

Against Irrealism

Nilanjan Das

University College London

nilanjan.das@ucl.ac.uk

Jan Westerhoff's *The Non-Existence of the Real World* is a beautifully written, rich and stimulating book. Its aim is to defend an ambitious form of *non-foundationalism*. Foundationalism is the conjunction of two claims: the claim that the world is a certain way fundamentally and objectively, and the claim that the aim of philosophy is to capture the way the world is fundamentally and objectively. The version of non-foundationalism that Westerhoff endorses involves denying both these claims. He denies that there are foundational truths about the world, and that philosophy can give us theories that capture such foundational truths. In this respect, arguably, he has as his allies Buddhist thinkers who endorse a similar strain of non-foundationalism, i.e., Mādhyamika thinkers like Nāgārjuna (150-250 CE).

For people like me (who are interested in contemporary metaphysics and epistemology as well as the history of Sanskrit philosophy), Westerhoff's book is especially instructive. On the one hand, it's a great example of how creatively engaging with the ideas and arguments of ancient South Asian philosophers can open up new avenues of thought. On the other hand, it shows that, for such creative engagement to occur, the philosophical work doesn't always have to be preceded by textual exegesis.

In the first chapter of his book, Westerhoff casts doubt on an assumption that many projects in ontology take for granted: the assumption that there exists an external world, i.e., a world of mind-independent objects and properties. Why should we accept this assumption? A simple answer is this. Our perceptual experiences can represent the properties of mind-independent objects. When the conditions are right, those perceptual experiences provide us evidence, or put us in a position to know, that those properties are instantiated. For example, when I undergo a perceptual experience as of there being a yellow, crescent-shaped banana before me, I am in a state that represents the yellow colour and the crescent shape of the banana as instantiated. When the conditions are right (i.e., my vision is functioning reliably), I am justified in believing, or in a position to know, that there is a yellow, crescent-shaped object before me.

This simple story, on Westerhoff's view, is wrong. He defends what he calls *irrealism*. This is a conjunction of two theses:

THE VIRTUAL WORLD THESIS. Our perceptual experiences only directly represent virtual objects and properties that are part of a virtual model or representational interface constructed by our brain.

THE NON-EXISTENCE THESIS. There are no mind-, language- or theory- independent objects that lie beyond this model or interface.

In what follows, I argue against Westerhoff's defence of irrealism. In §1, I consider whether we should accept the VIRTUAL WORLD THESIS given the evidence that Westerhoff presents in favour of it. In §2, I will raise broader concerns about the NON-EXISTENCE THESIS.

1. The Virtual World Thesis

While defending the VIRTUAL WORLD THESIS, Westerhoff appeals to two theories of perception: Donald Hoffman's interface theory, and Jakob Hohwy's prediction minimization theory. It's not clear to me that, in either case, we can easily get to the conclusion that our perceptual states don't directly represent a world of mind-independent objects.

1.1 Intentionalism

To fix ideas, I will begin with a form of *intentionalism*, the view that all perceptual experiences fundamentally consist in representational states, i.e., states with contents. To say that perceptual experiences fundamentally consist in such states is to say that the person-level features of these experiences, e.g., their phenomenal character, the role they play in justifying beliefs, or the way in which they guide behaviour, are explained by their being representational states. Unlike naïve realism, this view doesn't commit us to thinking that we are in a completely different kind of mental state when we are undergoing illusions or hallucinations from when we are undergoing a veridical perceptual experience. Unlike classical representationalism or indirect realism, intentionalism avoids the result that we don't directly represent mind-independent items through our perceptual experiences.

Some defenders of intentionalism claim that the contents of perceptual experiences consist in their accuracy conditions: on Siegel's (2010) view, for example, these accuracy conditions obtain in virtue of the fact that perceptual experiences represent clusters of properties as being instantiated. They are accurate just in case these properties are in fact instantiated. How do our perceptual experiences represent such clusters of properties?

