Dear Max,
Happy New Year! I now managed to go through the versions of the 'Introduction' and the 'Response' you sent me. I have incorporated all your suggestions in terms of typos etc, and have added some dates and contextualizations for the Indian names and terms. I also cut just over 1,000 words from the last version of the 'Introduction', *inter alia* by drastically shortening section 8, as you suggested. I hope the result present a more succinct response, while still covering all the essential points.

Let me know if you need anything else from me at this stage.

All best

Jan

Cut section 8 from 1012 to 356 = 656 words

Cut "A smaller […]" 113 words.

Cut "Nor is […]" 77 words.

Cut most of "Das notes that there might be […]" 172

Cut "An important principle Das's […]" by 46

1064

This version contains the sections (highlighted in yellow) that were cut or severely abbreviated from the published version to conform to the journal's word limit.)

The non-existence of the real world

Jan Westerhoff

Response to critics

Revised version January 2022

I was intrigued to see that there was surprisingly little overlap in the responses by the three critics. It therefore seemed most sensible to treat each critic's points separately, which is what I have done below. I hope I have managed to say something substantial on each of the points they brought up. While I harbour no illusions about my comments resolving all of the concerns raised to each critic's satisfaction, I hope that the response helps to elucidate unclarities in the earlier

exposition of my ideas and may, perhaps, provide resources for the further philosophical exploration of the views discussed.

Nilanjan Das

1. Prediction-error minimization

In relation to my discussion of Hohwy's prediction-error minimization framework Das rightly points out that nothing in this approach entails that our representations cannot carry information, and hence yield knowledge about the external world. Indeed, some interpreters see this approach as a way of discovering the "hidden causes" of whatever it is that brings about the changes in the outer nodes of our Markov blanket, arguing that

> the generative model 'recapitulates' the causal-probabilistic structure of the external environment that impinges on the organism's sensory apparatus. It constantly generates, in a top-down manner, a flow of virtual or mock sensory signals that predicts the unfolding of sensory signals by external causes.[1]

Similarly, some forms of Bayesianism see it as a way of discovering the true probabilities out there in the world (the 'Lord's prior'). What I am trying to argue is that neither appealing to the *existence* of some external item represented nor appealing to the *accuracy* of our representation (both the prediction-error minimizer and the Bayesian are likely to say that there are some structural features of the 'hidden causes' or the 'true probabilities' that our beliefs will capture) is necessary for the prediction error minimization theory; in fact that the theory would be in better shape if it dispensed with this idea. Considerations of parsimony aside, it is hard to make sense of the kind of connection supposed to exist between things inside and outside of the Markov blanket. On the one hand such a relation should be inside the Markov blanket for us to be able to think about it. Yet it also needs to reach beyond it in order to link up with the 'hidden causes'. But if we can conceptualize a relationship and one relatum, the status of the other relatum does not appear to be quite as inaccessible as the idea of the explanatory-evidentiary boundary makes out. As such, the notion of a 'connection' with something outside of the Markov blanket is something we should do without.[2]

2. Hoffman

Das notes that "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". According to Das's intentionalism, we represent a certain property in our environment if we are in an internal state that "has the biological function of indicating that property". Certain internal states of an organism might have the function of indicating the presence of

[1] Gładziejewski 2016: 562.

[2] For related worries about the idea of probabilities as objective characteristics of the world from a Bayesian perspective see Feldman 2013.

water. Does that mean that our internal state represents the mind-independent property of there being water in our environment? Not according to Hoffman's account. What Hoffman sets out to show is that our perceptual capacities are not geared towards indicating properties the environment has no matter what, but that they have been shaped by evolution to indicate instances of fitness arising from an organism *in* an environment. An organism's internal states do not have the function of indicating properties of the environment, but properties of the environment refracted through the lens of the organism's specific interests and needs, that is, instances of fitness. In these instances of fitness, properties of the organism and properties of the environment are so much bound up with one another that it is impossible to separate them. A realist critic might point out here that there must still *be* mind-independent properties of the environment, whether or not we can disentangle them from our representations or not. This move, however, pushes these properties into an epistemically inaccessible noumenal realm, and it becomes then unclear why the resulting position has any advantage over a more parsimonious rival who simply treats the mind-independent world as another part of the representational interface.

