

# How to Respond to the Simulation Skeptic

Anandi Hattiangadi

Institute for Futures Studies & Stockholm University

DRAFT

Comments and criticisms are Most Welcome!

Please do not cite or circulate without permission

## 1 Introduction

Imagine that you are looking out of a window, and notice that there is a cat stuck in a tree outside. You can see its crouched form on a high branch and you can hear it mewl as it contemplates the route back down. Indeed, you recognize it as your neighbor's cat, Margot, who has frequently been stuck in trees in the past. On the basis of your experience, your memories, and background beliefs, you form the belief that there is a cat in a tree. As it happens, your belief is true: Margot *is* stuck in a tree. It is plausible that under these conditions, you *know* that there is a cat in a tree. After all, your belief is both true and supported by strong evidence.

Now suppose that, out of the blue, a skeptic comes along and casts doubt on whether you know that there is a cat in a tree. Perhaps, the skeptic says, you are and always have been living in an artificially designed computer simulation of a world.<sup>1</sup> Perhaps the simulation you are living in is *pure*, in which case, all the beings in it are *pure sims*, including you. If you are a pure sim, your brain and body are simulated too, so you exist wholly inside the simulation; you are not merely connected to the simulation while being spatially located outside of it, like a brain in a vat, or Neo from *The Matrix* (Chalmers 2022, 30–31). Perhaps, the skeptic continues, the simulation you are in is *global*, so that the whole of the universe you inhabit is simulated, not just your local region, including Margot and the tree, the city you are in, or even the Earth (Chalmers 2022, 31). If you are in a global simulation, everything in the spacetime of your universe is simulated, from the Big Bang to the present moment, and from Earth to the far reaches of outer space. Finally, the skeptic muses, perhaps the simulation you are living in is *evidentially perfect*, in that it is evidentially indistinguishable from the world it simulates. If you are in an evidentially perfect simulation, there is no evidence available to you or any inhabitants of the simulation that would reveal that you are living in the simulation. There is no red pill, such as in *The Matrix*, that will cause the shades to fall from your eyes and see the world as it really is, outside of the simulation. There are no glitches. An evidentially perfect simulation simulates every single conscious experience and thought in the world it simulates. Thus, the totality of evidence potentially available to you if you are living in an evidentially perfect simulation is identical to the totality of evidence potentially available to you in the physical world the simulation simulates.<sup>2</sup>

---

<sup>1</sup> This is the *Simulation Hypothesis* as formulated by Chalmers (2022, 29).

<sup>2</sup> It is being implicitly assumed here that your evidence consists of perceptions and thoughts, such as your memories and beliefs. This is in line with Chalmers' characterization of evidence and the skeptical hypothesis. See section 2 for further discussion of this point.

In light of the hypothesis that you are living in a pure, global, and evidentially perfect simulation,<sup>3</sup> the skeptic goes on to argue that you can't know anything about the external world. His argument goes as follows (Chalmers 2022, 56):

THE MASTER ARGUMENT

- (1) You can't know that you're not living in a simulation.
- (2) If you can't know that you are not living in a simulation, you can't know anything about the external world.
- (3) So: You can't know anything about the external world.<sup>4</sup>

In support of the first premise, the skeptic argues that you would have exactly the same evidence that you have in the physical world that the simulation simulates, so the totality of your evidence is insufficient to support the belief that you are not living in a simulation. Since your evidence is insufficient to support the belief that you are not living in a simulation, you can't know that you are not living in a simulation. In support of the second premise, the skeptic argues that if you were living in a simulation, the external world would not exist—there would be no cats, no trees, no planet Earth—and all of your beliefs about the external world would be false. So, if you can't know that you are not living in a simulation, you can't know anything about the external world. The skeptic concludes triumphantly that you can't know anything about the external world.

As luck would have it, you have just finished reading David Chalmers' recent blockbuster, *Reality+: The Problems of Philosophy in Virtual Reality*, so you are armed with a novel and powerful response to the simulation skeptic. Following Chalmers, you reject the skeptic's claim that if you are living in a simulation, the external world does not exist, and your beliefs about the external world are false. You point out, as Chalmers persuasively does, that if you are and always have been living in a computer simulation, your beliefs about the external world are *true*. This is because, first, if you are in a simulation, your beliefs are *about* the virtual objects and properties with which you have been interacting in the simulation; and second, if you are in a simulation, the virtual objects and properties in your environment are *real*, though they are ultimately constituted by digital processes. Since your beliefs would be true even if you were living in a simulation, you do not need to know that you are not living in one to have knowledge of the external world. You conclude, triumphantly, that you *can* after all know many things about the external world.

Unfortunately, your dispute with the simulation skeptic is not yet settled. Ingenious and compelling though it may be, on closer inspection, Chalmers' response to the skeptic is inadequate. The skeptic's dialectical opening can be pursued against the background assumption of *semantic externalism*, the view that the contents of your beliefs depend on facts about your external physical and social environment (Burge 1979; Kripke 1980; Putnam 1975). As I will argue, this provides the skeptic with

---

<sup>3</sup> Though Chalmers (2022) considers many different kinds of simulation hypothesis, it is only pure, global, and evidentially perfect simulations that are relevant here. I will drop these qualifications hereafter.

<sup>4</sup> The conclusion of the skeptical argument is a bit of an overstatement. As Chalmers points out, the skeptical argument doesn't undermine some knowledge about the external world, such as that Margot is identical to Margot. Still, the skeptical argument undercuts many of the ordinary things that we think we know about the external world. I will set aside this complication here.

grounds to reject Chalmers' claim that your beliefs about the external world would be true, even if you were living in a simulation. Moreover, I argue that the difficulty cannot be resolved by appeal to a two-dimensional semantics, such as Chalmers has defended elsewhere (Chalmers 2002; 2004; 2006). However, ultimately, I would like to suggest that one of Chalmers' core insights—and his anti-skeptical stance—can be maintained. I propose a response to the skeptic that is inspired by Chalmers' own, though it only preserves some of his commitments, while abandons others.

