

The Simulation Hypothesis, Social Knowledge, and a Meaningful Life

Grace Helton

Draft of 12 February 2023

please cite official version and see the relevant issue for Chalmers' reply

For *Oxford Studies in Philosophy of Mind*

Symposium on David Chalmers' *Reality+: Virtual Worlds and the Problems of Philosophy*

Surely some of my views about what the world is like are wrong, perhaps even in rather surprising ways. But could my most basic assumptions about the world be incorrect? As I write this, I think I'm in a quiet room with a table, some lamps, and a laptop, rain pattering on the roof, my dog sprawled languorously by my side. I think my hands cradle a ceramic mug of hot tea and that a thin mist of steam, faintly smelling of jasmine, curls over the lip of the mug. Could I be wholly hallucinating this entire state of affairs, perhaps as someone in the throes of a psychotic episode? Or perhaps as the plaything of some evil demon who deceives me for its own amusement? And if I am undergoing a psychotic episode or am manipulated by a demon, would I thus be wholly mistaken about what the world is like?

This concern that reality might be radically different than it appears to be might seem to be a rather arcane one, the musings of someone who is either extremely paranoid or overly taken with the abstruse.¹ For, unless I have some *reason* to think that I am undergoing a psychotic episode or that I am being manipulated by some demon, why should I think twice about such speculations? As it turns out, there is at least one version of this sort of claim—that reality is radically different than we think it is—which enjoys at least some empirical support. This is *the simulation hypothesis*, the claim that we and everything in our environment is realized by a large-scale computer simulation, one implemented by superintelligent artificial intelligence (AI).²

The empirical argument for the simulation hypothesis is due to the philosopher Nick Bostrom, and it runs like this: There is some reason to think that AI will advance to the point of having incredibly fast processing speeds and incredibly powerful processing capabilities of the kind which would easily permit them to simulate full-scale human civilizations,

¹ Cf. Schwitzgebel (2017).

² This follows Chalmers' way of construing the hypothesis, as against a more minimal version which is silent as to who (if anyone) created the simulation (Chalmers 2022: 29, Cf. Bostrom 2003).

complete with billions of conscious creatures. There is also some reason to think that at least some such superintelligent AI would be interested in simulating many such worlds, for instance, for entertainment or research purposes. Putting these claims together, there is at least some reason to think that many full-scale simulations of human civilizations will ultimately be created, so many, in fact, that out of all conscious humans who have ever lived or who will ever live, the vast majority will be simulated. But if this is the case, then there is at least some reason to think that *we* and our world are simulated. For statistically speaking, if most humans who have ever lived or will live are simulated, what are the chances that we are among the minority who are not simulated?³

The simulation argument does not conclusively establish that we are living in a simulation. However, the argument gives us some reason to take the hypothesis seriously beyond its mere conceptual possibility. In this way, the simulation hypothesis is unlike the hypotheses that my experiences are wholly generated by psychosis or by an evil demon; I have no reason to take these explanations seriously beyond their bare possibility, so I am plausibly justified in dismissing these hypotheses out of hand. In contrast, I am not justified in dismissing the simulation hypothesis out of hand.⁴

It is in this broader intellectual context that David Chalmers takes up several philosophical questions about both the simulation hypothesis and VR more generally in his timely, extremely accessible, and impressively wide-ranging book, *Reality+: Virtual Worlds and the Problems of Philosophy*. While Chalmers does not endorse the claim that we are living in a simulation, he thinks the claim probable enough to merit the question: What would it matter if we were (Chalmers 2022, 102)? More particularly, Chalmers focuses on questions such as: If we are indeed living in a wide-scale computer simulation, would our views about the world be correct in at least some basic respects? Could these views further be said to amount to knowledge? And would the fact of our world being a simulation in any way diminish the value of our lives? Chalmers offers soothing answers to these questions. Namely, he argues that in the simulation: Our views about the world would still be correct in at least some basic respects; these views might further amount to knowledge; and our lives might still be deeply meaningful (Chalmers 2022, 105-224; 399-422; 440-62; 311-50).

