


Tim Crane

"TAKING SIMULATION SERIOUSLY": TIM CRANE REVIEWS REALITY+ BY DAVID CHALMERS



David Chalmers' new book is a *tour de force*. In about 500 pages of very readable prose, he introduces many of the philosophical problems of our age: the ultimate nature of reality, the problems of knowledge and scepticism, the problems of consciousness and artificial intelligence, the source of value, the structure of a good society; and even the nature of God makes a brief appearance.

Chalmers poses these traditional questions against the background of the idea of virtual reality (VR). If we assume that VR machines develop to be much more sophisticated than they are now, then we can ask questions like: how do you know you are living in the real world as opposed to a simulation? Does it matter whether the "people" you interact with in a VR machine are real? Can things have real value in VR? Does the source of value come from experience or from our actual relations and interactions with people?

The book also contains charming – and unusually pertinent, for philosophy – illustrations by Tim Peacock. Also notable is Chalmers' effortless weaving of his discussion of contemporary problems in Western philosophy with references to other non-Western traditions, and his unforced

introduction of ideas from thinkers (e.g. women) who have been ignored by the tradition. Readable, effortlessly up to date, handsomely produced, and well-structured, *Reality+* will be a boon for teachers of introductory philosophy courses and is sure to remain on syllabuses for many years. I know I will use it as soon as I can.

However, the book is not just an introductory text. It also contains a strikingly original and controversial theme: that we should take seriously the idea that virtual reality is, in itself, a kind of reality. Rather than being a fictional or fantastic representation of a reality which is not there, or a misrepresentation of what is outside the mind, the things we encounter in VR are real – they just have a different underlying nature than other things. The tables you encounter in VR are real tables – but they are made of “bits and bytes” rather than wood and metal. This is the idea he calls “simulation realism”.

The idea will be familiar to readers of Chalmers' 2005 paper, “The Matrix as Metaphysics” in which he developed simulation realism in the context of the famous movie, *The Matrix*, whose plot I do not need to repeat here. But the discussion of the idea in *Reality+* is much clearer, more wide-ranging and deeper. In the rest of this review, I will concern myself with the simulation realism thesis, and the argument for it.

We should begin by distinguishing, first, the hypothesis that we might be living in a VR simulation, from Chalmers' simulation realism. While Chalmers' thesis implies that VR is not a misrepresentation of reality, the mere idea of a VR simulation doesn't imply this. The VR simulation hypothesis is of course a contemporary version of Descartes' sceptical hypothesis in his 1650 *Meditations*, that your current experience is compatible with one's being deceived by a “malin génie”, an Evil Demon, who has (in some way or another) made it seem as if the objects around you are exactly the way they seem to be, even though they are not; and even that mathematical truths are true, when they are not. As Descartes put it, “I shall think that the sky, the air, the earth, colours, shapes, sounds and all external things are merely the delusions of dreams which the demon has devised to ensnare my judgement”.

Of course, Descartes did not think that this *was* our actual predicament, nor did he draw Chalmers' conclusion about the nature of reality; but the initial description of how the appearances are created by some external mechanism is strikingly similar to Chalmers' idea.

This raises the question: what is the *essential* difference between the VR simulation idea and the Cartesian idea? What does the VR hypothesis add to Descartes'? Asking this question is one way to approach the main theme of Chalmers' book. There seem to be two options. We could say that the VR idea is basically the same as the Cartesian one. Or we could say that the VR simulation appeals to new things – new facts, new interpretations – which were unavailable to Descartes, and so presents a radically different argument. Each option raises questions for Chalmers, and the answers will not be straightforward. Perhaps because of this, I have the sense that he himself vacillates somewhat on the answers.

As I shall illustrate below, Chalmers sometimes says that the point of the two scenarios is essentially the same: the VR simulation story is just an up-to-date way of making Descartes' point.

This has a number of consequences. One is that because the Evil Demon story is a sceptical scenario, the VR scenario should be too — and indeed, that is one of the uses to which Chalmers puts the simulation scenario.