I proceed as follows. In the next section, I present Chalmers' response to the simulation skeptic in more detail, drawing out some of his basic commitments. In section 3, I elaborate the retort the skeptic could reasonably make to Chalmers. In section 4, I argue that an appeal to a 2-dimensional semantics does not help to resolve the difficulty. In section 5, I propose a way out of these difficulties that preserves the anti-skeptical claim that you can know many things about the external world, but which is not compatible with all of Chalmers' commitments.

## 2 Chalmers' Response to the Simulation Skeptic

The key to Chalmers' response to the simulation skeptic is to reject the second premise of the MASTER ARGUMENT. However, Chalmers accepts the first premise of that argument—he accepts that you can't know that you are not living in a computer simulation. It will be instructive to take a closer look at his motivation for accepting this premise, since it reveals some of his background commitments.

Here is how the motivation goes. Suppose that you are not in fact living in a computer simulation, but in a physical world. Then, your belief that you are not living in a simulation is true. However, that is not enough for you to know that you are not living in a simulation, since knowledge does not just require true belief, it also requires justification (Chalmers 2022, 44). To qualify as genuinely knowing something, your belief must be appropriately based on evidence. And, as Chalmers argues, though “you might think you have definitive evidence that you're not' living in a simulation, ‘that's impossible, because any such evidence could be simulated.” (Chalmers 2022, 32) Indeed, he goes on, in “a good-enough simulation, the world would look and feel to you exactly as today's world looks and feels to you now. And if a simulation would look and feel the same as reality, it's hard to see how we could know we're in a simulation rather than reality.” (Chalmers 2022, 57) Thus, Chalmers concludes, you can't know that you are not living in a simulation.

The foregoing reflects two of Chalmers' philosophical commitments. The first is a commitment to *epistemological internalism*, the view that knowledge requires justification by states that are cognitively accessible to the knower, at least upon reflection (BonJour 2002, 243). Chalmers explicitly indicates sympathy to the view that knowledge requires justification (Chalmers 2022, 44). And while he never explicitly commits to the claim that justifying states must be cognitively accessible to the knower, this is suggested by Chalmers' implicit assumption that if everything in a simulation would look and feel exactly as it does in the physical world, one would have the same evidence in both cases. Since your sensory experiences—which represent how things look and feel<sup>5</sup>—are cognitively accessible to you, at least upon reflection, Chalmers

---

<sup>5</sup> I don't mean to suggest here that Chalmers restricts your evidence to your sensory experiences. Though Chalmers explicitly mentions how things would look and feel, it is implausible that he imposes such a restriction. For instance, given that he has elsewhere defended the view that

seems to be assuming that if your total evidence in the two cases is cognitively indistinguishable, then your total evidence in the two cases is identical, which amounts to assuming that your evidence is cognitively accessible to you.

The second commitment is an implicit assumption of his argument that because your evidence in a simulated world would be identical to your evidence in the physical world, you cannot know that you are not in a simulation. Here, he seems to assume that to know that you are not living in a simulation, your evidence must eliminate the possibility in which you are living in a simulation, where your evidence *E* eliminates a possibility *w* just in case you do not have *E* in *w* (Lewis 2000, 553). This suggests that Chalmers might be assuming a general principle along the following lines:

ALTERNATIVES: For all *S*, *p*: *S* knows that *p* if and only if *S*'s evidence eliminates every (relevant) possibility in which it is not the case that *p*.<sup>6</sup>

ALTERNATIVES is restricted to *relevant* possibilities, since Chalmers acknowledges that his response to the simulation skeptic may not generalize to other skeptical hypotheses, such as the Cartesian hypothesis that you are being deceived by an evil demon (Chalmers 2005). Since he argues that you can know many things about the external world, even though he concedes that you cannot rule out some skeptical hypotheses, he must be setting aside those skeptical hypotheses his response does not tackle as irrelevant.

Let us now turn to the second premise of the MASTER ARGUMENT. Though Chalmers ultimately goes on to reject this premise, he provides an intuitive justification for it, from the skeptic's perspective. By Chalmers' lights, the skeptic's key thought is that if you *were* living in a computer simulation, then the external world would not exist, and most of your ordinary beliefs about the external world would be false. Furthermore, if your evidence does not eliminate the possibility that you are living in a simulation, it equally doesn't eliminate the possibility that there are no cats or trees, and thus that your belief that there is a cat in a tree is false. Given ALTERNATIVES, if your evidence does not eliminate the possibility in which the proposition that there is a cat in a tree is false, then you can't know that there is a cat in a tree. Ditto for all of your beliefs about the external world. Thus, the skeptic concludes, if you can't know that you are not living in a simulation, you can't know anything about the external world (Chalmers 2022, 57).

Chalmers rejects this second premise of the MASTER ARGUMENT by putting pressure on the skeptic's assumption that if you are living in a simulation, most of your beliefs about the external world are false. As against this, Chalmers defends the following claim:

---

conceivability entails possibility, he would presumably allow intuitions about cases to be included in your evidence (Chalmers 2009). Though this evidence might not play a role in the evidence you have for your ordinary beliefs about the external world, it might

<sup>6</sup> ALTERNATIVES is familiar from Lewis (2000, 554). Though Lewis does not explicitly restrict the scope of 'every' to relevant possibilities, he allows context to restrict it. I have made the restriction explicit to accommodate Chalmers' restriction of the scope of his anti-skeptical argument to simulation skepticism. I am grateful to Anand Vaidya for the suggestion to explicitly restrict the scope of 'every' in this way.

TRUTH: If you are living in a simulation, most of your beliefs about the external world are true.

His argument for TRUTH has two prongs, which reflects the fact that the truth of a belief depends on both its *content*—what it is about—and *reality*—the way the world is in relevant respects. On the reality prong, Chalmers defends a view he calls *simulation realism* (Chalmers 2022, 106):

SIMULATION REALISM: The virtual objects and properties that one encounters in a simulation are just as real as the objects and properties that one encounters in a physical world.

