The Virtual and the Real

David J. Chalmers

Is virtual reality truly real? The most common view is that virtual reality is a sort of fictional or illusory reality, and that what goes in in virtual reality is not truly real. I will defend the opposite view: virtual reality is a sort of genuine reality, and what goes on in virtual reality is truly real.

The issue manifests itself in a number of questions. Are virtual objects, such as the avatars and tools found in a typical virtual world, real or fictional? Do virtual events, such as a trek through a virtual world, really take place? When we perceive virtual worlds by having immersive experiences of a world surrounding us, are our experiences illusory? And are experiences in a virtual world as valuable as experiences in a nonvirtual world?

Here we can distinguish two broad packages of views. A package that we can call *virtual realism* (loosely inspired by Michael Heim's 1998 book of the same name) holds:[1]

(1) Virtual objects really exist.

(2) Events in virtual reality really take place.

(3) Experiences in virtual reality are non-illusory.

(4) Virtual experiences are as valuable as non-virtual experiences.

A package that we can call *virtual irrealism* holds:

(1) Virtual objects do not really exist.

(2) Events in virtual reality do not really take place.

(3) Experiences in virtual reality are illusory.

(4) Virtual experiences are less valuable than non-virtual experiences.

[1]The glossary to Heim's book *Virtual Realism* characterizes virtual realism as "The pragmatic interpretation of virtual reality as a functional non-representational phenomenon that gains ontological weight through its practical applications" (p. 220). He also says that virtual realism affirms that "Virtual entities are indeed real, functional, and even central to life in coming eras". Although Heim's discussion largely focuses on social and technological issues distinct from those discussed here, perhaps these passages justify my adapting his term in the way suggested here.

The four theses in each package are separable from the others, and it is possible to hold just one or two of the theses in each package. But the theses in each package go especially naturally together. Each thesis needs clarification, which I will give in what follows.

The divide here mirrors a familiar ambiguity in the very label "virtual reality". On one common meaning of "virtual", "virtual X" means something like "as if X but not X" (consider: virtual kitten). On that reading, virtual reality is an unreal as-if reality, and virtual reality is no more reality than a virtual kitten is a kitten. On other more recent meaning, "virtual X" means something like "a computer-based version of X" (consider: virtual library). On that reading, virtual reality may be a sort of reality, just as a virtual library is a library.[2]

I have explored the philosophical status of virtual reality once before, in my 2003 article "The Matrix as Metaphysics". That article focuses on a perfect and permanent virtual reality such as the one depicted in the movie *The Matrix*. In that article, I argued that if we are in a Matrix, most of our ordinary beliefs (e.g. that there are tables) are true: if we discovered that we are in a Matrix, instead of saying that there are no tables, we should say instead that tables are computational objects made of bits. In effect, I answered questions (1)-(4) by saying that at least in the case of a permanent and perfect virtual reality:

(1) Virtual objects really exist and are computational objects;

(2) Events in virtual worlds are largely computational events that really take place;

(3) Experiences in virtual reality involve non-illusory perception of a computational world;

(4) Virtual experiences of a computational world are about as valuable as non-virtual experiences of a non-computational world.

We might call the combination of (1) and (2) *digitalism* about virtual reality. In this article, I will in effect extend the digitalist view that I have defended for permanent and perfect virtual reality to give the same answers even for the temporary and imperfect virtual realities that are possible with current VR technology. I will not presuppose any knowledge of the earlier article in doing this, but I will make occasional connections to it.

[2] A similar ambiguity is familiar with the expression "artificial intelligence", where "artificial" can be understood as "as-if" or as "synthetic".

1 Definitions

First, what is virtual reality? There is no universally accepted definition of virtual reality, and the concept exhibits some vagueness and flexibility. Still there is a common core to most uses and definitions of the term.[3] Capturing this core, I will say that a virtual reality environment is an immersive, interactive, computer-generated environment. "Virtual reality" simpliciter can be considered a mass noun covering virtual reality environments and/or the technology that sustains them. The three key notions of immersion, interaction, and computer generation can be explained as follows.