According to causal-teleological theories of representation, *indication* is a key component of representation. The following analogy is helpful. Thermometers represent temperatures insofar as they are in states that have the function of indicating those temperatures. Similarly, organisms like us perceptually represent the properties of things by undergoing states that have the (biological) function of indicating them. Very roughly, a state stands in the indication relation to a property just in case it *normally* occurs only when that property is instantiated (Dretske 1995, p. 48). But indication isn't all that there is to perceptual representation; in order to perceptually represent a property, not only must the relevant organism be in a sensory state that indicates and thereby carries information about that property, but it must also be able to rationally exploit the information that is carried by those states. That is why some say that an organism perceptually represents a certain property just in case it undergoes an internal sensory state that (i) has the biological function of indicating that property and (ii) can be fed into a cognitive system for the rational control of thought and action (Tye 1995; Dretske 2006).

In the remainder of this section, I argue that intentionalists can explain the data that Hoffman and Hohwy present without rejecting the possibility that perceptual experiences directly represent mind-independent properties of objects in the external world.

2.2 The Interface Theory of Perception

Take Hoffman's view. Hoffman et al (2015) assume that, in order to accurately represent the world, there has to be some sort of structural resemblance between our perceptual representations (probably understood as mental images) and the world. Hoffman's central claim is that, given that our perceptual systems are a product of evolution, our perceptual states needn't accurately represent the world. If the predominant function of perceptual systems is to enhance adaptive fitness, they needn't represent the way the world really is. Hoffman shows this by running game-theoretic simulations where two different kinds of organisms are pitched against each other: the perceptual strategy of one organism allows it to track the quantity of a resource, while the perceptual strategy of the other doesn't track that. Quite predictably, the latter is driven to extinction. What does this show? Hoffman et al make the following observation:

We must take our perceptions seriously not because they reveal the true structure of the world, but because they are tuned, by natural selection, to fitness. The distinction between fitness and truth is elementary, and central to evolutionary theory. Fitness is a function of the objective world. However, a fitness function depends not just on the objective world but also on the organism, its state, and an action. For a hungry fly, a pile of dung conveys substantial fitness. For a hungry human, the same pile conveys no fitness.

Fitness is, in general, a complicated function of the objective world that depends on an organism, its state, and its action. There's no simple relation between fitness and truth, although many perceptual researchers assume otherwise. Geisler and Diehl (2002), for instance, assert "In general, (perceptual) estimates that are nearer the truth have greater utility than those that are wide of the mark." This would be convenient, but unfortunately it's not true. Fitness functions are more complex and versatile than that and rarely track truth. (Hoffman et al 2015, p. 1489)

On the basis of this, Hoffman concludes that we don't perceive things as they are: the objects we perceive are items that are part of a representational interface (like the icons on the desktop of a computer). They don't accurately represent the mind-independent world. Rather, they simply make it easier for us to successfully navigate the world.

It's worth noting that, intuitively, Hoffman's theory seems incompatible with the NON-EXISTENCE THESIS, the thesis that there are no mind-independent objects and properties. Since Hoffman's view appeals to the evolutionary history of perceptual systems, it presupposes the existence of a mind-independent world; as Hoffman himself points out, the

fitness-rewarding properties that our perceptual experiences track are “a function of the objective world.” So, let’s consider whether Hoffman’s theory supports the weaker claim encapsulated in the VIRTUAL WORLD THESIS: namely, that our perceptual experiences don’t directly represent properties of mind-independent objects.

We can agree with Hoffman that our perceptual experiences don't represent the structure of the world as it is independently of us. But that doesn't tell us why they cannot *accurately* represent mind-independent properties of objects in our environment. Suppose our perceptual experiences only selectively represent certain fitness-rewarding properties of mind-independent items. Since these fitness-properties ultimately depend on biological relations between us as organisms and our environment, these properties may only be relational properties that depend on us.

However, it remains compatible with this story that our sensory states indicate mind-independent properties of objects in our environment. Under Hoffman’s own account, whether certain fitness-rewarding properties are instantiated by objects in our environment depends partly on states of the objective world, e.g., on how certain resources are distributed in the external world. So, suppose that, normally, certain fitness-rewarding properties are instantiated only when certain mind-independent properties of objects in the external world are instantiated. Then, according to the causal-teleological story about representation, our sensory states can indeed have the function of indicating such mind-independent properties in virtue of indicating fitness-rewarding properties. Thus, if the form of intentionalism sketched above is right, nothing in Hoffman’s account rules out the possibility that we can perceptually represent mind-independent properties of objects in our environment.

