

Manchester Lit & Phil

Philosophy Forum



Reality – veridical or virtual?

Seeing may be believing ... but believing what?

Focus paper for the Manchester Lit & Phil Philosophy Forum

Christopher Burke

August 2023

Are we each living in a virtual reality or do we perceive the world as it really is? Well it's obviously the latter: how else could we get on with our lives? So says our common sense. Unfortunately this strong intuition doesn't hold water. Our relationship with reality is a lot more interesting than that. Welcome to the world of fitness pay-offs and interfaces ... **your 'real' reality?**



This paper deals with a contemporary philosophy about what is going on when we perceive our world and try to understand it: an ancient philosophical topic.

The particular metaphysical philosophy explored here is proposed by **Donald Hoffman**, a world renowned Professor of Cognitive Studies at the University of California.



His philosophy consists of **three interconnected theories**:

Fitness Beats Truth (FBT) Theory
Interface Theory of Perception (ITP)
Conscious Realism



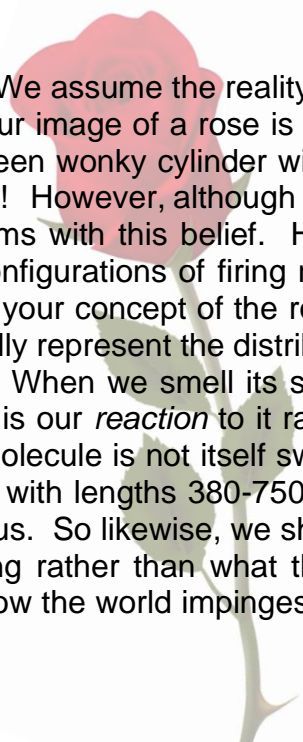
Below we take a critical look at Hoffman's **first two theories**; a companion paper for a following Philosophy Forum is devoted to the third. Each paper can be treated independently, but they do enhance one another.

The concepts are challenging, but technical terms are explained in side panels close to their first usage, where they are highlighted in a **bold orange** font. We will also clarify the ideas at the Philosophy Forum. They are **summarised with discussion questions in section 5**, and some people might find it useful to look at this section first for a quick overview of the topic.

1 A rosy view of the world?

"A rose is a rose is a rose is a rose"

adapted from a line by Gertrude Stein, 1913, in her poem *Sacred Emily*



A rose is a rose ... or is it? We assume the reality of the world around us is just as we perceive it. Of course our image of a rose is like the rose really is: red sheets curled round the top of a green wonky cylinder with little green sheets sticking out of it. That's the way a rose is! However, although this might seem obvious common sense, there are big problems with this belief. How can an image in your head, presumably consisting of configurations of firing neurones, possibly look like your observed rose? How could your concept of the rose not be overwhelmed if it was like *all* the data needed to fully represent the distribution of the billions of molecules constituting the 'real' rose? When we smell its scent or eat its hips, we normally think the smell or taste of it is our *reaction* to it rather than an intrinsic property of the rose itself. A sucrose molecule is not itself sweet; it just induces sweetness in us. Electromagnetic waves with lengths 380-750 nanometres are not themselves red; they induce redness in us. So likewise, we should surely think of *any* sense at all as *our reaction* to a thing rather than what the thing is like objectively. Our subjective reality is surely how the world impinges on us rather than how the world actually is.

2 The battle of the senses

2.1 The above argument is a very old one and some people reject it: they think that in many ways our senses *do* reveal the way the world really is. The shapes we perceive are the shapes of the object we are looking at. We feel the sharpness of an object is because it has a sharp point. Throughout the history of philosophy, a battle has raged about the nature of **perception** and thereby our knowledge of reality. Note that perception here is not just visual. On one side of the battle is **Empiricism**, a cluster of philosophies sharing the belief, put simply, that what we perceive is **veridical** – the way external reality is 'in itself'. Thus Empiricists consider themselves entitled to make **ontological** claims about what reality consists of. Its most noted ancient advocate would be Aristotle (384-322 BCE), but Scientism (the belief that reality is as science describes it) and more widely Physicalism (reality is as our physical representations describe it) could be considered as examples of its current forms.

Veridical:

Truthful. A veridical perception is one that represents things as they are, contrasted with an illusory or even delusory one that does not.