His argument for SIMULATION REALISM relies on the assumption that a simulation that is evidentially perfect is also *structurally perfect*—it mirrors the world that it is simulating in every respect, so it constitutes a *model* that is perfectly structurally isomorphic to the world that it models, down to the last detail (Chalmers 2022, 35). A structurally perfect simulation of our world will simulate every single particle in the universe, and each one of its properties; it will simulate every single macro-physical object and property as well as the structural relations of constitution between the micro- and macro-physical levels.<sup>7</sup> If you have a glass of physical water in front of you in the physical world, there will be a structurally isomorphic glass of virtual water in front of you in a structurally perfect simulation of our world. If the water in physical reality is composed of H<sub>2</sub>O molecules, the virtual water in the simulation will be composed of structurally isomorphic virtual H<sub>2</sub>O molecules. If the H<sub>2</sub>O molecules in the physical world are composed of hydrogen and oxygen, the virtual H<sub>2</sub>O molecules in the simulation are composed of virtual hydrogen and virtual oxygen. And so forth, all the way down to the smallest subatomic particles. The difference between the physical world and its simulated counterpart lies exclusively in the kinds of objects and properties which ultimately realize the structure of reality: In a physical world, everything is fundamentally realized by physical objects and properties, while in a simulated world, everything is fundamentally realized by digital processes or bits. And this difference, Chalmers argues, is immaterial to whether the objects and properties in a simulated world are real. Whereas in the physical world you have physical hands and there are physical trees, cats, and water, in a simulated world, you have virtual hands and there are virtual trees, cats, and water. It is important to note that Chalmers is not saying that virtual objects are *identical* to their physical counterparts—a virtual cat is not identical to a physical cat. All he is saying is that a virtual cat is just as *real* as a physical cat, though it is ultimately constituted by digital processes (Chalmers 2022, chap. 6).

On the content prong, Chalmers argues that if you are living in a simulation, your concepts and beliefs are about the virtual objects and properties that you interact with in the simulation. For instance, Chalmers would say that if you are living in a simulation, your external-world concepts—‘tree’, ‘cat’, ‘water’ and so on—are about virtual trees, virtual cats, and virtual water. And if you are living in a simulation and

---

<sup>7</sup> Since I will restrict my attention to pure, global, and perfect simulations hereafter, I will drop the modifiers.

believe that there is a cat in a tree, what you believe is the proposition that there is a virtual cat in a virtual tree, since that is what your concepts ‘cat’ and ‘tree’ are about. He subscribes, that is, to the *Virtual Content Thesis*:

VIRTUAL CONTENT THESIS: If you are living in a simulation, the contents of your external-world concepts are the virtual objects and properties you encounter in the simulation, and the contents of your beliefs are propositions representing virtual objects and properties.

One way to motivate the VIRTUAL CONTENT THESIS is by appeal to the semantic externalist view that the external environment in which you find yourself plays an important role in determining the contents of your beliefs (Burge 1979; Kripke 1980; Putnam 1975). Though Chalmers does not rely on semantic externalism to make his case for the VIRTUAL CONTENT THESIS, he suggests that the two are compatible (Chalmers 2022). Semantic externalism can be characterized with reference to Putnam’s (1975) familiar Twin Earth thought experiment. Imagine a planet that is just like Earth in most respects save that the clear, colorless, odorless liquid typically found in lakes and streams and so forth is not composed of H<sub>2</sub>O, but XYZ. Twin Hypatia is the Twin Earth counterpart of Hypatia on Earth. Twin Hypatia and Hypatia are alike in what goes on in their heads:<sup>8</sup> both have a concept they express using the term ‘water’, and that figures in many of their beliefs about the clear, colorless, odorless substance typically found in lakes and streams in their local environs. The difference between the two lies in the local environs: whereas the clear, colorless, odorless liquid in Hypatia’s environment is H<sub>2</sub>O, it is XYZ in Twin Hypatia’s environment. According to the externalist, this case goes to show that while the concept Hypatia expresses with the term ‘water’ designates H<sub>2</sub>O, the concept Twin Hypatia expresses with the term ‘water’ designates XYZ. By the same token, semantic externalism implies that if you are living in a simulation, your concepts ‘cat’ and ‘tree’ refer to the virtual cats and virtual trees that you interact with in the simulation, just as Twin Hypatia’s concept ‘water’ picks out XYZ in hers. Thus, when you form the belief that there is a cat in a tree in the simulated world, your belief has as its content the proposition <there is a virtual cat in a virtual tree>.<sup>9</sup>

Finally, combining the VIRTUAL CONTENT THESIS with SIMULATION REALISM, it is easy to see how Chalmers reaches the conclusion that even if you are living in a simulation many of your beliefs about the external world are true. Consider once again the belief you would express with the sentence ‘there is a cat in a tree’. If you are living in a simulation, the content of this belief is the proposition that <there is a virtual cat in a virtual tree>. Given that the proposition that  $p$  is true if and only if  $\not{p}$ , the proposition that <there is a virtual cat in a virtual tree> is true if and only if there

---

<sup>8</sup> The distinction between semantic internalism and externalism is commonly characterized in terms of intrinsic duplicates, so internalism is the thesis that content supervenes on intrinsic features of subjects (Goldberg 2007). However, this is somewhat problematic, since Hypatia and Twin Hypatia are not intrinsic duplicates, given that Hypatia is largely composed of H<sub>2</sub>O, while Twin Hypatia is largely composed of XYZ. So, I prefer to say that Hypatia and Twin Hypatia are duplicates with respect to what is going on in their heads.

<sup>9</sup> I follow the convention of using pointy brackets to designate propositions. I remain neutral here on whether propositions are to be understood as sets of possible worlds, or as structured entities containing objects and properties as constituents.

really is a virtual cat in a virtual tree. Given SIMULATION REALISM, and given that you are in a perfect simulation, this condition is met: there really is a virtual cat in a virtual tree in the simulated world, so your belief that <there is a virtual cat in a virtual tree> is true. Ditto for most of your other beliefs about the external world. Thus, even if you are living in a simulation, most of your beliefs about the external world are true. The skeptic appears to have been fended off.<sup>10</sup>

### 3 The Skeptic's Reply

As we have seen, Chalmers' response to the skeptic relies on the claim that the contents of your beliefs vary depending on whether you are living in a virtual world or a simulation. However, that very claim leaves a dialectical opening for the skeptic. The reason is that even if it is granted that the beliefs you would have if you lived in a simulation would be true, it is nevertheless the case that the beliefs that you *actually* have, with the *contents* they actually have, are *false* if they are evaluated at the world in which you live in a simulation. Moreover, whether your beliefs in the actual world are justified depends on whether the beliefs you *actually* have are true if evaluated at a simulated world, not on whether some *other* beliefs that sim-you has are true at that world. Thus, the skeptic could argue, if you don't know that you're not living in a simulation, you don't know anything about the external world.