The result is simple. We can agree with Hoffman that our perceptual experiences don’t simply represent the structure of the world as it is independently of us. But that doesn’t entail that they cannot directly represent mind-independent properties of objects in the external world.

2.2 The Prediction Error Minimization Theory of Perception

Hohwy's (2013, 2016) prediction error minimization (PEM) theory of mind and cognition may be more promising in this respect. The basic idea is this. When the brain receives sensory inputs, it uses those inputs to predict what kind of input it is going to receive in the future. When it receives the future inputs, there is often a difference between these and the predictions. So, the brain adjusts its predictions to minimise this error. The defenders of the PEM theory try to explain all aspects of mind and cognition in terms of prediction error minimization. An important claim of this theory is that the states of our biological sensory systems determine a Markov blanket.¹ Controversially, Hohwy thinks that the Markov

¹ What is a Markov blanket? Consider a complex system composed of many interacting nodes (or variables). A Markov blanket is a set of nodes such that, for some given node X, the state of X could be predicted given the states of the other nodes in that set. Thus, these other nodes form a ‘blanket’ that probabilistically screens off node X from the states of further nodes in the system.

blanket determined by our sensory states constitutes the boundaries of the mind,² and takes this to support a form of scepticism about the external world. If the mind is bounded by the sensory states in this way, it “will not be able to distinguish between possibilities where similar flows of sensory input are caused by two very different causal processes, beyond the blanket” (Hohwy 2017, §3).

Suppose we grant Hohwy's claim that our internal sensory states form a Markov blanket. Why does that entail anything about the boundaries of the mind? Presumably, Hohwy thinks that wide states of agents (i.e., states that depend on the extrinsic properties of an agent, e.g., her relations with her environment) play no role in psychological explanations (Hohwy 2016, pp. 274ff). Against externalists about mental states (i.e., people who think that mental states depend on extrinsic properties of an agent, not just intrinsic ones), Hohwy claims that, for the purposes of causally explaining the behaviour of our cognitive systems or us as agents, we don't need to appeal to external or bodily states beyond the Markov blanket created by our sensory states.

Notice that this is a version of an old argument given by internalists about mental states. The argument depends on two premises. First, we can causally explain actions like grabbing the phone without appealing to wide states; that explanatory work can be done simply by intrinsic properties of an agent (e.g., the internal representations of the phone and so on) in conjunction with external conditions (e.g., the accuracy of those representations). Second, if a state is genuinely mental, it must play a role in psychological explanations of action. Therefore, wide states cannot be regarded as mental states.³ Arguments of this sort can be resisted in two ways. One thing to say is that, even though only local or intrinsic properties of an agent may be relevant to psychological explanations, we should still individuate the contents of mental states in a way that makes them depend on extrinsic properties.⁴ The other move is to say that certain kinds of action can't be causally explained without appealing to wide mental states.⁵

I want to offer a different response. First, it's controversial whether anything in the PEM framework actually entails that we only encounter a world of internal representations in perception. In interpreting the framework, we don't necessarily have to construe the predictions made by the brain about future sensory inputs as internal representations. We can think of them (insofar as they determine probability distributions) as biases that skew the manner in which perceptual processes work.⁶ Second, by Hohwy's own admission, nothing in his account rules out the possibility that our internal sensory states can in fact indicate, or carry information about, properties of objects in the external world.⁷ So, if we accept

² For resistance to this claim, see Clark (2017a, 2017b) and Fabry (2017).

³ See, e.g., McGinn (1982, pp. 76-77) and Fodor (1987, p. 30).

⁴ See, e.g., Burge (1995).

⁵ See, e.g., Williamson (2000, pp. 61-64).

⁶ See, especially, Orlandi (2018).

⁷ See, for example, Kiefer and Hohwy (2018, p. 2393). Hohwy assumes that, while we are only directly encountering an internal model of reality when we undergo perceptual experiences, that internal model can

intentionalism, we can preserve the possibility that we can perceptually represent (in a direct manner) the mind-independent properties of objects in the external world.

