

[Home](#)

A review of David Chalmers' essay, "The Matrix as Metaphysics"

Reviewed by Peter B. Lloyd

David Chalmers' essay, "The Matrix as Metaphysics", is available online [here](#), as part of the philosophy section on the Warner Brothers' official Matrix web site. It is also visible on David Chalmers' [home page](#).

1. [Opening gambit: the envatted person is wrong, very wrong](#)
2. [Digression: How to tell whether you are in the Matrix](#)
3. [Chalmers claims to reject Berkeley](#)
4. [Envatting reconsidered](#)
5. [The implicit simpleness of elementary physical entities](#)
6. [Semantic chicanery](#)
7. [Confusion over real and virtual referants](#)
8. [How detailed must the simulation be?](#)
9. [Just-in-time reality generation](#)
10. [Simulations and Berkeley](#)
11. [Conclusion](#)

1. Opening gambit: the envatted person is wrong, very wrong

Chalmers' essay starts off in a worrying fashion, but then takes a plot turn (if we may speak of the 'plot' of an essay) worthy of *The Matrix*.

Given the sophistication of Chalmers' other writings on the philosophy of consciousness, it is surprising to see him, at first, seeming to fall into the trap of naive realism in the first section of this essay. After introducing the well-known thought-experiment of the brain-in-a-vat, he writes, "It has all sorts of false beliefs about the world.". (By "it", he means the brain in the vat, which Chalmers is identifying with the person whose brain it is.) This comment is immediately puzzling. The whole point of the brain-in-a-vat experiment is that the person whose brain it is, has experiences just the same as a person whose brain is embodied in a human body. In fact, Chalmers himself has already acknowledged that this person's experiences are normal, for Chalmers writes, "From the brain's point of view, things seem very much as they seem to you and me." Given that that is so, it seems very implausible to me that the person could possibly have any beliefs that would be rendered false simply by the brain's being in vat rather than a body.

So, what does Chalmers say these beliefs are, that are falsified by the brain's being in a vat?

The first example he gives is: "It believes that it has a body, but it has no body". (Again, 'it' is the person whose brain is in the vat.) But what exactly would this person mean if she said, "I believe I have a body"? What, precisely, does she believe? Well, to begin with, she believes that if she reaches down with her hand and touches her navel, she will get a kinaesthetic feeling of her hand lowering, following by tactile sensations in her fingers and her navel, and if she gazes at her navel, then she will have a visual sensation of her finger resting on her navel. She will also mean the set

of all such observations and interactions that she could carry out to obtain sensory impressions from her body. Now, by hypothesis, all of those observations are just the same whether her brain is in a vat or in a body. So, that part of her belief is true, Therefore, the most that Chalmers could claim is that her belief that she has a body is only partly true, not wholly true. But what other part of the person's belief in her body is there? Well, there is the belief that her pattern of observations of her body will continue in the same way for her whole life.

Now, here we run up against what is arguably an omission in the original statement of the thought-experiment. For, we *could* have either of the following two variants of the thought-experiment: (a) the person whose brain is in a vat remains plugged into the virtual reality throughout her life, from cradle to grave; or (b) at some point in time, the simulation is interrupted (like Neo being unplugged from the Matrix) and her sensory inputs are switched to some other source, which reveals to her that her brain is in a vat, not in a body. I am going to assume that (a) is the intended experiment, because that seems too be what Chalmers has in mind. Indeed, later on he says "I will ignore the complication that people sometimes travel back and forth between the Matrix and the external world". Therefore, we find that there is a further part of the person's belief that is true. Namely, that her pattern of observations of her body will continue throughout her life. (Yes, I acknowledge that this is slightly contrived restriction, but the whole point of this thought-experiment is to focus on this kind of situation.)