Chalmers employs a *structuralist* strategy to argue that we can enjoy *non-social* knowledge in the simulation; here, non-social knowledge is empirical knowledge of non-

³ Bostrom (2003). Chalmers cites Hans Moravec's (1993) work as a precursor to Bostrom's (Chalmers 2022: 83). Philosophers, computer scientists, and physicists who consider the simulation hypothesis include: Arvan (2014, 2015), Beane et al. (2014), Campbell et al. (2017), Dainton (2002, 2012, 2020), Johnson (2011), and Mizrahi (2017). For criticisms, see, e.g., Weatherson (2003) and Summers and Arvan (2022).

⁴ This is to say, one cannot dismiss it on 'Moorean' grounds (Chalmers 2022, 79-80).

minded things, such as atoms and shrubs.⁵ Structuralism says that since the causal structures of atoms and shrubs exist in the simulation, then atoms and shrubs exist in the simulation. Chalmers further suggests that when we interact with these causal structures, we can gain knowledge of the entities they comprise. Notably, this strategy does not extend to the psychological states of others, a limitation Chalmers acknowledges. In other words, *for all structuralism says*, the seemingly sentient creatures in your environment—your friends, neighbors, animal companions, and the like—might be non-conscious automata. These others might in fact be sentient, but structuralism gives us no reason to think they are.

Chalmers views the claim that we enjoy non-social knowledge in the simulation and the claim that our lives might be meaningful in the simulation as at least weakly connected, as follows: The former claim helps forestall a concern that if objects in the simulation are not genuine (and so not knowable), then life in the simulation is illusory and therefore, not as valuable as a non-simulated life (Chalmers 2022, 314).

In this paper, I will explore the fuller extent of the connection between non-social knowledge, on the one hand, and a meaningful life, on the other. I will suggest that, while non-social knowledge can contribute to the meaningfulness of otherwise *meaningless* lives, in at least many cases, non-social knowledge contributes either nothing at all or very little to the meaningfulness of otherwise meaningful lives. On the overall picture that emerges, for many lives, the value of social knowledge for a meaningful life *dramatically swamps* the value of non-social knowledge for a meaningful life. I call this *the social swamping view*.⁶

I first briefly describe Chalmers' structuralist approach and its limitations with respect to knowledge of other minds (§1). I then argue that in many cases, the value of social knowledge for a meaningful life dramatically swamps the value of non-social knowledge for a meaningful life. Along the way, I propose a *non-additive model* of the meaningfulness of life, according to which the overall effect of some potential contributor of value to a life depends in part on what is already in a life (§2). I close with some reflections on the prospects of vindicating social knowledge against a background in which the simulation hypothesis is treated as feasible (§3).

⁵ Chalmers sometimes uses the term *ordinary physical knowledge*, where I use *non-social knowledge* (e.g., Chalmers 2022, 500-501). I prefer *non-social knowledge* so as to avoid the suggestion that knowledge of other minds is not physical knowledge or else is not ordinary knowledge.

⁶ I say 'most agents' because some agents will prefer above all else to live lives of inquiry into the aspects of the non-social world, in a context of social isolation. While my broader assumptions about meaningfulness in life do not treat meaningfulness as straightforwardly a matter of desire satisfaction, it does accord agents' values a central place and so, I am open to the possibility that for some agents, this kind of life might be deeply meaningful. Thanks to Chris Register for this example and for discussion on this point.

Before proceeding, a point of terminology: I have already been using *social knowledge* to refer to knowledge of those aspects of the world which either are themselves made up of or which depend on other minds. Those who presume social reality to be wholly independent of other minds are free to substitute another term, such as *other-mind-dependent knowledge*. Nothing should turn on the term employed.

Section 1 Structuralism and The Vindication of Non-Social Knowledge

As mentioned, Chalmers exploits a structuralist view of entities in order to argue that we have non-social knowledge in the simulation. This is roughly the view that entities are equivalent to certain observable *causal roles*, i.e., to a kind of *causal structure*. According to this view, what it is to be a mug (or a quiet room or a quark) is to play a certain observable role or more particularly, to tend to instantiate certain patterns of cause and effect. For instance, what it is to be a mug is to be disposed to be usable for holding liquid and for drinking liquid (among other things). What it is to be a quiet room is to tend not to contain noises above a certain decibel (among other things). Being liquid and being above a certain decibel are in turn construed in terms of relevant observable patterns of cause and effect.⁷