The crucial consequence is this: if our internal sensory states can indicate mind-independent properties of items in the external world, then those indication relations can explain how we gain knowledge about the external world. No one denies that the internal sensory states of an agent (understood here as the local or intrinsic properties of that agent) can be replicated exactly in a sceptical scenario, e.g., a situation where the agent is merely a congenitally envatted brain or is being deceived by a Cartesian demon. But this is compatible with the possibility that, in some situations, those very sensory states can indicate mind-independent properties of objects in the external world. In a situation where those sensory states *reliably* or *safely* indicate that those properties are instantiated, the agent's sensory states may yield knowledge about them (in virtue of being rationally exploitable for the purposes of belief-formation). Thus, nothing in Hohwy's account shows that we cannot gain knowledge about the external world. So, his claim that the PEM theory supports scepticism isn't correct: as long as we accept a sufficiently externalist conception of knowledge, even though our epistemic situation may be subjectively or phenomenally indistinguishable from sceptical scenarios, we can still be in a position to know that we are not in those sceptical scenarios.

The lesson is this. The VIRTUAL WORLD THESIS isn't the only hypothesis that explains the evidence that Hoffman or Hohwy present. There are intentionalist theories that can explain the same data without rejecting possibility that our perceptual experiences directly represent mind-independent objects or their properties. So, the evidence presented by Westerhoff in favour of the VIRTUAL WORLD THESIS isn't decisive.

2. The Non-Existence Thesis

For argument's sake, let's grant that Westerhoff's arguments for the VIRTUAL WORLD THESIS succeed. What reason do we have for accepting the NON-EXISTENCE THESIS? In this section, I argue for two things. First, Westerhoff's argument for the NON-EXISTENCE THESIS depends on a questionable principle of theory choice. Second, if this thesis is true, irrealism itself is either false or unrepresentable or both.

2.1 Two Principles of Theory Choice

Let's assume that classical representationalism is true: our perceptual experiences only directly represent mind-dependent items and their properties, e.g., sense-data, items within a virtual model of reality, and so on. Such views can be of two kinds: realist and irrealist. Realist views entail that even though we don't directly perceptually represent mind-independent items or their properties, we can indirectly represent them (in virtue of certain causal and resemblance relations). Irrealist views deny this. Since the irrealist is committed to the NON-

accurately or inaccurately represent a mind-independent world. Kiefer and Hohwy (2018) explain this notion of accuracy in terms of structural resemblance.

EXISTENCE THESIS, she doesn't think we can perceptually represent (either directly or indirectly) any mind-independent object or its properties.

Westerhoff argues that irrealist theories are preferable to realist ones, because the former are simpler (Westerhoff 2020, p. 61). In doing so, presumably, Westerhoff endorses a relatively strong version of Occam's razor: the principle that, if there are two theories that explain a certain body of evidence, then, given that body of evidence, we should accept (or prefer to accept) the theory that posits fewer kinds of entities. This is a strong principle: it doesn't merely say that simplicity is one theoretical virtue amongst others; rather, it says that, when we face a choice two theories that explain the relevant body of evidence, considerations of simplicity take precedence over other considerations.

Is this principle correct? Contrast this with a different principle about (propositional) justification that we may call *phenomenal conservatism*: the principle that, if things appear to us to be a certain way, then, in the absence of defeaters against those appearances, we should believe that things are that way. Phenomenal conservatism predicts that we should accept or believe theories that predict the appearances to be true, rather than theories that don't (provided that there are no defeaters against those appearances). Phenomenal conservatism (arguably) has many virtues: it helps us resist scepticism about both the external world and moral knowledge; it provides a simple and unifying theory of justification for a range of different beliefs; it is the only theory of justification that isn't self-defeating.⁸

The strong principle of simplicity and phenomenal conservatism pull us in opposite directions when we face a choice between an irrealist theory of perception and a realist one. Suppose there are two incompatible theories of perception, T1 and T2, both of which explain the available evidence, i.e., the phenomenology of our perceptual and cognitive states. T1 is an irrealist theory that doesn't posit any mind-independent items. T2 is indirect realist theory that posits such objects and predicts that, at least under normal conditions, the external world resembles the way things appear in perception. While T1 posits fewer kinds of entities than T2, T2 is better at predicting the appearances to be true. If we take simplicity to be a decisive tiebreaker, then we should choose T1; if phenomenal conservatism is the true theory of justification, then we should believe T2. So, the challenge is this. If Westerhoff accepts the strong principle of simplicity stated above, he must tell us why we should reject a theory of justification like phenomenal conservatism.