In addition, an argument involving self-information entailed by non-self locating claims is presented by Das, suggesting as an example that being in the British Library entails the existence of the British Library. I don't think Hoffman would disagree with this, any more than he would disagree with saying that the claim that there is a blue folder on my desktop entails that there is something which is blue. However, this does not yet get us to "properties of mind-independent particulars", but only properties (and individuals) in the interface. Hoffman's interface theory of perception does not rule out "the possibility of there being self-locating truths or of self-locating truths entailing non-self-locating truths", but claims that any such truths are strictly truths about what goes on within the representational interface.

3. Externalism

In addition, Das points out that "as long as we accept a sufficiently externalist conception of knowledge […] we can still be in a position to know that we are not in those sceptical scenarios" irrealism raises. Externalism about mental states here means that such states "depend on extrinsic properties of an agent, not just intrinsic ones". As an extrinsic property of me is one that essentially involves entities other than myself, externalism about knowledge clearly avoids the irrealist predicament. But anybody inclined towards irrealism is unlikely to assume that cognitive states essentially involve entities beyond the cognitive interface. The regarding impasse is then prone to reduce to a debate about where the burden of proof should lie: does the irrealist have to show that the existence of a mind-independent world is *incompatible* with the epistemic situation we find ourselves in (as Das seems to think), or does the realist have to show that the assumption of such a world presents an explanatory advantage that the irrealist's representation-only view cannot match?

It seems that one of the reasons why the debate between the realist and his opponent is so difficult to resolve is that the two have different ideas of what such a resolution would look like. The realist demands that his opponent 'rules out' the possibility of a

mind-independent external world, while the opponent wants to see evidence of a unique feature of the realist conception of the world that cannot be replicated by 'realism in the image', that is, by the idea that the mind-independent world is part of the representational interface. The irrealist regards the realist's demand as excessive, as no amount of epistemic insulation would prevent it from *just happening to be case* that there was the realist's kind of world behind the representational interface (much like some events in the real world might *just happen* to correspond to events in some video game). The realist regards the irrealist's offer of 'exact replication' of his external world as insufficient, since his epistemic default is always the existence of something behind the representation. As long as we don't have proof that there *cannot* be such a thing we are justified in assuming that there is.

I agree with Das that even if we accept that perceptual experiences do not represent the world as it is independently of us, this does not entail that these experiences cannot be "true or accurate". The world behind appearances might indeed be structurally identical with our perceptions *by accident*. To the great surprise of the Kantian it might turn out that the structure of the noumenon just happens to be that of Euclidean geometry. But basing one's metaphysics on such a conjecture is hardly more satisfactory than basing one's finances on the non-zero chance of winning the lottery tomorrow. Unfortunately, the chances of the latter are extremely small. Hoffmann presents a corresponding probabilistic argument that the chances for a substantial structural similarity between representation and represented are even smaller.[3] If realism is as robust a theory as its defenders consider it to be, I believe that there must be a stronger argument to support it than the claim that, appearances to the contrary, the world might just accidentally turn out to contain a entities with the same structure as our perceptual representations.

4. Phenomenal conservatism

Phenomenal conservatism plays an important role in Das's argument. It is the claim 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". In the Indian philosophical arena this claim was endorsed by the Mīmāṃsā philosophers and opposed by Buddhist epistemologists,[4] so we have another instance here where the exchange between the irrealist and his critics traces some familiar philosophical territory. But how are we going to assess the truth of phenomenal conservatism in a systematic manner? One point that remains unclear to me is how much pressure the proviso "in the absence of defeaters" is supposed to sustain. If I look at the world from an unreflective armchair position it appears to me, and probably to most people, that there are real colours out there in the world, that the structure of space is Euclidean, that introspection reveals all aspects of my inner life to me with irrefutable authority, that it is an absolute fact whether two events happen or don't happen at the same time, that my visual field has the same resolution everywhere,

[3] Prakash et al 2020.