In drawing out the implications of this view for scenarios such as the simulation scenario, Chalmers has brought out something important and almost entirely overlooked in the vast philosophical literature on such scenarios: These scenarios might be populated by *genuine* quiet rooms, mugs of tea, and dogs, and not merely their simulacra. In particular, structuralism permits a different assessment than familiar semantic externalist approaches, on which the *terms* ‘quiet room,’ ‘mugs of tea,’ and ‘dog’ have references in the simulation. Because the structuralist maintains that what it is to be (say) a mug of tea just is some causal role, there are genuine mugs of tea—and not, merely ‘mugs of tea’—in the simulation.⁸ The fact that this solution is extremely simple in its basic form whilst affording a radical shift in thought is a testament to its philosophical power.⁹

Despite the power of the structuralist view, it has its limits and, as Chalmers himself repeatedly stresses, the view does not vindicate knowledge about the existence or nature of the psychologies of others. Very roughly, the reason is that, for reasons well-known from

⁷ This is ontic structuralism, not epistemic structuralism (Chalmers 2022, 145–82). See also Chalmers (2005, 2017) for discussion. Chalmers’ argument is strictly neutral between these versions, as he suggests that virtual realism might be grounded in either one (Chalmers 2022, 405–22). Elsewhere, he draws on a conceptual variant of structuralism to develop an argument against external world skepticism (Chalmers 2018).

⁸ That is to say, structuralism is not semantic externalism, though Chalmers himself thinks semantic externalism can help vindicate some forms of knowledge (see, e.g., Chalmers 2022, 372–84).

⁹ In order to find a true antecedent to this view, one must go back to the work of Bouwsma (1949). For discussion, see Chalmers (2022, 120–23).

20th century philosophy of mind, others' mental states are not reducible to third-personal observable roles, whether or not things such as mugs are so reducible.¹⁰ Rather, for all structuralism tells us, the seemingly sentient creatures around us might be automata, much like fully multi-modal, hologram versions of the iPhone's talking "Siri."

Elsewhere, I have argued that: the initial suspicion that structuralism cannot vindicate knowledge of other minds is correct; this is so *even if* structuralism is combined with a sophisticated kind of functionalist and wholly materialist view of the mind; this result has implications for wide swathes of beliefs across domains, including at least some beliefs about political history, aesthetic movements, and cultural practices; and as a result, structuralism cannot give us the wholly satisfying solution to skepticism we might've hoped for (Helton forthcoming).¹¹

Needless to say, the question of whether structuralism can vindicate knowledge of other minds is a fraught issue and not one I can properly draw out here. For present purposes, what matters is that both Chalmers and I both think structuralism is ill-suited to vindicate social knowledge, even if it can vindicate non-social knowledge. This shared presumption naturally sets up the question I consider next: What is the significance of non-social knowledge versus social knowledge for a meaningful life? Notably, this question is of general interest, regardless of one's commitments to structuralism.

Section 2 Social Knowledge and Meaningfulness in Life

'Nothing real can be threatened' -Beyoncé

In this section, I will argue that, for at least many lives, the value of social knowledge *massively swamps* the value of non-social knowledge in the meaningfulness of those lives.¹² I call this *the social swamping view*. I will explore a stronger and weaker version of this claim, concluding that at least one of these claims is true.

First, consider the strong version of this claim, specifically:

Social Swamping View (Strong)

¹⁰ This point is original with Lovelace (1842), as discussed by Turing (1950, 450-1; 454-60). See also Block (1978, 1981). Chalmers (2022, 459-60; 500-501).

¹¹ See also Helton (2021, 242-46); Cf. Chalmers (2022, 500-501).

¹² Special thanks to Liam Kofi Bright and Sarah McGrath for conversation on this point. For recent helpful overviews on meaningfulness in life, see Mawson (2013), Seachris (2019), and Thaddeus (2022).

There are at least some meaningful lives, meaningful partly in virtue of their socially rich aspects, which are such that lacking knowledge about non-social reality detracts *not at all* from the meaningfulness of those lives.