We can state this challenge in a much more general manner, independently of any commitment to phenomenal conservatism. The most plausible version of Occam's razor says that, other things equal, simpler theories are preferable to more complex ones. Take two theories of perception, a realist one and an irrealist one, both of which explain the phenomenology of perceptual and cognitive states. Even though the irrealist theory is simpler, we cannot directly apply Occam's razor to this case, because other things are not equal between these theories. The realist theory is better at preserving the truth of ordinary

⁸ For arguments for this view, see Huemer (2001, ch. 5; 2005, ch. 5; 2007).

appearances. Therefore, if someone thinks that considerations of simplicity give us an all-things-considered reason to prefer the irrealist theory to the realist one, that person owes us an explanation of why considerations of simplicity should outweigh considerations about preserving the truth of ordinary appearances.⁹

2.2 Is Irrealism Unrepresentable?

This brings me to my last claim: namely, that irrealism is either false or unrepresentable or both. My argument depends on a causal constraint on representation (Putnam 1981): namely, a linguistic expression or a concept, used by a certain agent, can refer to or represent a particular or a kind only if that agent is causally connected to that particular or kind, or other particulars or kinds that could be used to descriptively refer or represent to the relevant particular or kind. If this constraint is correct, then, in a radical sceptical scenario, the agent who finds herself in that scenario cannot represent the scenario she is in. Consider the brain in a vat: since the congenitally envatted brain is causally insulated from the external world, when it uses the expression “brain” or the corresponding concept, it cannot refer to real brains and therefore cannot linguistically or conceptually represent its own predicament, i.e., that it is a brain in a vat. Here, I show that an analogous result follows for the irrealist.

Our starting point is an objection that the irrealist faces. Westerhoff notes that, though the irrealist can assent to the truth of the statement, “There is an external world beyond my representation,” she must understand this statement quite differently from the direct or indirect realist.

Note, also, that the irrealist can assent (together with the direct realist and the representationalist) to the truth of the statement ‘there exists an external world beyond my representation’...An irrealist who equates physical objects with parts of the representational interface does not deny the reality of physical objects, any more than an economist who equates risks with potentials for loss denies the reality of risks. (Westerhoff 2020, p. 56)

Here’s one way of understanding the idea. If we are genuinely trapped in a virtual world that is causally insulated from an external world, then our talk or thoughts about external objects or properties can at best represent elements of the virtual model of reality that we have constructed. Since that virtual model includes the physical objects that we can think or talk about, our belief or statement that there exists an external world beyond our representations can be true relative to that model.

⁹ The challenge remains intact even if we accept other theories of propositional justification. For example, on a Bayesian theory, the degree to which an agent should be confident in a hypothesis H in light of a piece of evidence E is simply her prior rational conditional credence in H given E, which, in turn, is the ratio of her prior rational credence in H&E and her prior rational credence in E. Relative to this framework, the challenge for someone like Westerhoff is to explain why the rational constraints on priors are biased in favour of parsimonious theories rather than appearance-preserving ones.

A worry about this view is that, by including the world of physical objects within the representational interface, it blurs the distinction between irrealism and direct realism. Westerhoff is sensitive to this worry.

We might wonder whether the irrealist position, claiming that all external objects are simply part of the representation does not reduce to a form of direct realism. For the irrealist everything that was behind the veil of perception is now in front of it, and so directly, and no longer indirectly epistemically accessible. As the veil plays no theoretical role once there is nothing behind it any more, we may as well dispose of it, and are then left with a veil-less theory of perception in which everything is perceived directly. (Westerhoff 2020, p. 68)

In response to this challenge, Westerhoff points out that direct realists and irrealists disagree on what we have epistemically direct access to (Westerhoff 2020, p. 69). Direct realists are committed to the idea that, through our perceptual experiences, we directly represent objects that are mind- or language- or theory-independent. By contrast, irrealists deny this. On their view, people like us are simply wrong to believe that they have epistemically direct access to objects that are mind- or language- or theory-independent. My worry is that Westerhoff cannot help himself to this way of distinguishing irrealism from direct realism.