[4] In this debate between non-Buddhists and Buddhists, pursued vigorously at the time of Dharmakīrti (6th ct CE) we see an epistemic optimism, trusting our epistemic powers in all normal circumstances facing an epistemic pessimism, claiming that our regular epistemic access to the world is fundamentally flawed. For a brief account see Westerhoff 2018: 259-263.

and so on. Yet even some modest exposure to cognitive science, theories of perception, and relativistic physics shows that these appearances do not sustain the belief "that things are that way". Revised versions of all of these can of course be incorporated into our belief-system, but at that stage we appear to have lost much of what was initially appealing to phenomenal conservatism, namely the claim that the world you see is more or less the world you get. If all things considered the world as it appears and the world as our best theories describes it are radically different, what precisely is phenomenal conservatism trying to conserve? If the examples of the theory of relativity, quantum physics, or cognitive science are any guide it appears as if the deeper we dig, the stranger matters become. Even to committed realists, the structure of the world 'as it is in itself', to the extent to which we determine this on a scientific basis, is very different from the way things appear under 'normal conditions'. Whether or not this refutes phenomenal conservatism is a matter we need not settle here, but it seems to be at least clear that phenomenal conservatism needs to be backed up by some kind of argument. It is not a methodological maxim we can accept without some further idea of what might support it.

5. Parsimony

Another worry concerns the suspicion that considerations of parsimony may not unequivocally support irrealism. While it is qualitatively simpler than realism, as irrealism only requires one kind of entiy (representations), whereas realism requires two (representation and represented), realism has the advantage of "preserving the truth of ordinary appearances". Apart from the question whether ordinary appearances are indeed worth preserving *as a guide to how things are at the most fundamental level* (rather than simply as pragmatically useful heuristics employed in navigating our daily life) considerations of parsimony are not the only big-picture idea that speak in irrealism's favour. Another is the point that what we mean by 'the world' is saturated with concepts like causation, time, space, logical implication, probability physical, mental, abstract, concrete and so on (some of which essentially rely on another) to such an extent that it is unclear how much sense we can make of a world 'in itself' wholly beyond these concepts. This is probably less of a concern for Das, whose phenomenal conservatism presumably commits him to the claim that causation, space, time, and so on are all somehow 'out there' in the world, and more or less successfully mirrored by our concepts. But those not immune to the Kantian worries that many of the concepts our idea of 'the world' is built around might be no more than shadows thrown by the mind will not only believe that a theory that dispenses with the noumenal realm might be more parsimonious, but are likely to be unsure how such an entity denuded of properties corresponding to notions essential to our cognitive lives could be what the reality we live in is at the most fundamental level. As a consequence, once we question the existence of any substantial structural correspondence between our representations and the mind-independent world, and are unwilling to accept an undefined noumenal 'something' as the ground of all appearances, treating the world behind the representational framework as part of the representational framework seems to remain as the most sensible option.

6. We don't need evidence for the external world

An important principle Das's argument relies on is the assumption that in the absence of defeating evidence, we can simply assume the existence of an external world. We do not need to adduce any positive evidence for its existence. While at this point the argument can degenerate into simple burden-of-proof shifting, the reason I am somewhat wary of the principle that if there is no evidence defeating the existence of x, we are rationally permitted to take the existence of x for granted, is that it would allow us to postulate the existence of *anything* beyond the particle horizon. In this matter my philosophical intuition is precisely the opposite of the one proposed by Das: any hypostatization has to earn its keep in explanatory terms. For entities that do not (like the idea of a noumenal world we cannot interact with in any way), the response should not be agnosticism ("perhaps there really are moons made of blue cheese behind the particle horizon, or perhaps there aren't"), but atheism: explanatory idle wheels do not have doubtful ontological status, but no ontological status.

7. Irrealism as unrepresentable

Das raises the difficulty that the irrealist's denial to accept anything beyond the representationalist interface might be unrepresentable by the irrealist's own standards. According to him, the irrealist literally cannot say what he is setting out to say. The worry is that the irrealist cannot understand the expression "beyond the representationalist interface"as referring to something beyond the interface, since for him there is no such thing. As such any sentence containing this expression would be either false or meaningless. Alternatively, he could regard the expression as standing for something within the interface, such as the interface-bound representation of something beyond the interface. But now the irrealist and his opponent seem to be talking about different things. When the opponent talks about things beyond the interface, he means mind-independent things causally interacting with him in some way. As such nothing the irrealist can say about such things, now interpreted according to his interface-bound semantics could even disagree with what his opponent is saying.