On its face, this claim seems obviously false. For, one might think that this claim entails, rather implausibly, that non-social knowledge lacks value. And as against *this* claim, one might suggest a thought experiment along the following lines: Suppose there is some human who tragically lacks most candidate elements of a meaningful life. Her desires are routinely thwarted, her relationships are missing or disingenuous, and she cannot perform authentic actions. Still, despite all of this, her beliefs about non-social aspects of reality, such as her belief *there is a shrub over there* or *that's the ocean*, largely amount to knowledge. Intuitively, this knowledge makes her life at least a tiny bit more meaningful than it would be were she to altogether lack such knowledge. In light of these considerations, one might suggest the following claim:

No Wholly Meaningless Lives with Non-Social Knowledge

There are at least some lives which are such that knowledge about the non-social realm can contribute at least a bit to the meaningfulness of those lives.

As it turns out, I think that this thought experiment is apt, and that it is true that some lives which would otherwise be wholly meaningless gain a bit of meaning from their bearers having some knowledge of non-social reality. But, this claim isn't in conflict with the claim I started with, which is the claim that some meaningful lives are wholly undiminished by a lack of non-social knowledge. Understood as existential claims, these do not form a contradiction.

The appearance of conflict stems, I suspect, from an implicit model of how potential contributors of value create meaningfulness in a life. If we accept a model on which meaningfulness is a matter of combining valuable things together, where each contributor makes its contribution independently of what else is in a life, the claims considered conflict. For, on this *additive model*, if knowledge about non-social reality can ever make a difference in the value of a life, this knowledge always makes a difference, regardless of whatever else is in that life. So, on this view, if non-social knowledge can make a life which is otherwise devoid of value a bit more meaningful, then non-social knowledge can also make a very meaningful life a bit more meaningful.

But, one needn't adopt an additive model of the meaningfulness of life. Instead, one might adopt a *non-additive* model, according to which the overall effect of some potential contributor to a life's meaning depends in part on what else is already in that life. In this way, a meaningful life might be a bit like a stew. While adding a little salt might dramatically

improve the flavor of a plain broth, adding a little salt might not improve *at all* an already delicious stew, one with powerful and distinct flavors. While a dash of salt necessarily changes the *ingredients* in the stew, a stew's ingredients can change without any change to its taste. Likewise, the model under consideration is the non-additive or, if you like, 'stew' *model* of the meaningfulness of life. The value of each potential contributor depends on what else is there, and in some cases, an otherwise valuable contributor can be 'canceled out' by what else is there, such that it becomes wholly irrelevant to a life's overall meaning.¹³

Once we appreciate that life's meaningfulness might not be an additive matter, new possibilities come into view. For instance, consider someone who has a very good life. By this I don't mean an unceasingly pleasurable life, but a rich and authentic one. Let's call her Alya. Alya has safety, shelter, food, and other basic necessities, but also music, art, love, and community. Her work is valuable and rewarding; her relationships are not without challenges, but they are intimate and reciprocated. Maybe Alya is madly in love with someone she just met; maybe she enjoys the companionship of several decades-old, platonic relationships. Maybe she engages in the rituals of an ancient religion; maybe she follows no religion and surfs a lot.

Suppose that Alya's beliefs about other minds and broader social beliefs are not just true but constitute knowledge; her beliefs about the inner states of others are correct, and those creatures around her—whether human or some other species—have inner lives, just as she supposes they do. Suppose further that few or none of Alya's beliefs about the non-social realm constitute knowledge. So, for instance, beliefs of hers such as *that's a shrub* and *there are atoms* somehow fail to amount to knowledge.¹⁴ Would this lack of non-social knowledge necessarily detract from the meaningfulness of Alya's life, which is replete with authentic relationships, valuable work, and rich social knowledge?

If we presume both that non-social knowledge is valuable and that the additive model of life's value is correct, then we must say 'yes.' While Alya's life is meaningful, it's necessarily *less* meaningful due to this lack of non-social knowledge. But, if we dispense of the assumption that meaningfulness is additive, we have room to say 'no,' Alya's life is *no less meaningful* due to this lack of non-social knowledge. And, this might be so even if such knowledge is itself valuable.

¹³ See Moore (1903) for the claim that the value of a whole is not the sum of the value of its parts. Cf. Zimmerman (1999). Special thanks to Daniela Dover for discussion on this point. Notably, the non-additive model is consistent both with the view that knowledge in general has intrinsic value and with the view that the value of knowledge is merely instrumental. For relevant discussion, see, e.g., Kelly (2003) and Rinard (2019).

¹⁴ The structuralist will have to say either that the relevant causal structures do not obtain in Alya's environment or else that she for some reason does not track them properly.