<https://philosophy.en-academic.com/2467/veridical/>

Perception:

The meaningful interpretation of data from all the senses, not just vision.

Ontological:


Referring to the nature of reality itself.

2.2 On the other side is **Idealism**, a cluster of philosophies sharing the belief, put simply, that what we perceive is **virtual** – ie cognitive *constructions* in response to our sensorial input, which cannot be assumed to be anything like external reality 'in itself'. Thus Idealists can't know what reality consists of, but legitimately can only make **epistemological** claims about our *understanding* of reality, *not* reality itself. Plato (~425-348 BCE) is considered the major figure here, with Immanuel Kant (1724-1804) as the preeminent post-Enlightenment advocate. The notion – widely believed by current philosophers, neuroscientists and psychologists – that we create an internal world model, is essentially idealistic. Conscious Realism is definitely a current example of Idealism.

2.3 The two stances, Empiricism and Idealism, have very a different understanding of truth. Empiricists, perceiving reality veridically – as it is 'in itself' – think that truth can be simply checked against observations because experience *corresponds* directly with reality. So this theory which empiricists hold is called a **correspondence theory of truth**.

2.4 This avenue is of course not available to idealists since they cannot access reality 'in itself' to check their percepts and concepts directly. For them, the only means of establishing truth is that their cognitive constructions are consistent with incoming sense data, memories and logical norms. So Idealists hold a **coherence theory of truth**.

2.5



2.6

Edward Steed, *Philosophy Illustrated*

Virtual:

“That which is ‘not real’, but displays the full qualities of the real — in a plainly actual (i.e. not potential) — way. The prototypical case is a reflection in a mirror: it is already there, whether or not one can see it; it is not waiting for any kind of actualization. This definition allows one to understand that real effects may be issued from a virtual object, so that our perception of it and our whole relation to it, are fully real, even if it is not.”

<https://en-academic.com/dic.nsf/enwiki/365152>

Epistemological:

Referring to the nature of knowledge about reality.

It is not logically possible *ever* to adjudicate on which approach, Empiricism and Idealism, is right because it is not possible to step outside of our perceptions and check them against reality. And neither could someone else check them because they can't step out of their reality to do the checking either. Personally I adopt a representationalist stance: our percepts and concepts are only 'like' what they represent in that they contain essential information *about* the reality they represent.

But that information is stored in a completely different way to information stored in whatever thing is represented. Here a computer can provide a useful metaphor: a complex program accurately representing the economy of a country doesn't resemble it – it just has informational elements and operations

attempting to correspond to those in the real economy. Similarly we understand our world through a mental *model* of it that we hope corresponds in propitious ways with the world around us. The most plausible way of thinking about our experiences is that we essentially build them from *virtual* representations. So how do we do that?

- 2.7 This very simple sketch of millennia of metaphysical theories polarises the Empiricism-Idealism debate: greater space would allow the inclusion of many more subtle and plausible intervening positions. But it will suffice as a background for us to take a critical look at a relatively new theory – Fitness Before Truth (FBT), which explains *why* our representations cannot be veridical. A leading proponent is Donald Hoffman, who marshals considerable evidence to back that claim.

3 Fitness Beats Truth (FBT)

- 3.1 What FBT maintains is that we perceive the world in terms of its benefit or threat both to our individual survival and to the continuance of our genes through reproduction. We have evolved to perceive and understand the world not in terms of veridical truth, but as fitness-payoffs.

- 3.2 An obvious counter-argument to this would be that we need to know reality veridically to know what is best for our survival. But this objection is open to a number of challenges. Hoffman uses mathematical game theory to ‘prove’ that fitness strategies always beat truth strategies. However the quantitative axioms from which this conclusion is derived have been questioned and the modelling of the two strategies criticised as simplistic, eg:

“Hoffman’s sceptical argument to the effect that usefulness trumps truth only works if the signals produced by the sender are the sole available cue in the maximization of usefulness. If, on the other hand, the sender-receiver system needs to build its sense of what is useful by combining information coming from different sources, things are not so clear.”

