

Reviews

Reality+ : Virtual Worlds and the Problems of Philosophy by David J. Chalmers (W. W. Norton & Company, 2022).

doi:10.1017/S003181912200033X

Some philosophers are purists, thinking that the problems of philosophy float above the world of changing empirical circumstances. In *Reality+*, David Chalmers demonstrates the untenability of this purism by showing that technology raises new philosophical questions and changes old ones. The book is also successful as a relatively accessible, entertaining, and not entirely Eurocentric introduction to the problems of philosophy. It is a sprawling work covering many different topics, and a kind of manifesto which argues for Chalmers's sometimes controversial views, some of which are developed more fully in his earlier work, and which together form a general approach to reality in a technological age. Most strikingly, he proposes a 'structuralist' account of reality that can solve the traditional problem of global skepticism about the external world. This claim is the central, recurring theme of the book that holds the disparate parts together. Unsurprisingly, since it targets one of philosophy's enduring problems, it the most philosophically problematic claim in the book.

According to Chalmers, you and the world you see may well be part of a simulation – he thinks that there is at least a 25% chance of this (p. 101). His reasoning for this surprising estimate resembles Bostrom (2003)¹, but goes beyond it in some details. Throughout the history of the universe, there will probably be many advanced civilizations with the technology to create trillions of detailed simulations containing 'sims,' or simulated beings that resemble you. And some of these civilizations are enough like ours in their needs and interests to want to do so (pp. 90, 138-39). Of course, some may not bother. But if even one out of a million such civilizations does so, that one could well create trillions of sims, which would vastly outnumber the non-sims in the universe. Accordingly, you're probably a sim because most conscious beings in the universe are (Ch. 5).

A crucial step in his argument is that simulated beings can be conscious, just like you. His argument for this, though, could have used more discussion, and objections to it considered more fully. Still, one need not agree with all of Chalmers's arguments (nor his estimate of the chances) to appreciate the main upshot: the simulation scenario is a real possibility. The closer our technology gets to producing a

¹ Nick Bostrom, 'Are We Living in a Computer Simulation?', *Philosophical Quarterly* 53 (2003), 243-255.

Reviews

simulated world as rich as the one we currently experience, the more clearly it is *possible* that we are in a simulation. We cannot know that we are *not* in such a simulation.

Fortunately, Chalmers holds that this is not the skeptical disaster it is often taken to be. If you are in a global simulation – that is, your entire life is within the simulation – the simulated objects around you are perfectly ‘real’. (pp. 170-71). This follows from *structuralism* (Ch. 22), roughly the view that the physical properties described by science are equivalent to structural properties. Structure captures the mathematical and causal relations between phenomena, roughly the ‘role’ that each property plays in relation to other properties and observation (Ch. 22). Since fully simulated worlds can have the same structure as the non-simulated world we usually take ourselves to be in, the physical claims we hold true in the non-simulated world can also be true in the simulated world. These physical claims include ones that make many of your ordinary beliefs, such as that there is a table here, true.

Crucially, different substrates can manifest the same structure. In a global simulation, all the objects we see around us, including the particles they’re made of, are ultimately made of bits of information, or digital signals. The simulation hypothesis (i.e. that we are in a global simulation) is a metaphysical hypothesis about the fundamental nature of objects, rather than a hypothesis in which there are no such objects. Just as we would consider a table to be a table regardless of what it is ultimately made up of (quarks? quantum fields? ideas in God’s mind?), we should consider it to be table even if it turns out to be made of bits of digital information. Thus, a simulated table is just a table, with a digital ultimate nature (pp. 118-19). This is the doctrine of *simulation realism*: ‘If we’re in a simulation, the objects around us are real and not an illusion’ (p. 106).

Though structuralism receives an extended discussion in the book, some of that discussion may raise questions don’t get directly addressed (though perhaps they could). Reality seems to be reduced to physics, and physics seems to be reduced to mathematical, structural properties of a sort that current science deals with (pp. 177-79). But, why should we think that the ultimately *correct* physics is not substrate-specific? And why should we think that all there is to reality is what is scrutable to science? Couldn’t some aspect of what it is to be a real quark, for example, be hidden to human intelligence? The book seems to assume a thoroughgoing scientism, perhaps even a kind of verificationism. Couldn’t there be more to reality than mathematical structure, even if ultimately correct physics does not describe it?

At any rate, the same idea applies to global scenarios that don't involve a Matrix-style simulation. Any global skeptical scenario will do if it matches the causal structure of the ordinary world. This includes the famous brain-in-a-vat scenario, the Cartesian evil demon, and the life-long dream (pp. 181, 454). In Descartes' demon scenario, there is a table here, and it is ultimately an idea in the demon's mind. It's harder than philosophers have thought to come up with a (global) scenario in which this is *not* reality.