Is there any other part of the belief? Well, people who have been exposed to the philosophical doctrine of physicalism may be tempted to say that her belief that she has a body will also include a belief in the physical substance of her body. But is that belief *actually* held by the person whose brain is in the vat? Or is it only an artefact concocted by philosophers and mistakenly attributed to people? How can we tell? Well, in order for the person genuinely to have this belief, she would need to have the concept of the physical substance of her body as something that exists over and above all of her conscious perceptions of her body. How can such a concept be defined? Certainly not in terms of perceptions (for everything perceptual is the same for envatted and embodied persons), but only in terms of fundamental physical quantities such as mass and space. And how are those fundamental quantities defined? They are not. They are the undefined building blocks of physical theory. So, it turns out that the physical body is an abstraction that is denoted only within a linguistic system that is closed with respect to semantic reference. By 'closed with respect to semantic reference', I mean that everything in the discourse of physics cashes out in terms of fundamental physical quantities, which themselves are undefined. So, the 'physical body' turns out to be something that we cannot coherently believe to be real - because its mode of definition renders it necessarily abstract.

Let me put it another way. Because all of physics is defined in terms of the abstract and undefined fundamentals (mass etc), physics is just the same for envatted people as for people outside the vat. It is like mathematics. The laws of arithmetic are precisely the same for people in the vat, because the laws are abstract. Both mathematics and physics can be applied to observable situations, of course. If I have three pens on my desk, and two people take one each, then I am left with one pen: this works just as well in the Matrix as it does in the outer world. Likewise Newton's laws of motion apply equally inside the Matrix. In both cases, an abstract model is applied to sensory observations.

So, when the envatted person says, "I believe I have a body", her belief is exhaustively contained in her beliefs about her life-experiences. Her beliefs are not about a supposed 'physical substance'. Even if she applied an abstract physical model to her body, any conclusions that she draws from that model have to be cashed out in sensory observations. So, any predictions that she makes with her physical model will

be true in the Matrix as long as the Matrix correctly supplies the requisite sensory impressions. Therefore, her beliefs are not, in fact, falsified by her being envatted.

So, in the first section of the essay, it looks as if Chalmers has got it wrong. But things are not what they seem in Chalmers' essay! Before we go on to his second section, though, let us look at another point that Chalmers makes.

2. Digression: How to tell whether you are in the Matrix

Chalmers writes: "From the inside, there is no way to tell for sure that I am not in the situation of the brain in a vat. So it seems that there is no way to know for sure that I am not in a matrix." There is, however, a loophole. For, the way he has defined envatment is pretty much in line with it is set up in *The Matrix*. Namely, the computer controls all, and only, the sensory inputs. Now, a person who is in such a situation can easily discover it by carrying out surgical interventions in her own brain. For example, she could cut the corpus callosum and observe whether her visual field is split down the middle. In this split-brain experiment, each hemisphere (assuming that each hemisphere retains conscious awareness) will find that one half of the visual field has vanished. The 'person' in the left hemisphere will find that right-hand side of each eye's visual field no longer exists; and vice versa for the 'person' in the right hemisphere. Now, under Chalmers' assumptions (and indeed the Wachowski brothers' assumptions), the envatting apparatus cannot reproduce this effect because it controls the data flow in the optic nerve, before it is switched between hemispheres.

There are many other, more difficult but more dramatic, surgical interventions. The person could destroy the brain tissue that enables facial recognition, thereby inducing visual prosopagnasia: the person can still see everything but just cannot recognise anyone's face.

This is ultimately a minor point, because Chalmers could respond by inventing a further thought-experiment in which the envatted brain is wired up not only for sensory input and motor output but also for stimulating or inhibiting every internal piece of brain tissue.

We can iterate this further. For, the envatted person might ambitiously rewire her own brain, creating synaesthetic results: she might rewire her optic nerve into her olfactory cortex so as to smell her optic inputs. Well, to accommodate this, Chalmers could invoke a third thought-experiment in which the envatting apparatus has microsurgical mechanisms that will rewire the actual brain in the same way as the virtual brain is rewired. Then, the envatted person would get all the synaesthetic effects as expected.