Allow me to elaborate. Consider the belief you would express with the sentence 'there is a cat in a tree'. Assuming that you are living in a physical world, the content of your belief is the proposition that <there is a physical cat in a physical tree>. By hypothesis, this proposition is true when evaluated at the actual world. The question is whether this proposition is true when evaluated at the world in which you are living in a simulation. In the simulated world, though there are virtual cats and virtual trees, there are neither physical cats nor physical trees. Thus, the proposition that <there is a physical cat in a physical tree>, the proposition you actually believe, is *false* when evaluated at the simulated world. Since this proposition is false at the simulated world, to know it, your evidence must eliminate the simulated world. After all, ALTERNATIVES implies that if you know  $p$ , your evidence eliminates all of the possibilities in which  $\neg p$ . Since the simulated world is a possibility in which the proposition that <there is a

---

<sup>10</sup> It is worth contrasting Chalmers' response to the simulation skeptic with Putnam's (1981) response to the skeptic who hypothesizes that you are and always have been a brain in a vat. Assuming semantic externalism, Putnam argues that the sentence 'I am a brain in a vat' is false regardless of whether or not one is a brain in a vat. Crucially, Putnam pointed out that if you were a brain in a vat, then terms like 'brain' and 'vat' would not be connected to real brains and real vats, but to neural stimulations of some kind, which Putnam characterizes as 'brains-in-the-image' and 'vats-in-the-image'. So, if you were a brain in a vat, Putnam argues, and you asserted 'I am a brain in a vat', the content of your assertion would be the proposition that <you are a brain-in-the-image in a vat-in-the-image> which, if you were a brain in a vat, would be false. So, Putnam concludes, I know that I am not a brain in a vat. There are several respects in which Chalmers' response to the skeptic differs from Putnam's. First, while Putnam assumes semantic externalism, Chalmers views this as optional. Indeed, given the objection raised in the present paper, it may be that Chalmers' response to the skeptic is incompatible with semantic externalism. Second, whereas Chalmers holds that I *don't* know that I am not living in a simulation, but I can have knowledge of the external world nonetheless, Putnam argues that I *know* that I am not a brain in a vat. This makes Chalmers' response to the skeptic more powerful than Putnam's, since it applies to all of our beliefs about the external world, whereas Putnam's is limited to the belief that I am not a brain in a vat, and other similar beliefs.

physical cat in a physical tree> is false, for you to know that proposition, your evidence needs to eliminate the simulated world. But your evidence does not eliminate the simulated world, since you would have all the same evidence in the simulated world as you do in the actual world. Since your evidence does not eliminate the simulated world, you cannot know the proposition that <there is a physical cat in a physical tree>. Ditto for all of your beliefs about the external world. So, if you can't know that you are not living in a simulation, you can't know anything about the external world.

The skeptic puts pressure on the epistemic status of the beliefs you actually have, by pointing out that their *contents* would be false if you lived in a simulation. Such a skeptic is not moved by Chalmers' argument that your *beliefs* would be true if you lived in a simulation, because he points out that if you lived in a simulation, you would have *different* beliefs from those you actually have, with different contents than those you actually believe. And what matters for whether you can know the proposition that <there is a physical cat in a physical tree> is whether your evidence eliminates possibilities in which *that* proposition is false, not some other proposition you would believe in some other possible situation. What the skeptic questions is whether the beliefs you actually have, with the contents they actually have, amount to knowledge. He argues that they do not, because the propositions regarding the external world that you believe are false when evaluated at a simulated world, and since your evidence does not eliminate simulated possibilities, your actual beliefs, with the contents that they actually have, are not justified.<sup>11</sup>

Notice, moreover, that though your beliefs about the world would be true if you were living in a simulation, because they would have a different content than they actually have, many of them would not amount to knowledge. After all, sim-you believes that <there is a virtual cat in a virtual tree>, a proposition that would be false if evaluated at the evidentially indistinguishable physical world,  $w_0$ . By ALTERNATIVES, for sim-you to know that <there is a virtual cat in a virtual tree>, their evidence would have to eliminate  $w_0$ . But since  $w_0$  is evidentially indistinguishable from  $w_1$ , it is not eliminable. It turns out that sim-you is in the same unfortunate epistemic situation as you.

#### 4 2-Dimensionalism to the Rescue?

The foregoing difficulties with Chalmers' response to the skeptic were spelled out against the background assumption of semantic externalism, according to which our beliefs are associated with just one kind of content, *wide* content, which depends on the external environment in the manner just described. However, Chalmers himself is not a proponent of semantic externalism, but famously defends a two-dimensional (2D) semantics, which postulates two kinds of content: *narrow*, and *wide* (Chalmers 2002; 2004; 2006). Unlike wide contents, narrow contents, informally glossed as descriptive meanings, are thought to depend on the goings on in the heads of agents. Perhaps the difficulties outlined in the foregoing section can be resolved by appeal to 2D semantics.