If all of this is right, then the now standard skeptical argument fails:

1. You can't know you're not in a simulation
2. If you can't know you're not in a simulation, you can't know anything about the external world.
3. So: You can't know anything about the external world.' (p. 56)

Many responses to skepticism target 1, but Chalmers accepts 1 and rejects 2. According to simulation realism, you can know that there is a real chair in front of you even if you do not know whether you are in a grand simulation.

There is much more in *Reality+* than I could summarize here, and there are many controversies worth exploring along the way. But the biggest worry one might have concerns the main claim, simulation realism: it seems to address but one formulation of the skeptical problem, rather than the core problem itself.

If one rejects 2 because, according to structuralism, simulated objects count as real external objects, then what, exactly, have we avoided in avoiding 3? We have avoided the conclusion that we are ignorant about whether *something that causally behaves like external object* exists. But we never worried about that; we were trying to avoid a more specific conclusion. A skeptic who accepts 3 can attain the same knowledge about the external world as the simulation realist claims to have, simply by inventing a new term, 'tabby', for anything playing the causal role of a table. Any scenario in which there is a structuralist table is one in which there is a tabby, and vice versa, so the skeptic and the simulation realist posit the same knowledge about the world. Given this, it's hard to see how any problem posed by skepticism has been solved here: the skeptic who accepted 3 can come to know as much as the simulation realist, merely by inventing a new term.

We can see the problem by observing that what we can call *standard skeptical conclusion*:

Reviews

I don't know whether there is a table because it might be (part of) a simulation.

seems to indicate the same basic ignorance about the world as the *simulation realist conclusion*:

I don't know whether the table is (part of) a simulation.

These make all the difference for Chalmers, but they sound like reformulations of the same ignorance. In both cases, the same, *basic skeptical conclusion* holds:

I don't know whether the thing causing my table experience is a simulator, a demon's idea, part of a dream, etc.

This basic skeptical conclusion is not avoided by simulation realism, and yet it is why standard skepticism seems so problematic in the first place. If you don't know whether you're in a simulation, you don't know much at all about what the thing causing your table experience is, about your environment beyond the experiences. (This, incidentally, is exactly what Hume regarded as the skeptical problem in the *Enquiry* 12.1.)

Here, then, is the heart of the issue. Simulation realism and standard skepticism leave you ignorant about the same set of possibilities. Both agree that, as far as you know, there is either a simulation-caused table appearance, a dream-caused table appearance, an ordinary-process-caused table appearance, and so on. The only difference is that the simulation realist counts knowledge of this disjunction as knowledge that a massively disjunctive *thing*, 'table,' exists. The standard skeptic is stingier about how to apply 'table'. But what matters, in terms of your information about what is going on around you, is which possibilities you can rule out, not how you describe those possibilities.

Chalmers even makes the claim, in the concluding chapter, that 'conscious or not, other people exist.... Perhaps other people [who are not conscious] were only recently implanted in my memories. If so, you'll have told me only that other people are digital beings who are part of a simulation that produced memories.' (p. 461). This variation of Russell's example from *The Problems of Philosophy* is to illustrate a 'broad claim about the shape of the world and my life' that could not, on Chalmers's view, be undermined even by non-global skeptical scenarios. But, again, it is hard to find solace here. The use of 'people' in Chalmers's statement is so easy to satisfy as to be almost vacuous. It seems beside the point that 'other people' exist even if they aren't conscious. What one

worries about, if one ever worries about one's ignorance of whether there are other people, is that one is ignorant about whether *there are other people of the sort that are conscious*. We never put it that way, of course, because it never occurs to us to think that a non-conscious thing could be a person, as Chalmers is suggesting. It is as if the terms have been changed, and we are meant to take solace in how our old skeptical statements sound under these new meanings. The only way to take solace is to equivocate. Anyone who was ever worried about this:

I don't know whether there are other people, because I might be in a tricky simulation created minutes ago!

Should be just as worried about this:

I don't know whether other people are nonconscious parts of a simulation created minutes ago!

Nevertheless, that the currently standard formulation of the skeptical conclusion can be avoided, given structuralism, is still some progress in our understanding of the problem. We owe such progress not only to the nature of new simulation technology, but to one of the brilliant philosophical minds of our time working out its implications in *Reality+*.

Yuval Avnur

yuval.avnur@gmail.com

This review first published online 11 August 2022

Games: Agency as Art by C Thi Nguyen (Oxford University Press, 2020).

doi:10.1017/S0031819122000341

Many of us take great pleasure in playing games and spend considerable time doing so. For many people who play games, playing provides a sort of pleasure that feels unique, that other activities do not offer. At the same time, when we contemplate playing games, we can find ourselves in a state of mind from which games look like a waste of time. Playing games involves trying to do silly and pointless things, like putting balls through hoops, moving wooden tokens around boards, and pretending to manipulate physical objects within entirely virtual digital environments. When I