Let's grant that we are mistaken in holding the belief that we have epistemically direct access to mind-, language-, and theory-independent external objects. This means that we can represent objects as mind-, language-, and theory-independent, i.e., as objects whose existence and nature doesn't depend on our mental states, linguistic practices, and accepted theories. To represent objects in this way, we should at least have the capacity to represent the relevant dependence relation (or relations); call it R . If the causal constraint on representation is correct, then we can linguistically or conceptually represent this relation only if we are causally connected to it or have the representational resources (in virtue of other causal connections) to represent it descriptively. But, if we are not causally connected to anything that lies beyond the representational interface, then the relation R has to be a part of the virtual model of reality that we have constructed.

There are two problems. The first is this. *Ex hypothesi*, it's true relative to that virtual model that there are physical objects which lie beyond our representations and which don't depend by relation R on our representations. So, provided that we are accurately representing how things are within the virtual model, we simply cannot *falsely* believe that we are epistemically directly connected to mind-, language- and theory-independent objects. So, if the irrealist is committed to the view that ordinary people are mistaken about what they are epistemically directly connected to, she is wrong. A related, but deeper, problem here is the same as the one that arose in the case of the congenitally envatted brain. If irrealism is true, we as theorists don't have resources to represent a world of objects whose nature or existence is genuinely independent of our representational interface. So, it is impossible for us to occupy an Archimedean point outside of that representational interface and then, from that perspective, assess whether people's ordinary experiences or beliefs about what they're

epistemically directly connected to are mistaken. Thus, irrealism, if true, becomes unrepresentable.

I anticipate two possible responses from Westerhoff. First, he could argue that, in order to say that people are mistaken in their beliefs about what they have epistemically direct access to, we only have to assume that they have a concept of mind-, language-, and theory-independent objects. We don't have to commit ourselves to the existence of such objects or to any causal connection between ourselves and those objects (Westerhoff 2020, p. 305).

I am not sure this move succeeds. We have the conceptual resources to think about fictional objects like Sherlock Holmes. But that is only because we have more basic concepts which we acquire by causally interacting with a world of mind-independent objects and which allow us to construct the more complex fiction of which Sherlock Holmes is part. But it's not obvious that the more expansive fiction as of there being mind-, language-, and theory-independent objects is like this. Suppose all our primitive concepts represent elements of the virtual world we live in. Using those concepts, can we start thinking about a (fictional) world that lies completely beyond that virtual world? As I pointed out, if the concept of the dependence relation R (using which we represent mind-, language-, or theory-dependence) picks out some relation within the virtual model, it's unclear how we could use that concept to represent objects that lie completely outside of the representational interface. Here, we need some reason to think that such transcendental fictions are possible.

Another possible response is to resort to a Wittgensteinian saying/showing distinction. Westerhoff could say that, while we have no way of *stating* the falsity of people's belief that they have epistemically direct access to mind-independent objects, we can *show* that they are false.¹⁰ But I am not sure that this distinction alone can help. The problem isn't just that we can't say that those beliefs are false, but also that they can't be false. Given that people can only form beliefs about the elements of the virtual model that they directly represent, their representations of physical objects as mind-independent will be true relative to that virtual model. So, when they think that they are epistemically directly connected to mind-independent objects, they won't be wrong.

So far, I have argued that, given a familiar causal constraint on representation, irrealism is either false or unrepresentable or both. In reply, Westerhoff might reject the causal constraint I have been working with, or say that the relevant causal relation is itself a part of the virtual model of reality. These strategies are not going to help. If Westerhoff rejects the causal constraint, he will need an alternative account of representation that allows him to explain how our ordinary beliefs about what we have epistemically direct access to can be false, and how irrealism can be represented. Similarly, if Westerhoff says that the causal relation in question is only a relation between objects that are part of the virtual model of reality, then we should only be able to represent objects that are part of the virtual model of reality. Then, my reasoning above will apply again: the irrealist won't have the

¹⁰ This strategy is also discussed by Westerhoff (2020, p. 288, pp. 297-8) in a different context.

representational resources to represent their own thesis, and it won't be easy for her to show that our ordinary beliefs about what we are epistemically directly connected to are false.