Formally, 2D contents are modeled by pairs of intensions: a *primary* intension, which models narrow content, and a *secondary* intension, which models wide content. The primary intension of a concept or thought is a function from scenarios to

---

<sup>11</sup> Note that this also suggests that Chalmers' line of response to the skeptic may not be compatible with semantic externalism.



extensions or truth values, while the secondary intension of a concept or thought is a function from ordinary possible worlds to extensions or truth values. A scenario can be viewed as a *centered* possible world, that is, a triple consisting of a possible world, together with an agent, a time, and place marking the ‘center’ of the world.<sup>12</sup> To work out the primary intension of a concept, you need to consider what the concept would pick out at a given scenario given its descriptive meaning, if that scenario is imagined to represent the actual world. To illustrate, in the case of the concept ‘water’, the narrow content is a descriptive meaning that can be glossed as ‘stuff that occupies the water-role’, where the ‘water-role’ is an abbreviation of a descriptive meaning such as ‘clear, colorless, odorless liquid typically found in lakes and streams around here (etc.).’ The primary intension of ‘water’ is a function from the set of scenarios,  $V$ , to whatever it is that occupies the water-role at each  $v_i \in V$ , if anything does, while the descriptive meaning of ‘water’ is held fixed. For instance, when evaluated at  $v_0$ , the actual world centered on Hypatia, the primary intension of ‘water’ is H<sub>2</sub>O, since it is H<sub>2</sub>O that occupies the water-role at  $v_0$ . When evaluated at  $v_1$ , the Twin Earth world centered on Twin Hypatia, the primary intension of ‘water’ is XYZ, since it is XYZ that occupies the water role at  $v_1$ . Insofar as Hypatia and Twin-Hypatia attach the same descriptive meaning to ‘water’, the primary intension will be shared: both concepts are associated with a primary intension that picks out H<sub>2</sub>O at  $v_0$  and XYZ at  $v_1$ .

The secondary intension of ‘water’ is a function from the set of possible worlds,  $W$ , to extensions, which holds both the descriptive meaning and the scenario describing the actual world fixed. To work out the secondary intension of a term, you hold fixed your assumptions about what the term picks out in the actual world, given its descriptive meaning, and then evaluate its extension at every possible world  $w_i \in W$ . For instance, if we put ourselves in Hypatia’s shoes, and hold  $v_0$  fixed as the actual world, the secondary intension of ‘water’ will be H<sub>2</sub>O at all  $w_i \in W$ . If we put ourselves in Twin Hypatia’s shoes, and hold  $v_1$  fixed as the actual world, the secondary intension of ‘water’ will be XYZ at all  $w_i \in W$ . Thus, while the primary intension of ‘water’ is the same for Hypatia and Twin Hypatia, the secondary intensions come apart.

How might Chalmers appeal to the 2D framework to avoid the foregoing difficulties with his response to the skeptic? One plausible line of argument starts with the thought that since narrow contents depend on what goes on in one’s head, they are cognitively accessible to one, so if one’s justification must be cognitively accessible, as the epistemic internalist supposes, it too must pertain to the narrow contents of one’s beliefs. Moreover, Chalmers might argue, the narrow contents of your beliefs about the external world are not only shared between you and your simulated counterparts, but their primary intensions are true at all simulated scenarios, so your evidence need not eliminate simulated scenarios for you to have knowledge of the external world. To elaborate, consider the narrow content of the belief you would express with the sentence ‘there is a cat in a tree’. Informally, we can represent this as the content *<there is a cat in a tree>*;<sup>13</sup> formally, we can represent it as a function from scenarios to truth values, where it is true at  $v_i$  iff the descriptive meaning, *<there is a cat*

---

<sup>12</sup> There are other ways of characterizing scenarios (Chalmers 2006), but I will work with the centered worlds characterization here.

<sup>13</sup> I use pointy brackets and italics to represent narrow contents.

*in a tree*>, is satisfied at the world that corresponds to  $v_1$ . According to Chalmers, while physical cats and physical trees satisfy the descriptive meanings of our concepts ‘cat’ and ‘tree’ at our world, virtual cats and virtual trees satisfy the descriptive meanings of these concepts at simulated worlds. As a consequence, the narrow content *<there is a cat in a tree>* will be true at the scenario  $v_2$ , the actual world centered on you as you look out the window, and it will also be true at the simulated world,  $v_3$ , centered on sim-you looking out the window. Since *<there is a cat in a tree>* is true at  $v_3$ , your evidence does not need to eliminate it. By ALTERNATIVES, your evidence is sufficient to justify your belief that *<there is a cat in a tree>*, and thus you can know it. So, the skeptic’s challenge can be met: not only are the beliefs about the external world true, but they are justified, and you can know them.

Promising though this line of response may seem, I will argue that it ultimately fails. In particular, I shall question the claim that the narrow contents of our external world concepts are about virtual objects, on the assumption if we are living in a simulation. That is, I argue that if the 2D semantic framework is assumed, then if you are living in a simulation, most of your ordinary beliefs about the external world are *not* true. So, given ALTERNATIVES, your evidence needs to eliminate these simulated worlds, if you are to have knowledge of the external world. Since your evidence does not eliminate these worlds, the 2D line of response to the skeptic falls flat.

There are two platforms of the 2D framework that ultimately conspire to undermine the claim that our concepts and beliefs are about virtual objects and properties in a virtual world. The first platform is that wide contents are partly determined by narrow contents: the wide content of a concept or a thought is that which satisfies its narrow content. For instance, the wide content of ‘water’ at a world is whatever it is that satisfies the descriptive meaning of ‘water’ at that world.

The second platform of the 2D framework worth mentioning concerns the meta-semantics of narrow content, what it is that determines the narrow contents of a thought. Chalmers (2011) argues that the narrow content of a thought at a scenario is determined by what we would judge its contents to be, given ideal information about the scenario, and under suitably idealized conditions of rational reflection. For instance, the fact that our concept ‘water’ has the narrow content that it does is determined by our judgments about what its content would be at a given scenario, given idealized information about that scenario, and under conditions of ideal rational reflection. For example, if we suppose that  $v_1$  is actual, and we are told that in this scenario, the stuff that occupies the water-role<sup>14</sup> is XYZ, we typically judge that our concept ‘water’ picks out XYZ at  $v_2$ . In contrast, if we suppose that  $v_0$  is actual, and we are told that in this scenario, H<sub>2</sub>O occupies the water-role, we typically judge that ‘water’ picks out H<sub>2</sub>O in that scenario (Chalmers 2011, 396). It is worth emphasizing that the judgments that determine the narrow contents of our concepts and beliefs are based on idealized information, not necessarily available to the agents of the scenario. For instance, though Twin Hypatia does not know the microphysical constitution of water at  $v_1$ , we know that it is XYZ, and it is this information that forms the basis of the judgment that ‘water’ picks out XYZ at  $v_1$ , when we imagine it to be actual.