In order to resolve this issue it is important to be clear that *if* irrealism is accepted, any semantic theory that essentially relies on things beyond the interface needs to be abandoned. For this reason we cannot make such a theory part of our deliberations when trying to decide between realism and irrealism. Das is correct in claiming that for the irrealist "is impossible for us to occupy an Archimedean point outside of that representational interface", but that does not imply that the representability of irrealism depends on such an Archimedean point. The question we need to ask is: "If irrealism turn out to be true, is it representable according to the irrealist's own semantics?" The answer to this appears to be a clear 'yes',[5] for the irrealist has the resources to express that specific concepts are not instantiated (he can express that

[5] Of course the opponent could argue that you that it might *appear* to you as if you believed P (where P may be "there are no entities beyond the perceptual interface"), but that you do not really believe P, because the causal constraint is not satisfied. But if the realist's answer to the skeptic implies a skeptical scenario far more extreme than the original one (now we not only have to worry that there is no orange out there to correspond to our perception of an orange, we also have to worry that our belief "the orange is next to the apple" might not mean what we think it means) I am not sure whether the cure is not worse than the disease.

there are no yetis, for example), and so he can express that the concept 'entity beyond the representational interface' is not instantiated.[6]

8. The causal constraint

Das raises the worry that if we make the plausible assumption that successful linguistic reference depends on causal interaction with entities referred to[7] the irrealist seems to be unable to refer to ordinary entities like tables and chairs, since he has never interacted with them.

When speak about the non-mathematical realm such a causal constraint is indeed plausible, since it grounds our successful reference in causal contact with the entities referred to. The disagreement between the irrealist and his opponent is here not about the causal constraint as such, but how to interpret its reference to causality. It is clear that when Das speaks of "causal interaction" in his formulation of the causal constraint he takes this to refer to *real* causation, to a relation that connects us and our representational states with entities 'out there'. This, however, is an assumption that will appear unwarranted to the irrealist. If we are already assured that there is at least one relation that functions as a bridge between the representational inside and the outside of the *representanda*, the irrealist's claim that all our conceptual resources only manage to latch on to something within the representational framework is obviously false. Yet there is nothing stopping the irrealist from accepting the causal constraint as long as we understand it to speak about causation-in-the-interface.

The necessity of treating "causation" as referring to something within the representational interface rather than beyond it is hardly an original point. It is simply a version of Putnam's 'just more theory' move in the debate about the model-theoretic argument.[8] We cannot zero in on the intended interpretation of our theory by bringing in causation since the reference of terms like 'causation', 'to cause' and the like is affected by the same indeterminacy affecting other terms. The causal constraint is a principle that operates within the confines of the interface, not a bridge principle linking parts of the interface with entities that are not part of the interface.

[6] Das notes that there might be the remaining worry that we can no longer distinguish irrealism from naïve realism, since the naïve realist and the irrealist agree that there is no veil of perception. However, the two do not agree in all other respects (for example on the question of whether there would be observed objects without observers).

[7] Or causal interaction with other entities we can then use to descriptively refer to entities we have not causally interacted with.

[8] See Button 2013: 27-31.

Allison Aitken

1. Irrealism, nihilism, atheism

Aitken notes that "Westerhoff clarifies that irrealism is not a nihilism but an atheism about the external world." This might appear puzzling to some. The reason that I distinguish irrealism from nihilism is because in the absence of ultimate truths it cannot be an ultimate truth that there is nothing, as nihilism would have it. The alignment of irrealism with atheism is based on the idea that we can have an agnostic stance towards the external world ("maybe there is such a world, maybe not"), or an atheist stance ("there is no such thing"). Irrealism does not claim that we lack enough information to settle this question, but that on the basis of criteria accessible to us the more plausible position is to deny the existence of an external world. Nevertheless, atheism (theologically or ontologically conceived) is a view of how the world is at the fundamental level, and irrealism is not. This is where the parallel between irrealism and atheism ends.

2. The brain as self-grounding

Another worry concerns the question how precisely the brain, conceived along irrealist lines, fits into dependence structures. If we accept a 'virtual world' account of perception according to which the entire world we live in is a brain-based simulation, everything that surrounds us will be existentially dependent on the brain. The brain itself still depends on its parts, such as the two hemispheres, which will in turn depend on the synapses they contain, all the way down to whatever the fundamental particles are. These, again, will existentially depend on the brain, since everything in the world around us does. We therefore have a circle of dependence relations in which no entity is ontologically fundamental. Brain-based constructivism swallows up its own foundations. If the brain, too, is a brain-based construction then nothing is ultimately real (in the sense of being unconstructed). Some have argued that that brain-based constructivism is self-refuting, and indeed it is if taken to entail an ontologically fundamental level. If it is not, however, we simply end up in a position structurally identical to Escher's *Print Gallery* where we move from inside the picture to the outside world of the beholder and, continuing our trajectory, find ourselves inside the picture again.