Immersion: An immersive environment is one that generates perceptual experience of the environment from a perspective within it, giving the user the sense of "being there": that is, of really being present at that perspective. Typically this will involve at least a visual experience as of a three-dimensional environment, perhaps along with auditory and other sensory elements. In the present day, a paradigm of immersive VR technology involves a headset (such as the Oculus Rift) with a stereoscopic display. In the future one can imagine that glasses, contact lenses, or implants could accomplish the same thing.

Interaction: An environment is interactive when actions by the user make a significant difference to what happens in the environment. In current VR, this interaction takes place through the use of input devices such as head- and body-tracking devices, handheld controllers, or even a computer keyboard.

Computer generation: An environment is computer-generated when it is grounded in a computational process such as a computer simulation, which generates the inputs that are processed by the user's sensory organs. In current VR this computation usually takes place either in a fixed computer connected to a headset display or in a mobile computer (such as a smartphone) embedded in a headset using its own display.

We can also say that *virtual reality technology* is technology that sustains virtual reality environments. "Virtual reality" as a count noun is roughly synonymous with "virtual reality environment", while as a mass noun it covers both virtual reality environments and virtual reality technology.

[3] For example, Heim (1998) defines virtual reality as "an immersive, interactive system based on computable information." My definition is close to this one, but I think it is best to talk about environments rather than systems in order to exclude cognitive systems (a conscious AI system perceiving and interacting with a physical environment, say) from counting as VR.

The term "virtual reality" is sometimes used in looser ways than this – sometimes so loose as to capture almost any nonstandard means for generating experiences as of an external environment. To allow distinctions between grades of VR, we might say that "VR proper" is virtual reality that satisfies all three conditions above. We can then capture more inclusive notions of virtual reality by removing the three conditions one at a time.

Nonimmersive VR includes environments displayed on desktop computer or television screens, as with many familiar videogames. *Noninteractive VR* includes passive immersive simulations such as computer-generated movies presented on a VR headset. *Non-computer-generated VR* includes immersive and interactive camera-generated environments, such as the remote-controlled robotic VR sometimes used in medicine. The label of VR is also occasionally applied to environments satisfy just one of the three conditions: immersiveness (e.g. movies filmed with 360-degree cameras and displayed on a headset), interactiveness (e.g. remote control of a robot using a desktop display of its perspective), or computer generation (e.g. a computer-generated movie displayed on a desktop). The label is not typically applied to environments that satisfy none of the three conditions, such as ordinary (two-dimensional, passive, camera-based) movies and television shows. That said, it is interesting to note that the term "la realite virtuelle" was first introduced by Antonin Artaud (1938) to apply to the theatre, which is typically noninteractive and non-computer-generated and arguably nonimmersive.

There are also intermediate cases. So-called *augmented* or *mixed* reality involves immersive and interactive environments that are partly computer-generated. In these cases a physical environment is combined with computational processes to yield experiences of an augmented or mixed environment with both ordinary physical objects and virtual objects. Augmented or mixed reality is typically contrasted with VR, but it can also be considered as VR in an extended sense. Ordinary un-augmented physical environments are also immersive and interactive, but they are not usually considered to be VR, except perhaps by some who think that the external world is computer-generated or that it is a mind-generated construction.

What is a virtual world? I take a virtual world to be an interactive computer-generated environment, of the sort that we (seem to) inhabit when using virtual reality. On common usage, nonimmersive desktop videogames such as *World of Warcraft* involve virtual worlds even if they are not strictly virtual realities, so there is no immersiveness condition on virtual worlds. We will see that when it comes to ontological issues about virtual worlds, nonimmersive and immersive VR raise very similar issues, so it makes sense to drop the immersiveness condition in this domain for a broader analysis.

What is a virtual object? I take these to be the objects that are contained in virtual worlds and that we (seem to) perceive and interact with when using virtual reality. Paradigmatic virtual objects include avatars (virtual bodies), virtual buildings, virtual weapons, and virtual treasures.