On the view sketched, a life that is replete with, for instance, authentic expression of one's values, genuine agency, rich interdependence with other creatures, and vast swathes of social knowledge, might be *undiminishable* in a certain way. Lacking knowledge about things like shrubs and atoms cannot degrade such a life even one iota. This view treats certain aspects of a meaningful life as anchors, in the sense that, once present, these elements protect a life's value from certain forms of axiological unmooring.

One might object to this view by drawing a contrast case: Consider Alya's counterpart, Kalya. Kalya's life is just like Alya's except that Kalya's beliefs about non-social aspects of the world largely amount to knowledge. While, by stipulation, neither Alya nor Kalya much cares about whether their non-social beliefs amount to knowledge, we might still be tempted to say that surely, Kalya's life is a bit more valuable than Alya's. For one thing, not only does Kalya enjoy knowledge about the non-social world, she also enjoys certain forms of integrated ecological-social knowledge, which Alya lacks. For instance, Kalya might know that she went hiking with her boyfriend over the weekend and that together they sat on a large rock for a while, basking in the sunrise. Surely Kalya's life is at least a bit more meaningful than Alya's in virtue of having this sort of ecological-social knowledge.¹⁵

I think there is room to resist the suggestion that Kalya's life is necessarily more meaningful than Alya's, even whilst acknowledging the value of non-social knowledge in general. Merely think for a moment of the kind of rich life Alya has. She has close and genuine relationships, she can achieve many of her ends, she has means of expressing herself, and she doesn't much care whether she knows about things like shrubs and atoms. We might think that to suggest that Alya's life is, despite all this, improvable by non-social knowledge is to show a disrespectful attitude toward both Alya's own preferences and also toward the sources of objective value in her life.

I am not sure whether there are some lives whose meaningfulness cannot be improved at all by non-social knowledge, but I hope to have shown this view should not be dismissed out of hand. I will now argue for a more moderate cousin of this claim. If either of these claims is true, we should think, at a minimum, that the value of social knowledge for the meaningfulness of our lives *dramatically swamps* the value of non-social knowledge for the meaningfulness of our lives:

Social Swamping View (Weak)

There are at least some very meaningful lives, meaningful partly in virtue of their socially rich aspects, which are such that lacking knowledge about non-social reality *scarcely* detracts from the meaningfulness of those lives.

¹⁵ I thank Josh Armstrong for this example and for helpful discussion on this section.

This claim is consistent with the thought that Alya's life would be more meaningful if she had non-social knowledge, such as ecological knowledge. But, this claim is inconsistent with the thought that Alya's life could be dramatically improved by that knowledge. Due to the presence of things such as extremely rich social relationships and agency in her life, the addition of such knowledge would confer at most a modest increase in meaning. (This view requires a non-additive model of life's value, on which the effect of a potential contributor of value can be blunted by what else is there, even if not canceled out entirely).

In favor of this weaker claim, I would point to the likely emotional responses many of us would have in response to the loss of non-social versus social knowledge. Suppose God herself were to tell us that, while all of our beliefs about the psychologies of others and our own agency amount to knowledge, few or none of our beliefs about things like atoms and shrubs amount to knowledge, even though we can exploit these beliefs to carry out our aims. For many of us, we'd be shocked by this disclosure, and we'd likely have many questions. But, I suspect the overall emotional response after the initial shock would be, if not indifference, something like *curiosity*, a positively-valenced emotion. Indeed, some of us might be *delighted* at this disclosure, as it would introduce an element of wonder or mystery into our everyday life.

In contrast, suppose God herself were to tell us that, while all of our beliefs about things like atoms and shrubs amount to knowledge, few or none of our beliefs about the sentience or psychologies of others amount to knowledge. For many of us, wondering whether those around us—our friends, family, colleagues, neighbors, animal companions—are sentient at all would cause us to be grief-stricken and horrified. Indeed, it would be understandable for us to wonder whether our lives had ever been worth living or were worth continuing. This disclosure would be almost unspeakably horrific, cutting to the very center of the value of our lives.¹⁶

I submit that this dramatic emotional asymmetry has one good explanation: For very many of us, the value of social knowledge for the meaningfulness of our lives far outstrips the value of non-social knowledge. So, at a minimum, we should accept the weak claim described above: At least many lives made meaningful through their social richness can scarcely be diminished by a loss of non-social knowledge.