Manolo Martínez, 2019, Usefulness Drives Representations to Truth: [A Family of Counterexamples to Hoffman’s Interface Theory of Perception](#), *Grazer Philosophische Studien* 96: 319-41.

- 3.3 Hoffman uses many other more plausible arguments for FBT, including the non-veridical nature of vision revealed through illusions and the bodily need to efficiently minimise energy usage by just focusing on what is environmentally *relevant* to the agent. Veridical perception would require the processing of so much **extramental** information that it would be overwhelming. If a bio-agent had to consider *all* the possible information about a potential mate and not have the heuristic informational shortcut of beauty, their beloved would have lost patience long before they finally

There is an important philosophical duality in how we understand reality – between the phenomenal and the physical. This is very relevant to any exploration of selfhood. The philosophical definitions of these words are different from their everyday usage.

Phenomenal:

Relating to *subjective* reality, ie *first person* lived experience – *what it’s like to be* thinking and feeling.

Physical:

relating to a putative *objective* reality, ie what is revealed by careful *third person* observation and measurement.

got round to some action. And lost love would be the least of their worries if, while ponderously processing every bit of information about a looming predator, they had been eaten!

3.4 Hoffman doesn't refer to it, but FBT is strongly consistent with Karl Friston's compelling Free Energy Principle (cf panel). But he does quote some alternative heavy backing, for instance from the renowned cognitive psychologist Steven Pinker:

"We are organisms, not angels, and our minds are organs, not pipelines to the truth. Our minds evolved by natural selection to solve problems that were life-and-death matters to our ancestors, not to commune with correctness."
Steven Pinker, 1997, *How the Mind Works*

3.5 The non-veridical nature of our perception and thought is now widely accepted in most relevant academic disciplines. Our **intramental** model is constructed as a propitious tool for agency, not to be 'like' the extramental world it represents. The mind/brain doesn't take snaps.

3.6 Personally, I think that Empiricism is highly implausible for the reasons Hoffman and many others give, which boil down to it being far too cumbersome, unfocused and energy inefficient to have been selected by evolution. There is also a more profound philosophical objection to Empiricism which would favour FBT's renouncement of it: how could any intramental symbolic representation of an extramental entity be 'like' that entity? This certainly is not the case in any possible natural or artificial sensual or cognitive system. In **physical** terms, the extramental **representatum** consists of whatever 'stuff' constitutes it; but that is very different from the neurological or electronic structures which constitute its representation. Neither your phenomenal image of a rose nor the neuronal configuration which accompanies it are similar in form or content to a physical rose. Representation and representatum are different bits of reality, *unlike* one another.

3.7 But surely our visual system must preserve the basics of what it is looking at: edges, shapes, colours, textures, etc? These are certainly the ways we visualise the world, but that does not mean that the world is as we visualise it. And even if we accept a simple model of the retina preserving the objective distribution of light impinging upon it like a camera, visual processing from there to the visual cortex is

The Free-energy Principle (FEP)

FEP is a complex theory, developed mainly by Karl Friston, featuring much mathematics, about how animals survive through efficiently managing their available energy by processing intero-sensory and extero-sensory information to maintain the critical homeostasis (dynamic equilibrium) necessary for life. This involves the brain continually *predicting* events so as to minimise surprisal, which necessitates a greater expenditure of energy and potentially poses a greater risk to survival. Thus a bio-agent protects its highly ordered intrinsic systems against the prevailing and threatening entropy (tendency to decay and disorder) of its environment. It is posited that consciousness is a necessary emergent of this process, where cognition and affect combine to optimise the self's decision-making through a mix of adjusting expectations and adjusting the environment via motor control.

(cf Karl Friston, 2010, *The free-energy principle: a unified brain theory?* <https://pubmed.ncbi.nlm.nih.gov/20068583/>)

Extramental:

Existing independently of the mind/brain.

Intramental:

Referring to thought and feelings 'within' the mind/brain.

Representatum:

The thing being represented, ie its referent.

extremely complex. For instance, what we assume is an extramental edge is represented by a pattern of neurones (or logic gates) ... hardly edge-like. If Empiricism runs aground in physical terms, how does it fare in psychical terms? Surely our **phenomenal** experience when looking at a putative edge is 'like' the *physical* edge? To which I would answer: in what way could a thought possibly be 'like' an edge? The commonplace notion of our percepts and concepts about our world being veridical is rationally unsustainable.