3. Conclusion

Let's take stock. Here, I have raised concerns about the two theses that the irrealist endorses: the VIRTUAL WORLD THESIS and the NON-EXISTENCE THESIS. I have argued that the evidence that Westerhoff presents for the VIRTUAL WORLD THESIS doesn't adequately support it. Then, I have registered a worry about the principle of theory choice that Westerhoff takes to favour the NON-EXISTENCE THESIS. Finally, I have shown that, given a plausible constraint on representation or intentionality, irrealism is either false or unrepresentable or both.

In making these arguments, I haven't scratched the surface of the expansive vision of ontology and metaphysics that Westerhoff presents in his book. My only hope is that these remarks will be useful in assessing Westerhoff's view. For those who are already converted to the irrealist cause, these might help to clarify what's at stake in accepting irrealism. For others who are less friendly to irrealism, these remarks might show a way of resisting it.

References

Burge, Tyler. 1995. Intentional Properties and Causation. In: Macdonald, C., & Macdonald, Graham F. (eds), *Philosophy of Psychology: Debates About Psychological Explanation*. Blackwell.

Clark, Andy. 2017a. Busting Out: Predictive Brains, Embodied Minds, and the Puzzle of the Evidentiary Veil. *Noûs*, 51(4), 727–753.

Clark, Andy. 2017b. How to Knit Your Own Markov Blanket. In: Metzinger, Thomas, & Wiese, Wanja (eds), *Philosophy and Predictive Processing*.

Dretske, Fred. 1995. *Naturalizing the Mind*. MIT Press.

Dretske, Fred. 2006. Perception Without Awareness. Pages 147–180 of: Gendler, Tamar S., & Hawthorne, John (eds), *Perceptual Experience*. Oxford University Press.

Fabry, Regina E. 2017. Transcending the evidentiary boundary: Prediction error minimization, embodied interaction, and explanatory pluralism. *Philosophical Psychology*, 30(4), 395–414.

Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press.

Hoffman, Donald. 2019. *The case against reality: Why evolution hid the truth from our eyes*. WW Norton & Company.

Hoffman, Donald D, Singh, Manish, & Prakash, Chetan. 2015. The interface theory of perception. *Psychonomic bulletin & review*, 22(6), 1480–1506.

Hohwy, Jakob. 2013. *The Predictive Mind*. Oxford University Press UK.

- Hohwy, Jakob. 2016. The Self-Evidencing Brain. *i*, 50(2), 259–285.
- Hohwy, Jakob. 2017. How to Entrain Your Evil Demon. In: Metzinger, Thomas, & Wiese, Wanja (eds), *Philosophy and Predictive Processing*.
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.
- Huemer, Michael. 2005. *Ethical Intuitionism*. New York: Palgrave Macmillan.
- Huemer, Michael. 2007. “Compassionate Phenomenal Conservatism.” *Philosophy and Phenomenological Research* 74, 30-55.
- Kiefer, Alex, & Hohwy, Jakob. 2018. Content and Misrepresentation in Hierarchical Generative Models. *Synthese*, 195(6), 2387–2415.
- McGinn, Colin. 1982. *The Character of Mind*. Oxford University Press.
- Metzinger, Thomas. 2009. *The Ego Tunnel: The Science of the Mind and the Myth of the Self*. Basic Books.
- Orlandi, Nico. 2018. Predictive Perceptual Systems. *Synthese*, 195(6), 2367–2386.
- Putnam, Hilary. 1981. *Reason, Truth and History*. Cambridge University Press.
- Siegel, Susanna. 2010. Do Visual Experiences Have Contents? In: Nanay, Bence (ed), *Perceiving the World*. Oxford University Press.
- Tye, Michael. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.
- Westerhoff, Jan. 2020. *The Non-Existence of the Real World*. Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.