Section 3 Social Knowledge in the Simulation

Let's bring these reflections back to Chalmers and to the question of the relation between non-social knowledge and meaningful lives in the simulation. I think Chalmers is

¹⁶ See Schwitzgebel (2017, 280–82, 284) for a different but likewise sanguine assessment of the loss of non-social knowledge and Schwitzgebel (2017, 285–87) for a somewhat different take on the loss of social knowledge.

right that if our non-social knowledge is intact, this helps establish that in the simulation, our lives are not wholly *meaningless*. At the same time, I would suggest that the structuralist vindication of non-social knowledge, absent any correlative vindication of social knowledge, contributes not at all or *scarcely* at all to the claim that in the simulation our lives might be *very meaningful*.

In making this claim, I'm not sure whether Chalmers would disagree with it. Given his own emphasis on the value of interpersonal relationships and community in a meaningful life, it is possible that he would agree that the contribution of non-social knowledge to a meaningful life is relatively minimal (Chalmers 2022, 319; 329-30). However, Chalmers and I *might* disagree about the prospects of vindicating social knowledge in the simulation. For, he seems open to the view that knowledge of other minds in the simulation might be vindicated on broadly abductive grounds, for instance by generalizing from relevant neural or behavioral states (Chalmers 2022, 286-87).¹⁷ In contrast, I see this kind of strategy as encumbered, which isn't to say I think it could not possibly succeed.

Specifically, if the simulation hypothesis is at all feasible, certain other hypotheses are also feasible (albeit to a lesser extent) according to which at least some of those around us lack sentience. For instance, our AI creators might have had ethical quandaries about simulating eight billion creatures, many of whom will live foreseeably horrible lives; this possibility is heightened if we presume that pre-simulated humans solved the AI safety problem (Helton 2021, 237-38).¹⁸ Or, on purely practical grounds, our AI creators might have found it unduly burdensome to simulate the psychologies of eight billion humans, since doing so might have been costly in terms of processing power. So instead, they might have created some sentient creatures and rendered the rest as convincing but ultimately mindless automata. Thus, the simulation hypothesis introduces novel impediments to an abductive solution to other minds, which is not to say whether these impediments are ultimately unmovable.

I opened this paper with a description of my surroundings—a quiet room with a table, some lamps, and a laptop, steam rising from my mug, rain pattering on the roof, my dog sprawled by my side. I asked: Do I know I am really here, surrounded by these things? What I'd now like to ask is, in terms of the meaningfulness of my life, *does it matter* whether I know that I am here, surrounded by these things?

The answer I have suggested is that my knowledge of the tea's rising steam and my knowledge of my dog's relaxed psychological state are not on a par. If my life is otherwise devoid of meaning, my knowledge about the steam can contribute to the meaningfulness of my life. But, if my life is otherwise replete with sources of meaning, this knowledge about the

¹⁷ See also Schwitzgebel and Moore (2015).

¹⁸ See also Schwitzgebel (2013) and Schwitzgebel (2019, 431-33).

steam might not contribute at all to the meaningfulness of my life. In contrast, my knowledge that my dog is a sentient creature, one who currently feels relaxed, contributes to the meaningfulness of my life regardless of what other sources of value I have in my life.¹⁹