These definition are neutral on whether virtual worlds and virtual objects are real or unreal. I take it that realists and irrealists can both agree that virtual worlds are computer-generated, that we seem to inhabit them, and that virtual worlds contain virtual objects that we seem to interact with. For example, whether the world of Azeroth in *World of Warcraft* is a digital world or a fictional world, it is computer-generated, we seem to interact with it, and it contains virtual objects either way.

2 Virtual Fictionalism

The large majority of philosophers who have written about virtual worlds are virtual irrealists. More specifically, they hold that virtual worlds are fictional worlds. We might call this view *virtual fictionalism*.[4] On this view, virtual worlds have a status akin to Tolkein's Middle Earth, and virtual objects have a status akin to that Gandalf or the One Ring: they do not exist in reality, but only in fiction. Likewise, the things that are supposed to happen to them do not happen in reality, but only in fiction.

Virtual fictionalism can naturally be associated with the following cluster of views on our original question (though certainly not all virtual fictionalists need endorse all of these theses):

(1) Virtual objects are fictional objects.

(2) Virtual events take place only in fictional worlds.

(3) Experiences in virtual reality involve illusory perception of a fictional world.

(4) Virtual experiences have the sort of value that engagement with fiction has.

Fictionalism is an especially natural thesis for the virtual worlds that are present in many videogames. For example, there are many videogames based on Tolkein's works and set in Middle

[4]Varieties of virtual fictionalism are expounded by Juul 2005, Tavinor 2009, Bateman 2011, Velleman 2011, and Meskin and Robson 2012. Some of these fictionalists also distinguish special respects in which virtual realities are real: for example, they involve real rules (Juul) or agents who literally perform fictional actions with fictional bodies (Velleman). Aarseth (2007) denies that virtual worlds are fictional while nevertheless holding that they are not real: they have the same sort of status as dream worlds and thought experiments, which he also understands as not fictional.

Earth. If the Middle Earth of the books is fictional, so presumably is the Middle Earth of the games. There are also videogames set in historical periods such as the Second World War, while depicting events (such as the assassination of Hitler) that did not really take place then. A book or movie depicting these events would be fictional, and the same goes for a videogame.

It is misleading to take videogames as one's prime model for virtual reality, however. There is of course a close connection between any role-playing game and an associated fiction, but this connection holds whether the game is virtual or non-virtual. If a human in physical reality plays the role of Gandalf casting a spell in Middle Earth, the event of Gandalf casting a spell is fictional, but the underlying bodies and movements are real. Likewise, if an avatar in virtual reality plays the role of Gollum stealing the ring, the event of Gollum stealing the ring is fictional, but this is consistent with the underlying avatars and movements within the virtual realm being real.

Furthermore, videogames are just one among many possible uses of virtual reality technology. At the moment, videogame worlds are the most popular virtual worlds, but there are already many virtual worlds that are not especially game-like in character. When a virtual world is used for non-play purposes such as gathering information or communicating with colleagues, it is much harder to discern fictionality in the virtual world.

The well-known virtual world Second Life, for example, is generally characterized as a platform rather than a game. There is no special objective in the world of Second Life. Users can use the world for activities and interactions of all sorts. Suppose I enter the virtual environment of Second Life in order to have a conversation with a friend. In what sense is what goes on fictional? I am really having a conversation with my friend: this is not fictional at all. Presumably if there is a fiction here, it involves our avatars. For example, perhaps the virtual world depicts us as having certain bodies which we do not really have, and depicts our bodies as being a few meters apart when in fact we are thousands of miles apart.

I think this is the wrong way to think about Second Life and other virtual worlds. The right way is this. The virtual world of Second Life involves virtual bodies (avatars) in virtual space. Virtual bodies are distinct from physical bodies, and virtual space is distinct from physical space. We really have these virtual bodies, as well as having physical bodies. There is nothing fictional about this. These virtual bodies really inhabit virtual space, where they are really a few (virtual) meters apart. There is nothing fictional about this. If I pick up a virtual coin in Second Life, I really use my virtual body to take possession of a virtual coin. There is nothing fictional about this.