3.8 One could possibly rescue a weaker form of Empiricism by recasting its claims in informational terms: ie the *essential information* of the representatum is retained by the representation. This seems obviously plausible: how could a bio-agent survive if their senses didn't capture essential information about the world around them? But then what factors determine what is essential information and what is not? The most plausible answer would surely be: the information which matters to the bio-agent doing the observing ... that which helps it stay alive. So we are back to Fitness Beats Truth.

3.9 However, the Fitness Beats Truth theory begs an important question: how are the fitness payoffs represented? Here Hoffman introduces his next theory: the Interface Theory of Perception. We need to look even further beyond our commonsense notions of reality ...

4 Interface Theory of Perception (ITP)

4.1 FBT raises an important question: if we don't veridically perceive extramental 'reality', what do we perceive? Hoffman's answer is to rely heavily on the metaphor of a computer desktop – the interface between user and hardware. So he maintains that we experience 'icons' (percepts) that allow us to manipulate the 'hardware' (reality). His use of this metaphor is obviously reinforcing the idea of experience as *virtual* not *veridical*, notwithstanding the vividness of our perceptions:

“A red Maserati looks so shiny, artistic, aerodynamic, so real. But the FBT Theorem tells us that it's just a sensory experience—an icon—that is not objective and depicts nothing objective.”
Donald Hoffman, 2019, *The Case Against Reality*

4.2 He attempts to dispel our initial incredulity about such a statement by appealing to our sense of taste rather than vision, as I did above with the scent of a rose. He challenges us by pointing out that we don't normally think that the taste of a strawberry is like the real entity we call a 'strawberry'. Most people would concede that its taste is firmly part of our conscious sensory experience with the strawberry itself merely triggering it. But if we don't think of taste as veridical, then why should vision or any other sense be veridical? This argument is compelling. Whatever the qualities of a mind-independent object are, we can only know how they *affect* us ... not how they *are in themselves*. Hence our perception delivers, in Hoffman's terms, 'icons', not real entities. Brian Martin summarises it baldly:

“[Hoffman] says there is a reality, but that we don't know what it is. Our senses are designed to interpret reality using a code that is advantageous for fitness; our senses, without extra assistance, simply don't have the capacity


to get beyond this code in any systematic way, any more than we can grasp the reality in computer electronics by analysing what we see on the desktop. To mistake the interface for reality is like thinking that the letters CAT are the same as the animal.”

Brian Martin, 2019, [Do We See Icons or Reality?](#) A Review of Donald Hoffman's *The Case Against Reality*, *SERRC (Social Epistemology Review and Reply Collective)*

4.3 An accusation against ITP is that it is parasitic on ‘metaphysical reality’ (ie the notion that reality consists of mind-independent discrete objects in spacetime) which is the very philosophical position ITP attacks:

“Hoffman’s argument, in fact, presupposes the existence of such mind-independent objects for it to work; that is, just those objects that provide the selective pressures for fitness to evolve and to evolve in.”

Leslie Allan, 2022, *Hoffman's Conscious Realism: A Critical Review*



Laleh Golafshani, [The Lighter Side of Academia](#)

Allan is not the only one who considers ITP (and by inference FTB) to be self-refuting:

“Hoffman’s two theorems privileging ‘evolutionary fit’ perception over veridical perception are incompatible with Darwinism itself. As such, his ‘case against reality’ is a particular egregious example of what philosophers call ‘pragmatic self-refutation’, in which an argument put forward to support a position actually undermines its own premises.

Raymond Tallis, 2023, [An Encounter with Radical Darwinitis](#), *Philosophy Now*, Issue 154

4.4 I don’t think this charge that ITP depends on the very theory it refutes is valid. Hoffman could easily respond that his use of the computer desktop metaphor and evolution doesn’t imply ‘metaphysical realism’: both could themselves be regarded as convenient ‘icons’ with which we model the world. Hence ITP *could* remain consistent.

