
- Rosenthal, David (2010). How to think about mental qualities. *Philosophical Issues* 20 (1):368-393.
- Rosenthal, David (2011). Exaggerated reports: reply to Block. *Analysis* 71 (3):431-437.
- Rosenthal, David and Weisberg, Josh (2008) Higher-order theories of consciousness. *Scholarpedia*, 3(5):4407
- Sebastian (forthcoming) Experiential Awareness: Do you prefer it to me? *Philosophical Topics* (40):2
- Tye, Michael (2000). *Consciousness, Color, and Content*. MIT Press.
- Weisberg, J. (2011) Misrepresenting Consciousness *Philosophical Studies* 154:409-433