3. Appearance-reality distinction and error theory

One might also be concerned that irrealism undermines both the appearance-reality distinction, and the error theory about ordinary cognitions. The first raises a metaphysical question: If there is no reality behind the appearance, how can we still make sense of the distinction between the two? The second raises an epistemological question: If there is no ultimately true theory, how can we argue that people are massively deluded about the way things are?

In response to the first, note that the appearance-reality distinction only obtains in the interface. There is nothing beyond the interface that the concept 'reality' applies to. Can 'reality' still be a meaningful term despite the fact that "it is never

instantiated, and we are never epistemically connected with it"? I believe it can. Assuming there are no yetis, the term 'yeti' is never instantiated, and we are never epistemically connected with one. Still, sentences like "Reinhold Messner believes yetis are really bears" is still perfectly meaningful.
The irrealist is hence not rejecting his cake and having it too, since there are no cakes. When rejecting foundationalism in favour of non-foundationalism he is not saying that there are no ultimate truths, and that foundationalism fails because it misrepresents the ultimate nature of reality. Doing so would indeed by contradictory. Rather, his point is that when conducting an ontological analysis we might first end up with a foundationalist position. Analyzing this further, however, shows that foundationalism faces considerable problems, so that we will ultimately reject it. This is not because if 'fails to get ultimate reality right' but because by the standards for evaluating ontological theories rejecting foundationalism fares better than embracing foundationalism. And while we might refer to the rejection of foundationalism as anti-foundationalism, this should not be understood as embracing an ontological theory that says the opposite of foundationalism. As the Prāsaṅgika interpretation of Madhyamaka is keen to point out, rejecting some position does not entail implicit affirmation of the opposite.

In this context we should also consider Aitken's point that phenomenal transparency presupposes something behind the veil of appearances (namely the unconscious processes that bring about mental appearances). At the same time I seem to argue that no good sense is to be made of a reality behind the veil of appearances. So how can I still help myself to the idea of phenomenological transparency, of there being cognitive processes we 'see straight through', since they are introspectively inaccessible?
I do not believe that the parts of our cognitive processing that are introspectively inaccessible to us form a special problem for the claim that all is "mere appearance". The same problem is raised, for example, by things that are too small to observe (subatomic particles) or too far away (outside of observable universe, as determined by the age of the universe and the speed of light). We can meaningfully talk about such things, and they can play important roles in our theories, even though they never 'appear' to us in the way tables and chairs do. We postulate them in order to explain regularities amongst things that do appear, and we postulate introspectively inaccessible mental processes in order to construct coherent theories of the mental processes we can access introspectively.

As for the second question, it is clear that finding out that we are 'massively deluded about the way things are' cannot be the result of stepping outside of the realm of appearances in order to have an objective look at the world. Rather, any error must be detected from within the interface. By applying Madhyamaka analysis, we realize that certain assumptions (such as a world behind appearances) do not even make sense conventionally. Once we have realized this, we have not achieved an objectively true view of the world, having previously held an objectively false one. Rather, our present view is better than the previous one, better, not in terms of its

approach to ultimate truth, but in its ability to facilitate the eventual elimination of existential suffering.

4. A lightweight ultimately true theory?

The discussion in the book proceeded by providing negative answers to the four questions raised in the four chapters. Aitken raises the question whether it still incorporates a positive theory making any ultimately true assertions. This is an interesting point, which is in itself an *avatar* of the Indo-Tibetan debate between two different interpretations of the Madhyamaka school, Prāsaṅgika and Svātantrika. The exact difference between these two sub-schools is a matter of considerable complexity, but one relatively uncontroversial point is that Svātantrika-Madhyamaka is often characterized as accepting some philosophical position, while Prāsaṅgika-Madhyamaka is taken to be a wholly negative approach, exclusively engaged with refuting the opponent's position, without advancing any position of its own. Aitken also aligns the Prāsaṅgika/Svātantrika distinction with a positive/negative answer to the question whether ultimate truth is reducible to conventional truth.