---

<sup>14</sup> Of course, to avoid triviality, the description of the scenario cannot be stated using the term ‘water’, but using terms such as ‘clear, colorless, odorless liquid’ which give the term ‘water-role’ was said to abbreviate.

Taking those two platforms of the 2D framework together, if we want to know if ‘cat’ picks out virtual cats on the hypothesis that we are actually living in a simulation, we need to know whether virtual cats satisfy the descriptive meaning associated with the narrow content of ‘cat’; and in order to know whether virtual cats satisfy the descriptive meaning associated with the narrow content of ‘cat’, we need to know whether we would judge ‘cat’ to pick out virtual cats given sufficient information about a simulation scenario, on the supposition that the simulation scenario is actual. I will argue in the next two sections that we should not judge ‘cat’ to pick out virtual cats under these conditions, and more generally, that we should not judge our external-world concepts and beliefs to be about virtual objects and properties, on the supposition that we are actually living in a simulation.

#### 4.1 *Quasi-Perfect Simulations*

As previously noted, there are two aspects to the perfection of a simulation. A simulation can be evidentially perfect in the sense that there are no red pills, no glitches, no evidence that one could possibly obtain that would reveal that one is living in a simulation. A simulation can also be structurally perfect in the sense that it constitutes a perfect mirror or model of the reality that it simulates, down to the last detail.

Clearly, these two kinds of perfection can come apart. At any rate, the fact that a simulation is evidentially perfect does not logically or even a priori entail that it is structurally perfect. It is conceivable for a simulation to give rise to exactly the same sensory experiences and thoughts that you would have if you lived in the world it simulates, without being a perfect model of the world it simulates. It is conceivable for a simulation to give rise to the sensory experiences you have when you see a cat stuck in a tree without the computational processes that give rise to those experiences modelling cats or trees. For example, it is conceivable that you are in a simulation in which the computational processes that give rise to your visual experience of the cat are not co-located with the computational processes that give rise to your auditory experience of the cat. It is conceivable that the computational processes that give rise to your visual experience of the tree are not themselves structured in the way the tree is structured—perhaps the computational processes that give rise to your experience of the whole tree are not composed of parts that give rise to your experience of its leaves or trunk, and are not in turn composed of parts that constitute the cells of the tree, which would give rise to suitable experiences when viewed under a virtual microscope, and so forth. It doesn’t seem as if any of these simulations is inconceivable or a priori impossible. Thus, the possibility that a simulation is evidentially perfect but not structurally perfect cannot be ruled out a priori.

Let’s call simulations that are evidentially perfect but not structurally perfect *quasi-perfect simulations*. Since quasi-perfect simulations are evidentially perfect, they serve the skeptic’s purposes just as well as perfect simulations. You cannot rule out the possibility that you are living in a quasi-perfect simulation, since the total evidence available to you if you lived in one is identical to the total evidence available to you if you live in the physical reality it simulates. However, quasi-perfect simulations are problematic for the 2D response to the skeptic, since—as I will argue—the narrow contents of our external world concepts and beliefs do not pick out virtual objects and properties, on the assumption that we are living in a quasi-perfect simulation.

Consider once again the belief you would express with the sentence ‘there is a cat in a tree’. Now suppose that the scenario you actually find yourself in is a quasi-perfect simulation, QUASI. In QUASI, of your thoughts and experiences are caused by computational processes, but these computational processes do not perfectly model the objects and properties in the physical world. Rather, in QUASI, the computational processes that give rise to your experiences are being run on different servers that are in different geographical locations in the world outside of the simulation. Your visual experience of the cat’s smile is caused by computational processes taking place on a server in Cheshire, while your visual experiences of the rest of the cat are caused by computational processes taking place on a server outside Reykjavik, and your auditory experiences of the mewling are caused by computational processes taking place on a server outside Tokyo. In a scenario such as this, it is difficult to find a virtual object for ‘cat’ to pick out, since we tend not to think of cats as spatially distributed in this way. Rather, we tend to think of cats as organisms with a certain degree of functional cohesion, at the very least involving a single concrete individual, occupying a single, roughly continuous, spatiotemporal region, that both looks like a cat and sounds like a cat. Thus, it seems that if this quasi-perfect simulation is assumed to be actual, our concept ‘cat’ does not pick out anything at all, but is empty. Assuming that the structural deviance is significant and comprehensive, the same will of course hold for ‘tree’, ‘in’, and so forth. So, if you are living in a quasi-perfect simulation, and you form the belief with the narrow content *<there is a cat in a tree>*, your belief is not true, because there is nothing in the world in which you find yourself that satisfies it.

The conceivability of quasi-perfect simulations thus places some limitations on the 2D response to the simulation skeptic, since if a 2D semantics is assumed, our beliefs about the external world would not be true were we to be living in a quasi-perfect simulation. Now, Chalmers does acknowledge that his response has limitations. For instance, he acknowledges that his response does not address Descartes’ evil demon scenario, since all of your thoughts and experiences in that scenario might be caused by the mind of the evil demon without the evil demon’s mind structurally simulating the external world. And we set aside these skeptical scenarios as irrelevant. However, Chalmers does seem to suggest that his response to the skeptic will succeed in responding to all skeptical hypotheses involving *simulation* scenarios. If it turns out that his response doesn’t extend to quasi-perfect simulations, this imposes further limitations on his response beyond those he initially anticipated.

Moreover, as I will argue in the next section, even in perfect simulations, the narrow contents of our external-world concepts and beliefs are not satisfied by virtual objects and properties.