Why stop there? The envatted person could surgical wire her brain into other people's brains. Heck, she could wire her brain into *everybody* else's brain - billions of brains wired together in the virtual world. (I am supposing that the Agents do not stop her from doing this.) And then wire in every animal, too. To prevent her from discovering that she is envatted, Chalmers must suppose that every person and animal in the virtual world is also envatted and that when the envatted person wires her virtual brain into all the other brains, the physical brain is likewise wired into the other brain.

This is taking us somewhat away from the main philosophical point, but I thought it was worth noting that the brain-in-vat experiment that Chalmers uses would be somewhat simplistic if we really want to stop the envatted people from finding the truth.

3. Chalmers claims to reject Berkeley

So, back to Chalmers' essay. In his second section, Chalmers alludes to George Berkeley's argument, which is essentially the same as the one I have given above (in section 1). But, for some reason, he rejects it: "I don't find the view that appearance is reality plausible, so I don't endorse Berkeley's reasoning."

There are two important comments to make here. First, it is extraordinary that Chalmers considers it appropriate to reject Berkeley's internally consistent solution of the mind-body problem with a comment as inconsequential as "I don't find the view ... plausible". Berkeley is a major philosopher whose theories have been scrutinised in academia for generations and he is still on the syllabus. Even those who disagree with him generally regard his arguments as pretty hard to refute. He also has the distinction of having put forward an internally consistent solution to the mind-body problem. Very few, if any, other philosophers have done this. Therefore, even if Chalmers rejects Berkeley, there is an onus on him to give some reason. On what grounds does he find Berkeley implausible?

Second, he has to some extent set up a straw man in describing Berkeley's philosophy as claiming that "appearance is reality". This is an ambivalent statement, which most people will read as meaning that the world as it appears to us is the whole of reality. This is more of a Humean phenomenalism than Berkeleian idealism. Berkeley explicitly maintained that the natural world that we observe is driven by an external, non-human mind, which he called 'God'.

As we shall see, the other odd thing above Chalmers' superficial dismissal of Berkeley is that Chalmers then goes on to defend a position that is remarkably close to Berkeley's.

4. Envatting reconsidered

Chalmers then reconsiders the brain-in-a-vat experiment. He presents his thesis that our everyday claims about the physical world are true, even if the physical world is ultimately a virtual reality generated by some external intelligence. His grounds for saying this are that the meaning of any statement we make about the physical world is limited to physical facts, and does not extend to metaphysical facts about the underlying reality. This is, in fact, a special case of the semantic argument that Berkeley put forward. I can only guess that it's Berkeley's bad press that prevents Chalmers from confessing to being a Berkeleian.

(I once put it to David Chalmers, in an internet discussion, that he was a closet Berkeleian. He denied this, but did not give an argument to refute it.)

Chalmers presents his thesis in the form of three hypotheses:

1. The 'creation hypothesis', that physical space-time and its contents were created by beings outside physical space-time.
2. The 'computational hypothesis', that microphysical processes throughout space-time are constituted by underlying computational processes.
3. The 'mind-body hypothesis', that each person's mind is (and has always been) constituted by processes outside physical space-time, and receives its perceptual inputs from and sends its outputs to processes in physical space-time.

The main weight of this thesis is carried by hypothesis no. 2. (Hypothesis no. 1 plays a housekeeping role, and hypothesis no. 3 is not strictly necessary.)

If we accept his thesis, then Chalmers claims that we are not committed to scepticism about the everyday world in the Matrix. He says: "There are still tables, chairs, and bodies: these are made up fundamentally of bits, and of whatever constitutes these

bits." This, however, is dependent on which language-game is being used. Chalmers fails to make explicit the separation of language-games, and in this respect his argument is essentially a less rigorous version of Berkeley's argument.