She suggests that the claims an irrealist could be regarded as making might be part of a 'lightweight' ultimately true theory, a theory that does *not* refer to ontologically independent or fundamental entities, but which holds of whatever there is, and is the termination point of our ontological analysis. This point is simply "the fact that our ontological analysis never meets with an end". For Aitken, this is not a fact about the world, but a fact about our analytical endeavours. If we carry them out properly (i.e. according to the logical rules that form part of the conventional reality in which we live) we find that we never reach the end. Each entity that might have looked like ontological rock-bottom has another layer beneath it.

I don't think there is anything here a Prāsaṅgika would disagree with, but I also find it hard to see that this is an approach according to which "ultimate truth is non-reducible to conventional truth". Our analytical endeavours are part of conventional reality, and statements about their limits are obtained by reference to inferential procedures which are also part of conventional reality. Why, then, should this statement about conventional truth, obtained by conventionally true means be taken to be an ultimate truth? The reason, I believe, is that for Aitken there are two different sets of epistemic procedures. One, the result of a reflective equilibrium of coherence, intersubjectivity, and efficacy, is used to determine ordinary truths about the world (London is the capital of Britain, water contains oxygen, and so forth), the other, "ontological analysis" is used to ascertain such supposedly lightweight metaphysical truths as "appearances have no ontological foundations" etc. Why, however, should we accept this bifurcation of our epistemic practices? Is it not the same kind of epistemic instrument that we employ to investigate all kinds of truth (even though we might use them in somewhat different ways)? I would prefer to believe that when we establish whether or not the perceived puddle on the road is real water or a mere mirage we use the same conceptual toolkit we employ when trying to settle whether all, or merely some things are appearances. We are equipped with the epistemic instruments we have, and have to tackle all questions requiring settlement with some of these instruments. There is no privileged set that could

allow us to differentiate the matters decided by it from those decided by 'ordinary' epistemic instruments outside of this set.

What, then, about the question whether the irrealist holds any *positive* ontological theses? It seems as if the irrealist's "*prima facie* unlocalized ontological claims" Aitken mentions can all be re-interpreted as claims about the dialectical exchange of the Mādhyamika with his opponents along the following lines:

1a. There are appearances only.
1.b Every attempt to discover a non-appearance leads to a contradiction.

2a. Appearances lay no claim to mind-, language-, or theory-independence.
2b. Every attempt to present such a claim can be shown to be inconsistent.

3a. Appearances have no ontological foundations.
3b. Every possible candidate for an ontological foundation can be shown to be problematic.

4a. Things in the world depend all the way down or round and round.
4b. The only other alternative, foundationalism, is inconsistent.

5a. We are inhabitants of a virtual world.
5b. If we assume anything is non-virtual we end up with a problem.

Moreover, the two criteria of a lightweight's theory holding of whatever there is, and of constituting a termination point of our ontological analysis could also be phrased in terms of exchanges between the Mādhyamika and his opponent. Suppose T is the conjunction of all the claims made by the lightweight ultimately true theory. To say that T holds of whatever there is simply means that for every counterexample *c* the opponent comes up with the Mādhyamika can show that accepting the existence of *c* leads to a contradiction. Similarly, any entities referred to in an ontological analysis inconsistent with T can be further analysed to produce an analysis consistent with T.

The key difference between this and the original phrasing is that the latter is about the world, and despite avoiding reference to any fundamental entities still incorporates fundamental truths about the world, truths that are supposed to hold no matter what, or of "whatever there is"[9]. The former refers to what is and is not possible in argumentative exchanges. It is based on fundamental principles accepted by both parties (such as the law of non-contradiction) and argues that certain entities cannot be consistently incorporated in a theory of the world. It does not, however, suggest a description of the world *sub specie aeternitatis* that might explain why this is the case.
If these reformulations are successful then the fundamental question is no longer whether or not the Mādhyamika has a thesis that holds of whatever there is and

[9] This is, if I understand it correctly, what is behind Aitken's suggestion that emptiness is not a "first-order foundation" (an ultimately real entity) but a "second-order fact" (a truth about the first-order entities).

constitutes the termination point of our analysis. I believe that Aitken makes a convincing case in favour of an affirmative answer. But this in itself does not settle how "having a thesis that holds of whatever there is …" is to be interpreted. Does this mean having a thesis about the world, or is it having a thesis about the procedures we employ for analysing the world? This question, it seems to me, remains open.