#### **4.2 Perfect Simulations**

Suppose that there is a physical cosmos that is just the way we take our physical cosmos to be. Suppose further that at a certain time in the history of this physical cosmos, Maya decides to run a perfect simulation of her own world, and she runs that simulation on an iPhone that she carries around in her pocket. She decides to speed up the time in the simulation so that at a certain point during Maya’s lifetime, sim-Maya is born and lives in the simulation that Maya carries around in her pocket. We can call this whole scenario, including the physical cosmos and the perfect simulation nested in it NESTED.

Now consider the narrow contents of our external-world concepts and beliefs at NESTED, on the supposition that it is actual. Consider, for instance, the concept ‘cat’. When we view NESTED from Maya’s perspective, knowing that she is a physical being living in a physical world, we should have no qualms judging that the narrow content of ‘cat’ picks out the physical cats in Maya’s environment. What about the virtual cats in sim-Maya’s environment? If Chalmers is right about the contents of our concepts and beliefs in simulated worlds, then we should be equally prepared to judge that ‘cat’ picks out the virtual cats in sim-Maya’s environment, if NESTED is considered to be actual. However, this is implausible. After all, we have an *alternative* concept available, the concept ‘virtual cat’, which has a distinct narrow content from the concept ‘cat’, and which has descriptive conditions that are satisfied by the virtual cats in sim-Maya’s environment. Given that we have both the concepts ‘cat’ and ‘virtual cat’, there is no reason why we should be rationally obligated to view the narrow content of ‘cat’ as picking out virtual cats as well as physical cats in NESTED, considered as actual. That is, we are at least rationally permitted to view the narrow content of ‘cat’ as picking out only the physical cats in this scenario, while the narrow content of ‘virtual cat’ picks out the virtual cats in sim-Maya’s simulated environment. Indeed, it is plausible that we are rationally *required* to make this distinction, given that it is ready to hand, and given that ‘virtual cat’ is the concept that more accurately applies to virtual cats.

In response to this point, it might be objected that though ‘cat’ and ‘virtual cat’ have different narrow contents, their application conditions overlap: virtual cats satisfy the application conditions of both ‘cat’ and ‘virtual cat’, so the fact that both concepts are available does not on its own show that the narrow content of ‘cat’ does not pick out the virtual cats in sim-Maya’s environment. Similarly, we have the concept ‘black cat’, which is distinct from the concept ‘cat’, and yet black cats satisfy the narrow content of ‘cat’ just as much as other cats do. However, this analogy breaks down, because the concept ‘virtual cat’ is not the concept of a *kind* of cat, as the concept ‘black cat’ is. For instance, while there may be a conceptual entailment from ‘Margot is a black cat’ to ‘Margot is a cat’, there is no conceptual entailment from ‘Margot is a virtual cat’ to ‘Margot is a cat’, since a virtual cat is not a sub-class of the class of cats. Indeed, it is arguable that ‘Margot is a virtual cat’ conceptually entails ‘Margot is *not* a cat’, since we tend to think of virtual cats as a different kind of beast entirely.

It might be objected that the foregoing argument begs the question against SIMULATION REALISM, since it denies that virtual cats are real cats. However, all the simulation realist is committed to is that the objects and properties one encounters in a simulation are just as real as the objects and properties one encounters in the physical world. SIMULATION REALISM does not entail that the objects and properties one encounters in a simulation are *identical* to the objects and properties one encounters in the real world. Even though we have assumed that the virtual cats in NESTED are structurally isomorphic to the physical cats in that scenario, that is not sufficient for their identification, because we know that in NESTED, the physical cats are constituted by physical processes, while the virtual cats on Maya’s i-phone are constituted by digital ones.

Or perhaps it will be objected that I have been chauvinistic in taking Maya’s point of view, and haven’t paid sufficient attention to sim-Maya’s point of view. Since sim-Maya is interacting with virtual objects and properties, perhaps the narrow content

of *her* concept ‘cat’ picks out virtual cats, even if our concept does not. Similarly, if we were to consider a universe that contains both a planet Earth and a planet Twin Earth, we would say that the narrow content of ‘water’ picks out H<sub>2</sub>O on Earth and XYZ on Twin Earth. However, the two cases are not parallel. One reason is that we do not have an alternative concept ‘Twin water’, which we take to be distinct from the concept ‘water’ in its narrow content. Moreover, though sim-Maya does not know that she is living in a simulation, and thus does not know that the objects and properties she interacts with are virtual, she is not restricted to relying on the information she has in the simulation when she is evaluating the narrow contents of her concepts. Rather, when she evaluates the narrow contents of her concepts, she does so under *ideal* conditions, in particular, given *full* knowledge about NESTED. Thus, assuming that sim-Maya shares the concepts ‘cat’ and ‘virtual cat’ with us, it is plausible that if *she* were presented with a full description of NESTED, she too would be rationally required to judge that her term ‘cat’ picks out the physical cats in the world outside the simulation and not the virtual cats in the simulation, reserving the concept ‘virtual cat’ for this purpose. After all sim-Maya cannot know that she is living in the simulation described in NESTED, since her evidence, which is the same as Maya’s, does not eliminate NESTED. Thus, even the narrow contents of sim-Maya’s external-world concepts do not pick out the virtual objects and properties in her environment. Against the background of the 2D semantic framework, it follows that her external-world concepts such as ‘cat’, ‘tree’, ‘water’, and so forth are empty, and thus that many of her beliefs about the external world are not true.

Moreover, our intuitions in relation to the narrow contents of our terms when evaluated at NESTED are plausibly replicated in *all* skeptical scenarios involving perfect simulations. This is because, it is plausible that for a simulation scenario to be a skeptical scenario, there must be a world outside of the simulation, in which the simulation is run. This is, as it were, what makes a simulation scenario *skeptical*. Intuitively, in scenarios involving some kind of contrast between the world as it appears to us to be and the world as it in fact is, we are *deceived*, and only these kinds of scenarios can really threaten our knowledge. In contrast, scenarios in which the whole world is a simulation—in which there is no world outside of the simulation at all—are not really skeptical scenarios, since we are not deceived in such worlds. Rather, scenarios in which there is no world outside the simulation really are non-skeptical scenarios that differ metaphysically from our world (or from how we assume our world to be). For instance, consider DIGITAL COSMOS, a scenario in which the whole cosmos is fundamentally composed of bits. DIGITAL COSMOS is not a skeptical scenario, but a scenario in which the fundamental constituents of reality are non-physical bits. Indeed, DIGITAL COSMOS is not even strictly speaking a *simulation* scenario, since the digitally constituted cosmos is not artificially designed. So, though it may well be that we would judge the narrow content of ‘cat’ to pick out virtual cats in DIGITAL COSMOS, this takes us no further in response to the skeptic, since DIGITAL COSMOS is not a skeptical scenario. This suggests that in all skeptical scenarios, the narrow contents of our concepts of physical objects and virtual objects will come apart, and thus that our ordinary concepts and beliefs about the external world will not be true.