Suppose I were to enter a virtual reality in which I sit at a desk and type at a laptop. Then there are two language-games to be differentiated. (a) One language-game in which I can properly state that the desk and laptop do exist -- namely the language-game that I am using when talking to other people (i.e. avatars) inside the virtual reality. For example, if a virtual waitress comes by, I can ask her to put the coffee on the desk. It would make no sense for her to reply that there is no desk there. (b) Another language-game, in which I can properly state that the desk and laptop do not exist -- namely the language-game that I use when I log out from the virtual reality and go back the real world.

Of course, Berkeley did not use the term 'language-game', because that was invented by Wittgenstein 250 years later. Instead he used the terms 'vulgar usage' and the 'philosophical usage'.

Now, when Chalmers says that tables and chairs still exist, that statement is true only within the language-game that is internal to the virtual world. Outside that world, we use a different language-game, in which the virtual tables and chairs can properly be said not to exist.

Chalmers says that "the picture is strange and surprising" but that is only because he is failing to make explicit that there are two separate language-games, and that the assertion of the table and chairs' existence is valid only within one of the games. That this is not a particularly subtle point is evidenced by the fact that the myriad users of video games who put themselves into very weak forms of virtual reality have no difficulty at all in comprehending and articulating this difference. I find it surprising that Chalmers wants to say that virtual objects really do exist, when in fact it's clear to everybody that they do *not* exist -- even though we may speak as if they do exist when we are immersed in the virtual world.

So, as I see it, (a) Chalmers is right that beliefs within the Matrix are true, but (b) Chalmers is wrong to imply that they are on an equal footing with beliefs outside the Matrix. In fact, they are not an equal footing because they exist in different language-games.

As long as we are dealing with the language-game inside the virtual world, Chalmers is also right in claiming that we can correctly make statements about the past before we (or our avatars) were born. But, again, in my view, this is limited to utterances made within the Matrix.

5. The implicit simpleness of elementary physical entities

There is an unstated assumption in physics that the elementary entities of physics are genuinely elementary -- that is, they are the basic building blocks out of which the world is made. If Chalmers identifies the elementary entities of physics with computational entities in a Matrix-like simulation, then that assumption no longer holds. For, in that case, obviously those apparently elementary entities are not basic building blocks but are the products of complex underlying processes.

So, Chalmers' project of showing that everyday beliefs and statements are true by virtue of the *identity* of physical entities and computational entities will, in fact, fail. The computational entities remain *only simulations* precisely because they are not the elementary building blocks of the world.

Part of what we mean when we talk about elementary physical entities is that they really are basic. So, as Chalmers identifies subatomic particles with computational entities in a simulation, he is saying that they are not genuinely elementary after all. So our existing beliefs and statements about the physical world are no longer be entirely true.

This flaw can be avoided by recasting Chalmers' argument in terms of language-games. Within the language-game of physics inside the Matrix, it is true to say that elementary physical entities are real and really are the basic building blocks of the world. It is only when we leave the Matrix and rejoin the language-game of the outer world that we are able to state that what seemed to be elementary physical entities within the Matrix are, in fact, only simulations.

6. Semantic chicanery

Later in the essay, when addressing his objection no. 1, Chalmers writes:

*The **brain** is alone in a dark room. But this does not imply that the **person** is alone in a dark room.*

This rather goes against what he wrote at the start of the essay, where he regarded the brain and the person as one and the same.

To be sure, there is indeed a difference between the brain and the person whose brain it is. But that does not help Chalmers deal with his objection no. 2. For, objection no. 2 is essentially the following intuitive conviction: an envatted person is deluded if she believes she is walking around in sunny Tucson when in fact she is stuck in a vat. Now, when that individual believes, and asserts, that she is walking around in Tucson, then she is not referring to her person as such (that is, her mind) but to her body. She is asserting that her feet are pacing the pavement and that the Tucson breeze is blowing through her hair. So, by trying to make out that she is referring to her person, not to her body, Chalmers is engaging in semantic chicanery that has nothing to do with his main argument.