5. The question of solipsism

Aitken raises the interesting question of the relation between irrealism and solipsism. If solipsism is understood as the position that I alone am real, while all other beings and things are a product of my mental creation, it is clear that the positions are very different: irrealism is incompatible with the existence of any fundamentally real person. However, we might also understand solipsism as the position that other people are not (fundamentally) real, independent of any pronouncements about the reality of ourselves. It is this form of solipsism that is relevant in the Yogācāra discussion about the reality of other mental continua we find, for example, in the works of the two important Yogācāra thinkers Dharmakīrti[10] and Ratnakīrti,[11] it is also the one that we should have in mind in the present discussion. When we ask questions like "are we a solitary representational interface, or are different interface linked up as in a multi-player video game?" or "am I the only one experiencing phenomenal subjectivity, or are the other minds that do so as well?" we ask about what is happening beyond our representational interface. Answering such questions either positively or negatively is problematic since it presupposes something the interface theory says we cannot have: a view from outside of our interface into reality as it is in itself, in order to determine whether or not there are other interfaces we could at least see from the outside. The position that there are no other interfaces beyond our interface is as problematic as the nihilistic position that there is nothing beyond the interface, since it presupposes the kind of ultimately true theory the arguments in chapter 4 set out to show to be impossible. The only response to the question of solipsism the irrealist can give is to say that from the perspective within the interface (aka relative truth), which is the only perspective we can occupy, there are other persons, even though this perspective does not provide us with any information about the furniture of reality at the ultimate level. Can irrealism accommodate "multiple first personal perspectives and phenomenally conscious experiences"? In one sense the answer is obviously 'yes': our view of the world incorporates the existence of other people besides ourselves. But, one may then ask, are they also real from the inside? As long as an answer to this question presupposes anything more than information accessible within the interface (how people behave etc) it cannot be a meaningful question. If 'reality from the inside' can only be determined by us through leaving our interface, and if leaving the interface undermines our very existence, as leaving a liquid undermines the existence of a bubble in the liquid, the question cannot be settled at all.

[10] 6th ct CE. See Stcherbatsky 1969.

[11] 11th ct CE. See Ganeri 2015: 202-212, see also Perrett 2017.

Christian Coseru

1. The 'false dilemma'

Coseru argues that my argument for irrealism relies on the false dilemma of either reducing the world to the representational interface, or as regarding it as a noumenal something behind the interface, something, to use Putnam's inimitable phrase 'one-knows-not-what which solves our problems one-knows-not-how'. This dilemma, he claims, ignores alternatives such as Husserlian phenomenology or enactivist accounts of cognition.

Coseru questions the representationalist view of perception that forms the background to much of the discussion in chapter 1, arguing that we should instead adopt an enactivist account of perception. According to such an enactivist position, "perceptions are not representational contents *about* the world; rather they are engagements and involvements *with* the world." What this means depends crucially on what we mean by 'the world'. If the world is an assortment of physical and non-physical things existing in a mind-independent manner in a mind-independent space-time continuum a realist understanding is baked into this statement of the enactivist view of perception. This makes it of limited value in a discussion where the very viability of such a view of perception is at being considered.
On the other hand, the world could also be equated with the representational interface. As for the irrealist our perceptions clearly are "engagements and involvements" with the representational interface: the statement expressing the enactivist position thus understood is clearly true from an irrealist perspective. Accepting enactivism 'in the image' does not reduce to enactivism, no more than accepting realism in the image reduces to accepting realism, since the enactivist and the realist presumably accept the existence of an ultimately true theory of the fundamental structure of the world, while the irrealist does not. Depending on how the expression of enactivism is understood, then, the irrealist would either want to see an argument for the external world realism presupposed by it, or he could unproblematically endorse the enactivist sentiment as speaking from *within* the representational interface.[12]

Similarly, to argue that a phenomenological account of perception "confronts us […] with the objects themselves" may or may not be construed as tension with the irrealist's position. It is evident that the 'objects themselves' here cannot be mind-independently real physical entities laid out in a spatio-temporal arena, otherwise we would not be dealing with a phenomenological but with a naïve realist view of perception. But if the ontological status of these objects is either bracketed (so that they could turn out to be anything, including parts of the representational interface), or if they are simply identified with whatever our perceptual relation is taken to relate us to (in the irrealist's case: parts of the representational interface) there is no need for the irrealist to object to this account of perception.

[12] See Button 65–7: 2013.