In sceptical scenarios, we typically do not have to work out the details. When Bertrand Russell says that our experience is consistent with the world being created five minutes ago, he does not have to explain *how* this might have come about. He does not have to explain who made the world five minutes ago, how they made it, why they would want to, and so on. As Chalmers himself says, when talking about the demon, “Descartes is not clear on the details” (p.53).

It is not a good objection to Descartes’ Evil Demon story to say that that there are no demons; that misses the point.

That’s true. But for sceptical purposes, the details do not matter. If a given sceptical scenario is coherent, then objections to particular parts of the story are irrelevant. It is not a good objection to Descartes’ Evil Demon story to say that that there are no demons; that misses the point. Similarly, it’s not a good objection to Russell’s sceptical argument about the past to say that no one could actually, realistically, create the world with all its apparent history. The point of these arguments is not to create a realistic hypothesis about how the world is, but to ask what it is about our current knowledge and experience that rules out these unbelievable hypotheses. And saying they are unrealistic is not a good objection to scepticism.

One striking thing about Chalmers’ simulation realism – the external world does exist in VR, it’s just very different from the way we think it is – is that it can give a general answer to scepticism about the external world:

Reflection on virtual reality technology can help us respond to the problem of the external world... I argue that if indeed we’re in a simulation, tables and chairs are not illusions but perfectly real objects: they are digital objects that are made of bits. (p.xix)

But if the right conclusion to draw from the VR scenario is that reality might not be what we think it is – it might be fundamentally “digital” – then why shouldn’t we say a similar thing about Descartes’ Evil Demon scenario? Why shouldn’t we say that tables and chairs are not illusions but perfectly real objects, made of parts of the demon’s consciousness?

Chalmers sometimes seems perfectly happy with this parallel argument. At one point he makes the comparison between his simulation scenario, Descartes’ scenario, and Putnam’s famous brain-in-a-vat:

In a perfect simulation, things are perfectly real. The same goes for other Cartesian scenarios, such as Descartes’ evil-demon scenario and Hilary Putnam’s brain-in-a-vat scenario. Generalizing simulation realism to these scenarios, we arrive at the no-illusion view of Cartesian scenarios. (p.119)

And he also spells out the consequence explicitly later on, when he says that “the subject in Descartes’ evil-demon scenario is not undergoing an illusion” (p.122).

But this surely can’t be the right conclusion to draw about Descartes’ Evil Demon scenario. Remember that Chalmers thinks that the mere possibility that we might be living in a VR simulation is a reason to take it seriously as a theory of reality – the idea that reality might be made of bits is supposed to be a serious metaphysical option. But this can’t be the right thing to say about the demon hypothesis. The mere coherence of the scenario gives us no reason to believe in demons; if we were sceptical about demons before (as we should be) then the scenario gives us no reason to change our minds. That isn’t its point.

Of course, some philosophers have argued that the way to fight scepticism is to deny that the kind of error hypothesised by the sceptic is really possible. George Berkeley is an example. Like Chalmers, he thought that tables and chairs “are not illusions but perfectly real objects” – only they are not constituted by bits, but by ideas in the mind of God. Chalmers acknowledges the similarity: “you could think of [Berkeley’s view] as a cousin of my view on sceptical scenarios, with God playing the role of the simulator or evil demon” (p.120).

Berkeley was an idealist, not just in the sense that he thought that reality is fundamentally mental (“ideas”), but also in the sense that he thought that there cannot be a massive gap between appearances and reality (“esse est percipi”). These are two distinct doctrines. In the first sense, idealism is contrasted with dualism or materialism, and in the second sense it is contrasted with realism. The second contrast is about the extent of the possible gap between appearance and reality. Realists say that we could in principle be very wrong about how reality is – it is in principle coherent to suppose that the ordinary objects around us do not exist. Scepticism presupposes realism, so understood.