## 5 Another Response?

I would like to close by suggesting a way that to preserve the anti-skeptical conclusion Chalmers reaches, while avoiding the foregoing difficulties. The proposal involves accepting Chalmers' VIRTUAL CONTENT THESIS, according to which your external world concepts and beliefs are about virtual objects in simulated worlds. However, it departs from Chalmers' response to the skeptic by assuming semantic externalism in place of the 2D semantic approach, while rejecting both epistemic internalism and ALTERNATIVES in favor of an epistemic externalist account of justification. In particular, it seems to me that the anti-skeptical thrust of Chalmers' approach can be preserved if we replace ALTERNATIVES with an account of justification in terms of *safety* (Williamson 2000, 147):

SAFETY: If one knows, one could not easily have been mistaken in a relevantly similar case.

Given semantic externalism and the VIRTUAL CONTENT THESIS, SAFETY offers a powerful response to the simulation skeptic, since it implies that many of one's beliefs about the external world are safe, and thus amount to knowledge. Allow me to elaborate. Suppose that in the actual world you believe the proposition that <there is a physical cat in a physical tree>, and that proposition is true. Now consider the world,  $w_3$  once again, which is a perfect simulation of the physical world. Given the VIRTUAL CONTENT THESIS, in  $w_3$ , 'cat' picks out virtual cats, and 'tree' picks out virtual trees. So, if you were in  $w_3$ , under circumstances relevantly similar to those in which you find yourself in the actual world, you would believe the proposition that <there is a virtual cat in a virtual tree>. Since  $w_3$  is a perfect simulation of the actual world, the proposition that <there is a virtual cat in a virtual tree>, which you believe at  $w_3$ , is true at  $w_3$ . So, you would not be mistaken at  $w_3$ . It follows that your belief that <there is a physical cat in a physical tree> is safe—at least if we focus our attention on perfect simulations. And by SAFETY, it follows that you know that <there is a physical cat in a physical tree>. Ditto for many similar beliefs about the external world.

Perhaps this response to the skeptic can even deal with quasi-perfect simulations. Assuming semantic externalism, in a quasi-perfect simulation, what determines the contents of your concepts will not be their narrow contents, but the relations to the external environment. If that is the case, then even in a quasi-perfect simulation, concepts such as 'cat' and 'tree' will arguably refer to the strange, gerrymandered, quasi entities that are causally connected to your concepts, and the content of the belief you would express with the sentence 'there is a cat in a tree' would be the proposition that <there is a virtual quasi-cat in a virtual quasi-tree>, a proposition that arguably is true in QUASI. Indeed, this response may even extend to Cartesian skeptical scenarios in which you are deceived by an evil demon. Perhaps, in such a world, one's concepts of external objects refer to ideas in the mind of the demon, since these are the entities that are causally connected to one's concepts, and perhaps they qualify as just as real as physical objects, for all the same reasons that simulated objects qualify as real. Perhaps, if that is so, then one's beliefs about the external world would still be true of the goings on in the mind of the demon.

Now, it might be objected that the extension to quasi-perfect and Cartesian skeptical scenarios places implausibly weak constraints on reference. Perhaps our external world concepts are empty at those worlds because even though our thoughts

have causes in those worlds, their causes are not cohesive enough determine reference. I myself am sympathetic to this worry. However, this only shows that the present response has the same limitations as Chalmers' own. And if we set those cases aside, the present proposal appears not to face the difficulties that Chalmers' original response to the simulation skeptic faces.

## 6 Works Cited

- BonJour, Laurence. 2002. 'Internalism and Externalism'. In *The Oxford Handbook of Epistemology*, edited by Paul K. Moser, 234--264. Oxford University Press.
- Burge, Tyler. 1979. 'Individualism and the Mental'. *Midwest Studies in Philosophy* 4: 73--121.
- Chalmers, David J. 2002. 'On Sense and Intension'. *Philosophical Perspectives* 16: 135--82.
- . 2004. 'Epistemic Two-Dimensional Semantics'. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 118 (1/2): 153--226.
- . 2005. 'The Matrix as Metaphysics'. In *Philosophers Explore the Matrix*, edited by Christopher Grau, 132. Oxford University Press.
- . 2006. 'The Foundations of Two-Dimensional Semantics'. In *Two-Dimensional Semantics*, edited by Manuel Garcia-Carpintero and Josep Macia, 55--140. Oxford: Oxford University Press.
- . 2009. 'The Two-Dimensional Argument Against Materialism'. In *Oxford Handbook to the Philosophy of Mind*, edited by Brian P. McLaughlin and Sven Walter. Oxford University Press.
- . 2011. 'Revisability and Conceptual Change in "Two Dogmas of Empiricism"'. *The Journal of Philosophy* 108 (8): 387--415.
- . 2022. *Reality+*. W.W. Norton & Company.
- Goldberg, Sanford. 2007. *Internalism and Externalism in Semantics and Epistemology*. Oxford University Press.
- Kripke, Saul A. 1980. *Naming and Necessity*. Harvard University Press.
- Lewis, David. 2000. 'Elusive Knowledge'. In *Knowledge: Readings in Contemporary Epistemology*, edited by Sven Bernecker and Fred I. Dretske. Oxford University Press.
- Putnam, Hilary. 1975. 'The Meaning of "Meaning"'. *Philosophical Papers* 2.
- . 1981. *Reason, Truth and History*. Cambridge University Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.