Works Cited

- Arvan, Marcus. 2014. "A Unified Explanation of Quantum Phenomena? The Case for the Peer-to-Peer Simulation Hypothesis as an Interdisciplinary Research Program." *The Philosophical Forum*, 45(4), 433-446.
- Arvan, Marcus. 2015. "The Peer-to-Peer Simulation Hypothesis and a New Theory of Free Will." *Scientia Salon*. <https://scientiasalon.wordpress.com/2015/01/30/the-peer-to-peer-hypothesis-and-a-new-theory-of-freewill-a-brief-overview/>? Accessed October 8, 2021.
- Beane, Silas R., Zohreh Davoudi, and Martin J. Savage. 2014. "Constraints on the Universe as a Numerical Simulation." *The European Physical Journal A* 50, 9: 148.
- Block, Ned. 1978. "Troubles with Functionalism," In *Perception and Cognition: Issues in the Foundations of Psychology*, edited by C. Wade Savage, Univ. Minn. Press, pp. 261-326.
- Block, Ned. 1981. "Psychologism and Behaviorism." *The Philosophical Review*, 90(1), 5-43.
- Bostrom, Nick. 2003. "Are We Living in a Computer Simulation?" *The Philosophical Quarterly*, 53(211), 243-255.
- Bouwsma, Oets Kolk. 1949. "Descartes' Evil Genius." *The Philosophical Review*, 58(2), 141-151.
- Campbell, Tom, Owahdi, Houman, Sauvageau, Joe. & Watkinson, David. 2017. "On Testing the Simulation Theory." *International Journal of Quantum Foundations* 3: 78-99.
- Chalmers, David J. 2005. "The Matrix as Metaphysics," in Grau, C. (ed.) *Philosophers Explore the Matrix*, Oxford: Oxford University Press.
- Chalmers, David J. 2017. "The Virtual and The Real." *Disputatio: International Journal of Philosophy*, 9(46).
- Chalmers, David J. 2018. "Structuralism as a Response to Skepticism." *The Journal of Philosophy*, 115(12), 625-660.
- Chalmers, David J. 2022. *Reality+: Virtual Worlds and the Problems of Philosophy*. Penguin UK.
- Dainton, Barry. 2002. "Innocence Lost Simulation Scenarios: Prospects and consequences." Retrieved March 31, 2021, <<http://www.simulation-argument.com>>
- Dainton, Barry. 2012. "On Singularities and Simulations." *Journal of Consciousness Studies*, 19(1-2), 42-85.
- Helton, Grace. 2021. "Epistemological Solipsism as a Route to External World Skepticism." *Philosophical Perspectives*, 35(1), 229-250.

¹⁹ For extremely helpful comments on this paper, I am indebted to: Josh Armstrong, Daniela Dover, and Chris Register.

- Helton, Grace. forthcoming. "On Being a Lonely Brain-in-a-Vat: Structuralism, Solipsism, and the Threat from External World Skepticism." *Analytic Philosophy*.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66, no. 3: 612-640.
- Lovelace, Ada. 1842. "Translator's Notes to an Article on Babbage's Analytical Engine." In *Scientific Memoirs*, Edited by R. Taylor, vol. 3., pp. 691-731.
- Mawson, Tim J. 2013. "Recent Work on the Meaning of Life and Philosophy of Religion." *Philosophy Compass*, 8(12), 1138-1146.
- Metz, Thaddeus, 2022. "The Meaning of Life," *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta & Uri Nodelman, Eds. URL = <<https://plato.stanford.edu/archives/win2022/entries/life-meaning/>>. Accessed Feb. 1, 2023.
- Mizrahi, Moti. 2017. "The Fine-Tuning Argument and the Simulation Hypothesis." *Think* 16 (47):93-102.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Moravec, Hans. 1993. "Pigs in Cyberspace." NASA. Lewis Research Center, *Vision 21: Interdisciplinary Science and Engineering in the Era of Cyberspace*.
- Rinard, Susanna. 2019. "Equal Treatment for Belief." *Philosophical Studies* 176: 1923-1950.
- Schwitzgebel, Eric & Bakker, Scott R. 2013. "Reinstalling Eden." *Nature*, 503(7477), 562-562.
- Schwitzgebel, Eric, & Moore, Alan T. 2015. "Experimental Evidence for the Existence of an External World." *Journal of the American Philosophical Association*, 1(3), 564-582.
- Schwitzgebel, Eric. 2017. "1% Skepticism." *Noûs*, 51(2), 271-290.
- Schwitzgebel, Eric. 2019. "Kant Meets Cyberpunk." *Disputatio*, 11(55), 411-435.
- Seachris, Joshua. 2019. "Meaning of Life: Contemporary Analytic Perspectives." *Internet Encyclopedia of Philosophy*. <<https://iep.utm.edu/mean-ana/>> Accessed Feb. 2, 2023.
- Summers, Micah & Arvan, Marcus. 2022. "Two New Doubts about Simulation Arguments." *Australasian Journal of Philosophy*, 100(3), 496-508.
- Turing, Alan. 1950. "Computing Machinery and Intelligence." *Mind*, 59(236), 433-460.
- Zimmerman, Michael J. 1999. "Virtual Intrinsic Value and the Principle of Organic Unities." *Philosophical and Phenomenological Research*, 653-666.