The comparison with Berkeley suggests that Chalmers’ simulation realism is a kind of idealism in this second sense. This is because a certain kind of systematic error is ruled out. Curiously, he seems to be unaware of this, showing more tolerance of bizarre doctrines like panpsychism than of this kind of idealism:

I’ll argue that some forms of idealism should be taken seriously at least as a speculative hypothesis. We can’t rule out that consciousness underlies the universe. However, I think that any version of idealism that rests on equating our appearances with reality is doomed. (p.71)

But it is something like this version of idealism to which simulation realism seems to be committed. This form of idealism does not exactly *equates* appearances with reality, true. However, if the simulation, brain-in-a-vat, and demon scenarios are all explained in the same way (in terms of the “no illusion view”), then Chalmers’ preferred metaphysics does not allow the possibility of a radical break between appearance and reality that realism allows. In fact, it is hard to see how Chalmers can construct any global sceptical hypothesis which is not answered by his “no illusion view”. His strategy removes a central support for the sceptical argument: the possibility of a radical gap between appearance and reality. Chalmers writes that “if the analysis works, it dissolves what is perhaps the Western tradition’s prime reason for doubting that we can know

anything about the external world" (p.xx). But in fact, it removes all such reasons. Idealism makes sceptical doubt incoherent.

This is not supposed to be an objection to Chalmers' "no illusion" view – it should not be an insult to call someone an idealist, even these days. My point is rather to clarify the extent of his commitments, and the extent to which he does not seem to be aware of them.



So far I have been assuming that Chalmers has taken the line that the VR simulation scenario is essentially the same as the Cartesian scenario. But maybe this is wrong. Maybe the VR story *does* add something genuinely new, perhaps a more detailed and more plausible story about how our underlying reality might be structured. As Chalmers himself says,

The simulation idea does more than illustrate the problem, however. It also sharpens the problem by turning Descartes' far-fetched scenarios involving evil demons into more realistic scenarios involving computers – scenarios we have to take seriously. (p.xix)

When he says "take seriously" I assume he means *take seriously as something that could really happen to us* (unlike being deceived by an evil demon). This moves it beyond the category of sceptical hypotheses – since, if I am right, these hypotheses do not have to be taken seriously in that sense. If so, then the problem posed by VR is not just a sceptical one, but the sort of problem forced on us by changes in technology and the non-philosophical world, like the problems posed by climate change or the internet, and so on. For Chalmers, this does mark out a difference with the earlier sceptical scenarios:

You might think that the switch from evil demons to brains in vats to simulations is a mere change in packaging, but there is one respect in which the use of modern technology makes the argument more powerful....The simulation hypothesis may once have been a fanciful hypothesis, but it is rapidly becoming a serious hypothesis. (p.53-54)

After all, we know that VR machines exist; all Chalmers is asking us to imagine is that they get a lot more sophisticated than they are now. So we may have to take seriously – not just as a sceptical hypothesis, but *really* seriously – the idea that we are living in a simulation. In this, Chalmers follows the speculations of Nick Bostrom and others.

We know that VR machines exist; all Chalmers is asking us to imagine is that they get a lot more sophisticated than they are now.

But if this is so, then the details do matter. More precisely, it matters that VR should be able to create a realistic simulation of all aspects of our experience. And this is a substantial empirical assumption which needs justification. For example, if a VR experience includes the experience of interacting with other people, then a computer must be capable of producing experience that is in all respects indistinguishable from the interaction with a real person. This is a big assumption about the potential of Artificial Intelligence (AI). As researchers in AI know, this is a huge task, and it is not obvious that AI, as it currently is, can produce such a thing – the impressive efforts of the text producing programme GPT-3 notwithstanding (there is an enjoyable discussion of GPT-3 in Chapter 13 of *Reality+*. but I will not give out spoilers here).

Just consider, for example, the experience of interacting with a voice assistant like Siri or Alexa. One thing that is so impressive about these machines is how they can recognise spoken language. What is less impressive is the “conversations” that result. Often such conversations end up with a suggestion like “here are some web pages related to your inquiry”. This is fine if the point of conversation were always to find information, or the answer to a question. But that plainly isn’t the general point of conversation — in fact, it’s not obvious that conversation has a single point at all.