7. Confusion over real and virtual referants

Chalmers' thesis that statements and beliefs within the Matrix are true and on an equal footing with those outside the Matrix leads to a mess of confusion over proper nouns. In his objection no.2, he considers an envatted brain that is held in New York, when the envatted person believes she is in Tucson. How can it be that the person is in Tucson and New York at the same time?

Chalmers' response to this is to backtrack and admit that the envatted person is in a merely simulated world after all:

The envatted being's concept of "Tucson" does not refer to what we call Tucson. Rather, it refers to something else entirely: call this Tucson, or "virtual Tucson".*

But if proper nouns have to be distinguished in this way, then surely every other word must also be similarly qualified. If it is unacceptable to say "I am in Tucson and in New York", then it is also unacceptable to say "I am walking along the street and I am floating in a vat". If we have to replace "Tucson" with "virtual Tucson", then by the same logic we have to say "My virtual self is virtually walking virtually along the virtual street". But that is tantamount to marking the whole language of as a different language. In fact, Chalmers admits this in his response to his objection no. 3, where he writes:

Neo's concept of "hair" does not refer to real hair, but to something else that we might call hair ("virtual hair").*

This undermines the position that Chalmers is trying to defend, which is that the beliefs and statements of the envatted are already true without having to be qualified.

I believe that the correct analysis (on the language-game theory that I am advocating) is that the assertion "I am in Tucson" is true within the envatted person's language-game. But it is false in the language-game of a technician looking after the envatted brain.

As Chalmers works his way through his list of objections, he inevitably finds himself moving way from his main position (that in-Matrix statements are on an equal footing with out-of-Matrix beliefs) and towards a language-game analysis. In response to his objection no. 4, he writes:

*If we are **not** in the matrix, but someone else is, we should say that their term "hand" refers to virtual hands, but our term does not.*

Finally he states the precise opposite of his original thesis:

Our words refer to different things, depending on whether our perspective is inside or outside the matrix.

The way in which this 'perspective' is manifested is through differential language-games. It is not just that individual terms such as 'hair' and 'hands' are switched between referring to real or virtual items. Rather, there are two entire and self-contained language-games. The use of language, and its integration into our social lives, forms a coherent pattern of life. If Trinity says, "Neo, there's an Agent behind you!", then Neo's understanding of this is relative to the language-game: he understands that Trinity's avatar has virtually seen the avatar of an Agent, and that Neo needs to shift his own avatar to dodge the incoming kung-fu chop. Terms from the real and virtual worlds cannot mixed. If Trinity were to say, "Neo, there's a sentinel behind your virtual body!" then she would be speaking nonsense, because the sentinel exists only outside the Matrix. Statements have to belong to one language game or the other.

But this contradicts Chalmers' original thesis that ordinary beliefs are true *simpliciter*. They are not. They are true *within* the appropriate language-game. Chalmers attempt to approach the problem at the level of individual terms (rather than at the level of a whole language-game) leads to such confusion as the following, in his response to objection no. 7:

For various reason I think it is not plausible that words like "action" and "friend" can shift their meanings as easily as words like "Tucson" and "hair". Instead, I think one can say truthfully (in our own language) that the envatted being performs actions, and that it has friends.

Chalmers is in a pickle here, because his logic requires him to insist on the envatted person saying "virtual action" and "virtual friend", and yet he knows that this will cause problems when real brains move between operating in the Matrix world and the outer world. This plausibility problem disappears if we shift language-games rather than trying to shift terms.

8. How detailed must the simulation be?

There is, however, one important detail where Chalmers lays down too strict a condition. He assumes that the simulation must model the physical world down to the

level of atoms and sub-atomic particles:

For an abstract computation to qualify as a simulation of physical reality, it must have computational elements that correspond to every particle in reality (likewise for fields, waves, or whatever is fundamental), dynamically evolving in a way that corresponds to the particle's evolution.