While Coseru sees me in danger of ignoring a possible 'third way' between an interface-only position and positing an ineffable noumenon beyond our epistemic reach I am worried about the shape-shifting nature the route between the two horns of the dilemma might be taken to possess. In some contexts the proposed alternative can be understood as a hard-nosed realism robustly committed to a world 'out there', in others it morphs into a kind of Husserlian transcendental idealism hard to distinguish from the irrealist position proposed in all but terminological framing. But of course we cannot have it both ways, and care must be taken that our response does not slide from one to the other in order to avoid the impact of any given argumentative pressure.

2. Differentiating experiences

Another worry Coseru raises is that the irrealist model may be unsatisfactory as it fails to account for the ways our experiences appear to us. Nevertheless, it seems to me that irrealism has sufficient resources to account for such features. There is, for example, no reason why Coseru's "richly detailed visual experience" of the Empire State Building should not be something that can take place within the representational interface, nor would the claim that all perception takes place within the confines of the interface entail that we cannot distinguish between experiencing, dreaming, imagining, or remembering some object. For sure, according to the irrealist account seeing the Empire State Building and remembering it does not link us up with entities of fundamentally different kinds (a building in one case, a memory trace in the other) — in both cases we are related to bits of our representational interface. But these bits are related to other bits in different ways, and it is *this* difference in relational structure that allows us to tell one from another. Seeing the Empire State Building in front of us is likely to trigger the expectation of being be able to walk to Grand Central Station in under twenty minutes, remembering seeing it will generally not trigger this expectation. Our perceptions relate to the rest of our cognitive lives in a way that is quite different from the way our imaginings do, and this is what makes them qualify as different kinds of experiences. For this reason Coseru is right to point out that hallucinated water cannot quench our thirst, but it does not follow from this that perceptions of non-hallucinatory water are "engagements and involvements *with* the world", while perceptions of hallucinatory water are not. Rather, the former are related to cognitive states of one kind (quenching waking-state thirst) while the latter are related to cognitive states of another kind (quenching thirst-in-the-hallucination). The existence of these different relations allows us to differentiate veridical and hallucinatory experiences, but it does not on its own establish that only the latter are "involvements with the world".

Coseru also notes that we are "aware of at least some features of the perceptual field that are unattended to", suggesting that the irrealist could account for this by including these features within the representational interface, but that this would face "significant explanatory costs associated with modelling unexplored and unattended features". Yet this seems to overlook that there is no more necessity for

explicitly "modelling" such features than there is — to use Dan Dennett's example — to fill in for the blind spot when observing a wallpaper with hundreds of identical portraits of Marilyn Monroe. In the same way in which the brain does not have to reproduce the Marilyns "as if by photocopying across an external mapping of the expanse of the wall",[13] unattended features need not be explicitly modelled within the perceptual interface and thereby increase its computational complexity beyond plausible limits. It is only when the unattended features become attended to that the perceptual interface has to move from something like the mere label 'unattended features here' to something more specific.

[13] Dennett 1992: 354.

Bibliography

Tim Button: The Limits of Realism, Oxford University Press, Oxford, 2013.

Daniel Dennett: Consciousness Explained, Penguin, London, 1992, 354.

Jacob Feldman: "Turning your priors to the world", Topics in Cognitive Science 5:1, 2013, 13-34.

Paweł Gładziejewski: "Predictive coding and representationalism", Synthese 193, 2016, 559–582

Jonardon Ganeri: The Self. Naturalism, Consciousness, & the First Person Stance, Oxford University Press, Oxford, 2015.

Roy Perrett: "Buddhist idealism and the problem of other minds", Asian Philosophy 27:1, 2017, 59–68.

Chetan Prakash, Chris Fields, Donald D. Hoffman, Robert Prentner, Manish Singh: "Fact, Fiction, and Fitness", Entropy 22, 514, 2020, doi:10.3390/e22050514

Theodore Stcherbatsky: "Establishment of the Existence of Other Minds: Annotated Translation of Saṃtānāntarasiddhi. A Free Translation of Dharmakīrti's Santānāntarasiddhi with Vinītadeva's Santānāntarasiddhiṭīkā" in Papers of Th. Stcherbatsky: Translated for the First Time into English by Harish. C. Gupta, edited with an Introduction by Debiprasad Chattopadhyaya, Indian Studies Past & Present, Calcutta, 1969, 81–121.