If a VR machine really can replicate all aspects of our experience, then it must be able to replicate genuine conversations (among many other things). And in order to do this, it needs to solve the biggest problem for AI research, the problem of artificial general intelligence (AGI). While AI machines have succeeded in exhibiting specific intelligence in clearly defined tasks (e.g. chess, GO), no one currently has any idea about how to create the kind of intelligence we have – general intelligence, also known as common sense.

Chalmers acknowledges this in *Reality+* when he writes that “no existing digital system is yet close to general intelligence” (p.279). Elsewhere, however, he is more sanguine. In a *New York Times* interview he is quoted as saying that “artificial general intelligence is possible” and that although “there are a lot of mountains we need to climb before we get to human-level AGI”, “it’s

going to be possible eventually, say in the 40-to-100-year time frame". However, most of the discussion of AI in *Reality+* concerns whether an AI machine can be conscious, and not AGI and its relevance to the real possibility of a total VR simulation. Yet the question for the uncommitted outsider is why we should believe that there realistically can be such a simulation, if we do not know how to solve the problem of general intelligence?

The dilemma for Chalmers, then, is this. Either the VR simulation is just an updated version of the Cartesian scenario, in which case we do not have to work out any details. That's fine, but then we are in familiar territory, and all the interest lies in the consequences of simulation realism. Alternatively, the VR simulation *does* add something new, but then the philosophy is hostage to the problems surrounding AI and the creation of AGI. These seem like very different directions; I think Chalmers does not sufficiently distinguish them in *Reality+*.

Chalmers is always an inventive philosopher. He has helped shape the current philosophical landscape in the philosophy of mind and consciousness. With a vivid image – the zombie – or a phrase – the “hard problem” – he has set the agenda for recent debates about the mind-body problem, as well as about other things, like the relationship between possibility and conceivability. While his effort to recreate Rudolf Carnap's *Aufbau* in *Constructing the World* (his published John Locke lectures given in Oxford in 2010) was less influential than his work on consciousness, I expect that *Reality+* will set various agendas in metaphysics just as *The Conscious Mind* (1996) did in the philosophy of mind. And it is all the more impressive that he does this in a book much of which can be understood by beginners.

I will end with a critical comment arising from this comparison with *The Conscious Mind*. The arguments of his 1996 book depended on starting points which might seem simple or obvious or uncontroversial at first – such as the distinction between hard and easy problems, or the very idea of a zombie – but which in fact embody a lot of philosophical assumptions, some of them very complex. Interrogating these assumptions should be part of the foundational work in the philosophy of mind. But very often they are just accepted – “we all know what a zombie is” – and then the complex dance through the various positions begins. To paraphrase Wittgenstein, the problem is that we have already committed ourselves to a certain way of looking at the matter.

The same is true of *Reality+*. While it might seem a fairly straightforward thing to imagine living in a simulation – “it's just like the Matrix!” – there are a number of questions that need to be settled before the consequences can be properly seen. My reason for focussing on the foundational questions about simulations and scepticism is that a lot turns on whether you actually accept these starting points. Before looking at the consequence of the VR simulation hypothesis, we need to know what kind of hypothesis it is, what assumptions it is based on, and what independent plausibility they have. This requires that we move a bit more slowly than Chalmers sometimes allows us.

Reality+: Virtual Worlds and the Problems of Philosophy by David Chalmers is published by Allen Lane.

*Tim Crane is professor of philosophy at the Central European University (CEU) in Vienna. He works in the philosophy of mind, and his research attempts to address questions about the most general nature, or essence, of the human mind, and about the place of the mind in the rest of nature. His books include *The Mechanical Mind* and *The Meaning of Belief*.*

Website: timcrane.com

Twitter: [@timcrane102](https://twitter.com/timcrane102)

From *The Philosopher*, vol. 110, no. 2 ("The New Basics: Society").

If you enjoyed reading this, please consider [buying a copy of this issue](#) or [making a donation](#).

We are unfunded and your support is greatly appreciated.