This is neither necessary nor likely. The Matrix computer needs to simulate only enough to render the observations that people make. Unobserved parts of the universe can be omitted from the simulation, and even the observed parts can be simulated at an aggregated level.

This is actually a key plank in Chalmers's argument. For, his central argument is this: the reason that all of our statements about the world carry over into the Matrix world (without loss of truth) is precisely that every physical entity, down to the lowest level of subatomic physics, is present in computational form inside the simulation. If that is false, and if the Matrix simulates unobserved details in an aggregated manner, or simulates only the local area (or does both) then his argument will fail.

Nevertheless, we can still reach the same conclusion by means of a modified argument. For, if we were to examine our actual use of language more carefully, we would find that the fine-grained structure of the physical world is not actually incorporated in what we mean in everyday utterances. When I say that my laptop is on my desk, the meaning of that utterance does not include any statements about the molecular constituents of either the laptop or the desk.

(To be sure, I might infer from what I know of physics certain experimental predictions about what I would find if I were to examine the desk and the laptop more carefully. If I were to scrape off bits of laptop and desk and examine them in an electron microscope, then I would discover blobs that I could interpret as atoms. Now, if we track our actual use of language, we will see that the information carried by the expression "My laptop is on my desk" is not (and cannot be) about the hypothetical unobserved atomic facts, it is (and can only be) about my actual and potential observations. I might, in the future, examine my laptop and desk microscopically. That potential observation is included in the information conveyed by the expression, and is hence part of the meaning. What is not part of the meaning is the proposition P, that the laptop and desk consist of atoms right now. For, that proposition P is not about anything observable. A proposition that is specifically about 'the atoms that I am not now observing' is, self-evidently, not a statement about something observable, and therefore cannot be part of the information that we wish to convey.)

On this view, Chalmers can still assert the truth of our ordinary statements about the world -- even if the simulation is not computed at the microphysical level -- because the microphysics is not part of the meaning of everyday beliefs and statements.

9. Just-in-time reality generation

In July 1999, I published [Paranormal Phenomena and Berkeley's Metaphysics](#), a book in which I defended George Berkeley's theory that our world is virtual, and that all our conscious experiences are generated by a vast external consciousness. (Berkeley called it 'God', I called it the 'metamind'.) In section 4.2.4, I proposed that this virtual world operates by "just-in-time object generation", in which the metamind generates objects as and when we are about to observe them.

Chalmers has made use of this term in describing an alternative to the full-scale physics simulation.

Extendible Local Matrix Hypothesis: *I am hooked up to a computer simulation of a local environment in a world, extended when necessary depending on subject's movements. ... This sort of simulation is quite unlike simulation in an ordinary matrix. In a matrix, the whole world is simulated at once. There are high start-up costs, but once the simulation is up and running, it will take care of itself. By contrast, the extendible local matrix involves "just-in-time" simulation. This has much lower start-up costs, but it requires much more work and creativity as the simulation evolves.*

(I am not entirely convinced by Chalmers' accounting of the costs. A full physics simulation will not simply trundle along by itself. The underlying computer will actively have to execute the simulation of every part of the universe at every moment, which must surely be a vast undertaking. Furthermore, the total cost of simulationship (to paraphrase Microsoft) is much lower for the just-in-time simulation because human beings collectively will ever observe only a minuscule part of all the facts of the physical universe.)

The just-in-time model is where we see the weakness of Chalmers' approach to the analysis of truth and meaning. In this extendible simulation, he says that beliefs and statements about physical states of affairs that happen to lie outside an observed space are necessarily false. Certainly, this will make no sense to any who has always lived inside such a simulation. For that person, beliefs about local things are on a par with beliefs about remote things. Moreover, what Chalmers is proposing here is quite at odds with how people do successfully and meaningfully talk about real and virtual worlds.

