

Simulation Scenarios: Reply to Chalmers

Peter Godfrey-Smith

This is a brief response (3 pages) to David Chalmers's comments on my "Simulation Scenarios and Philosophy," which is a commentary on his book *Reality+*. His comments appear in "Taking the Simulation Hypothesis Seriously," which also includes replies to commentaries by Eric Schwitzgebel and Susan Schneider. The whole collection is in *Philosophy and Phenomenological Research* (currently in "Early View," 2024).

I'll reply to points in the order of their presentation by Chalmers. The opening quote in each numbered passage is from the Chalmers paper.

1. "Godfrey-Smith takes the simulation hypothesis seriously enough to offer a number of reasons not to take it seriously" (p. 1).

I don't take the simulation hypothesis seriously as a view about our actual situation – how things might be for us. I am more interested (interested enough to write) in the fact that other people do seem to take it quite seriously. Chalmers gives the hypothesis that we are living in a simulation a probability of "at least 25 percent or so" (p. 100 of the book). That is a big number, surely big enough to affect practical decision making in low-feedback contexts. If we're in a simulation, would it make sense to give money to environmental charities to protect places you are sure you'll never visit? I guess this will depend on what kind of simulation a person might take themselves to be living in (how that 25% is divided between simulation possibilities). But a lot of low-feedback decisions (voting, charities, wills...) will surely look different.

2. "Godfrey-Smith suggests that some experiences may be hard to simulate this way, mentioning hot showers and skateboard crashes. Perhaps the thought is that the body plays a special role in these experiences [yes - PGS], and it affects the brain in a way that goes beyond standard sensory inputs and outputs [no - PGS]. This thought might be fleshed out by asking: to undergo these experiences, does the brain have to be warmed up and jolted around, or is it enough for it to represent warmth and motion?"

Whole-body experiences are special in this setting, yes (here we are discussing the embodied or "biosim" possibility). The skateboard crash includes jolts, scrapes, pains, the feeling of tangled limbs and being momentarily upside down, of skin being lost... These are "sensory inputs," but they are reflecting influences on a great range of

transducers spread all over the body. Maybe this this can be simulated, but the model of vision that informs so much discussion of VR ("Soon, we won't be able to tell it from ordinary reality!") is misleading here. (Along with the skateboard crash and hot shower, I add as another VR-resistant experience the feeling of getting too close to a fire and breathing in smoke.)

3. "[M]y arguments about the simulation hypothesis do not depend on simulations that are perfect in this sense [perfect simulation of some actual person]. What matters is that simulations be humanlike. This does not requires perfectly simulating any existing human being" (p. 3).

The discussion of "perfect" simulation in my paper did not say that the simulation hypothesis requires that there be perfect simulations of actual people. (Obviously it does not.) The discussion in my paper (pp. 3-4) criticized claims made about perfect simulations by Chalmers himself, which are used by him as premises in arguments for the possibility of creating human experience in pure sims. These are arguments for the "substrate-neutrality" of the mental. I am criticizing passages like this one in Chalmers's book (quoted in my paper):

We can suppose that every neuron is simulated perfectly, as is every glial cell and other cells throughout the brain. The interactions between neurons are simulated perfectly, too. All the electrochemical activity is simulated, and so is other activity, such as blood flow. If there's a physical process in the brain that makes a difference in how the brain functions, it will be simulated. (pp. 287-88)

Also (not quoted in my paper):

Once we have fine-grained simulations of all the activity in a human brain, we'll have to take seriously the idea that the simulated brains are themselves conscious and intelligent. After all, a perfect simulation of my brain and body will behave exactly like me. (p. 23, emphasis added here and above)

Chalmers: "In any case on[c]e we recognize that perfect simulation is not required for the purposes of the simulation hypothesis, Godfrey Smith's worries about that notion fall away." My aim, rather, is to show that some (probably the main) reasons people give for believing in substrate neutrality are not good reasons. This bears on whether we might be "pure" sims in Chalmers's sense, rather than biological beings in a Matrix-like situation.

4. "[T]he reasons that Godfrey-Smith gives for not taking the simulation hypothesis seriously come down to two sim blockers: (1) *Conscious humanlike sims are impossible* (deriving from worries about brain simulation), and perhaps (2) *Nonsims will not create many conscious humanlike sims* (deriving from worries about simulator motives). I think these sim blockers are less likely than Godfrey-Smith does. But even he should allow that the denials of (1) and (2) are both serious and grounded hypotheses that deserve some substantial nonzero credence. If we give these denials even 10% credence each, then (assuming independence) this still leaves us with a 1% credence that most humanlike beings are sims" (p. 4).

I'd give 10% probability to the view that "conscious humanlike sims," in a broad sense, are possible – that roughly humanlike experience can be produced in this way. What is the chance that it is feasible, and someone bothers, to create many billions of simulated people and their apparent environments in a fine-grained enough way to include all their hot showers, skateboard crashes, smoke inhalations, and the like? (The two conditions, feasibility and motivation, are distinct but connected. As the project gets harder to achieve, more commitment – along with the peculiar malaria-embracing malice that Schwitzgebel discusses in his own commentary on the book – would be needed.) How likely is someone to create "many conscious humanlike sims" *of the relevant kind*, and in the huge numbers needed to mount a probabilistic argument? I think this is *extremely* unlikely. Not 10% or 1% unlikely.

5. "Godfrey-Smith says "To seriously suspect you are in a simulation is to suspect that many or all of your memory traces might be cooked up, and your background knowledge is no good at all". Of course I reject this claim. Once one accepts simulation realism, one can suspect that one is in a simulation without doubting one's memory or background knowledge" (p. 4, emphasis added).

You *can*, certainly, but is there any reason to do so – to trust memory traces and background knowledge that, by your own lights, are under the control of hidden agents with unknown motives?