4.5 Unfortunately Hoffman is not content with this riposte. He and his fellow researchers don’t seem to regard ITP itself as an ‘icon’, but instead assume that ITP indicates the way that the reality *actually is*, an *ontological* claim which would be barred by ITP itself. Look back at his claim about the Maserati: “...it’s just a sensory experience—an icon—that is not objective and depicts nothing objective.” By stating that our icons depict nothing objective, he is straying into making an ontological claim: ie claiming to know what objective reality ‘really’ is. How can he know by just experiencing his icons that they depict *nothing* objective? This borders

on Solipsism; the idea that there is no independent *objective* reality, merely an individual's own *subjective* reality.

- 4.6 In an earlier paper, Hoffman et alia double down on this by reasoning from the interface metaphor. By implication, they posit this complete lack of correspondence between perceptions and reality by pointing out that ...

“No features of the icon are identifiable with any features of the file in the computer.”

Donald Hoffman, Manish Singh and Chetan Prakash, 2015, *The Interface Theory of Perception*, *Psychonomic Bulletin and Review* 22/6

But how can ITP possibly grant this magical ability to know the inaccessible features of the referent thing in itself (represented by the underlying file) to compare it with the percept (represented by the icon on the interface)?

- 4.7 I think Hoffman et alia have extended the interface metaphor beyond its usefulness here. Icons on a desktop may contain no information about the underlying files they refer to, but percepts surely must contain information about what they refer to ... the bit of observed reality which has induced them. Otherwise how could bio-agents build efficacious mental models that help them navigate and manipulate the world.

- 4.9 Personally I think Hoffman is right for all the reasons presented above to reject the idea that our percepts are veridical. Logically, we can *never* know what unobserved extramental reality actually is. But since some of our percepts and concepts actually work, it is plausible to believe that they must contain some relevant information about mind-independent reality. When defending his Fitness Before Truth theory, he seems to accept this. But when he moves on to his Interface Theory of Perception, he explicitly denies this by stating his icons depict nothing objective. To me, the denial of directly knowing objective reality is cogent, but the denial of objective reality itself is a step too far in the idealistic direction. But Hoffman could reply that as the subjective world is all we know, it is me that is making invalid assumptions by claiming the existence of an objective reality. Indeed he develops a whole theory based on this: Conscious Realism, which proposes that consciousness is the only reality. This is explored in a companion paper to this one.



Whether or not you agree with him, Hoffman's latest book is an exciting read:

Donald Hoffman, 2019, *The Case Against Reality*, Allen lane.

This can be downloaded via [Oceanofpdf](#).

There is also a lengthy 2022 discussion: *Lex Fridman Podcast #293*:

Reality is an Illusion - How Evolution Hid the Truth.

Those without the time to tackle these can look at his 22 minute **TED Talk from 2015: [Do we see reality as it is?](#)**

5 Summary and questions

5.1 **Fitness Before Truth (FBT)** theory proposes that we don't know the world as it *really* is. Our understanding of it – the mental model we construct to get by in life – is based on the benefits or threats to our survival. A convenient illusion. The main justifications for this theory are:

- taking in all the information needed to represent the world as it truly is would not be a viable survival strategy
- evolutionary selection could never favour such a waste of time and resources
- no cognitive system could cope with the vast inflow of information
- no living thing could satisfy the impossibly high demand for precious energy.

Our subjective experience of the world is therefore not veridical.

Do you think that FBT is:

Q1

- **Completely correct?**
 - **Somewhat correct? If so, what else is going on besides FBT?**
 - **Not at all correct? If so, what is a credible alternative to FBT?**
-

5.2 The **Intercept Theory of Perception (ITP)** is a response to the question thrown up by FBT. If our experiences aren't the truth, then what are they? ITP suggests they are similar to the icons on our computer desktop: a convenient interface. And just as computer icons are not like the files they link to, so our percepts and concepts are not like the bits of the 'true' world they refer to. **Our subjective experience of the world is virtual.** Donald Hoffman's version of ITP goes further by maintaining that our icons/percepts, don't informationally correspond with *anything* objective.

Do you think that our percepts are:

Q2

- **Exactly like the real objects they represent? If so, how can you justify that?**
 - **Somewhat like the real objects they represent? If so, in what way?**
 - **Nothing like the real objects they represent? If so, how could you know that?**
-





Your notes

**Manchester
Lit&Phil**