When discussing events and things within a virtual world (be it a present-day video game, or some Matrix in the future), people talk about them just as if they were real. That manner of speaking is, however, limited to that language-game. The language-game will be demarcated by clear social cues. Imagine, say, a group of VR designers who spend all morning discussing the plot for a new VR product -- talking all the time as if the events were real. When they break for lunch, they shift into a different language-game, in which they talk about going to a (real) Starbucks. There is no confusion. Yet, there is also no concern with whether the simulation will model the whole world or whether it will be an extendible simulation.

Thus, the validity of ordinary beliefs and language with the Matrix is unaffected by the nature of the underlying simulation, be it a whole-world simulation or a just-in-time simulation.

Similarly, a Matrix-like simulation that is limited to aggregated, macroscopic things -- rather than microphysical things -- supports the ordinary beliefs and statements of the people who live in that simulation (including physicists). Again, Chalmers' approach requires him to say that microphysical beliefs are false inside such a simulation. So, according to Chalmers, a physicist living inside the Matrix has a true belief that he has a pen in his hand but a false belief that the pen is made of atoms; although if he scrapes off bits of the pen puts them under an electron microscope then he will have a true belief that there are atoms in his pen sample. I do not think that this captures the meaning of 'true'. I think we must adopt the language-game approach: within the physicist's language-game, it is true to assert that the pen is made of atoms.

10. Simulations and Berkeley

I was glad to see that Bishop Berkeley gets a positive mention later on in Chalmers' essay, as the good Bishop has been rather neglected in commentaries on *The Matrix*. Chalmers has, however, co-opted Berkeley into the whole-world simulation that Chalmers advocates.

A hypothesis like this was put forward by George Berkeley as a view about how our world might really be. Berkeley intended this as a sort of metaphysical hypothesis about the nature of reality. Most other philosophers have differed from Berkeley in regarding this as a sort of skeptical hypothesis. If I am right, Berkeley is closer to the truth. The God Hypothesis can be seen as a version of the Matrix Hypothesis, on which the simulation of the world is implemented in the mind of God. If this is right, we should say that physical processes really exist: it's just that at the most fundamental level, they are constituted by processes in the mind of God.

In fact, Berkeley explicitly denied that we could know the extent or level of detail of thinking in God's mind. Whether God was running a whole-world, full microphysical simulation, or whether he was running an extendible, macroscopic simulation, is not something that we can ascertain without somehow having access to God's mind. In his *Principles*, section 71, Berkeley compares God's internal database to a musician's score. Just as a musician renders music from the glyphs written in the score, so God renders our sensory perceptions from his private mental glyphs. What those glyphs are, we cannot tell, but Chalmers is jumping the gun in assuming that God's glyphs make up a complete simulation of everything in the physical world.

In fact, in [Paranormal Phenomena and Berkeley's Metaphysics](#), I argued that God probably did not simulate all physical things. Like most things in nature, God's mind is probably economical, and therefore simulates only what it needs to, in order to provide our conscious perceptions. Furthermore, when that assumption is made, certain paranormal phenomena (such as retrokinesis) become amenable to a naturalistic explanation.

11. Conclusion

Chalmers has written a pro-Matrix essay which attempts to show that everyday beliefs and statements held by people inside the Matrix are indeed true. They are true in so far as they refer to observable 'physical' processes. Their ultimately virtual nature is not something that can be observed by physical experimentation inside the Matrix.

I have argued that, whilst this conclusion is broadly correct, certain criticisms can be made. (a) The argument is, in a sense, too strong -- as it implies that in-Matrix statements are on a par with out-of-Matrix beliefs. Yet, as even Chalmers acknowledges, a distinction must be made between beliefs and assertions in the Matrix and those outside the Matrix. (b) The argument, in another sense, is too weak, as it fails to imply the truth of everyday beliefs and assertions in local extendible simulations and macroscopic simulations.

I have suggested that a Berkeleian semantic argument, in which we differentiate language-games inside the Matrix from those outside it, enables us to reach the same conclusion with the flaws (a) and (b).

© 2003, Peter B Lloyd