I will defend this picture in what follows.

3 Virtual Objects

What are virtual objects? In my view, they are digital objects, constituted by computational processes on a computer. They are perhaps best regarded as data structures, which are grounded in computational processes which are themselves grounded in physical processes on one or more computers. In some cases, multiple data structures may be associated with a single virtual object, in which case the virtual object will be a higher-level entity constituted by these data structures.

Corresponding to each avatar in *Second Life*, there is a data structure on the *Second Life* servers (perhaps distributed redundantly across many servers). When I see an avatar, it is this data structure that brings about my perception. What I perceive directly reflects the properties of this data structure: the perceived location of the avatar reflects one property of the data structure, while the perceived size, color, and so on reflect other properties. When my avatar interacts with a coin, the two data structures are interacting. Whenever two virtual objects interact in Second Life, there is a corresponding interaction among data structures. Data structures are causally active on real computers in the real world; the virtual world of Second Life is largely constituted by causal interaction among these data structures.

This gives rise to the first argument for digitalism: the argument from causal powers.

(1) Virtual objects have certain causal powers (to affect other virtual objects, to affect users, and so on).

(2) Digital objects really have those causal powers (and nothing else does).

---

(3) Virtual objects are digital objects.

Of course this is not a knockdown argument against the fictionalist. Fictionalists will probably deny the first premise by saying that virtual objects do not have causal powers, or better, that they have causal powers only in the sense that Gandalf has causal powers. That is, they have causal powers within a fictional world, and any effects on the real world are brought about not by the object but by a representation of the object. Still, even the nonconclusive argument from the premise that virtual objects seem to have these causal powers and that digital objects really have those powers is a reasonably strong one. If there are real objects that have all the apparent properties of virtual objects, there is not much reason to suppose that virtual objects really belong to a separate layer of fictional objects.

A closely related argument is the argument from perception:

(1) When using virtual reality, we perceive (only) virtual objects.

(2) The objects we perceive are the causal basis of our perceptual experiences.

(3) When using virtual reality, the causal bases of our perceptual experiences are digital objects.

---

(4) Virtual objects are digital objects.

Here premise (1) is intuitively plausible, and premise (2) is a widely accepted claim in the philosophy of perception. Premise (3) seems to be empirically correct. A data structure in the computer is causally responsible for generating my experience. One might suggest that an image on the display screen is the relevant causal basis, but a moment's reflection suggests that this cannot be right: multiple people may see different images on different displays while they all perceive the same virtual object. Just as many people can see the same actor by watching TV on different screens, because that actor that is the causal basis for all the images, many people can see the same digital object by experiencing virtual reality with different displays, because that digital object is the causal basis for all the images.

Once again, this is not a knockdown argument. A fictionalist can reply that this is really a case of hallucination in which no real object is perceived. In some cases of hallucination, there is a causal basis for the perception: for example, a chair in the environment might trigger an auditory hallucination of a voice, but one does not hear the chair. Still, the fact that in this case there are objects that serve so systematically as the causal basis of the experience makes this line harder to maintain.

It is widely accepted that when we look at a photograph or a film clip of Winston Churchill, we see Winston Churchill. We may see the photograph or the screen as well, but we see Churchill when we see the screen (seeing him *in* the photograph or screen, as Richard Wollheim has put it). The reasons for saying this include that Churchill was the causal basis of our experience, and the features of or experience depend systematically on the features of Churchill when he was filmed. Both of these reasons apply to seeing digital objects in virtual reality. Furthermore, in at least three respects virtual reality is *more* like ordinary seeing than seeing a film. First, in typical VR, one need have no sense of seeing a screen, and it can perhaps be argued that one does not really see the screen at all. Second, in VR one has immersive three-dimensional perceptual experience from a perspective. Third, in typical VR one can move around in response to what one sees, change

one's perspective, and act on the world. On the other side of the ledger, it might be objected that in ordinary perception, the experience matches the object, in that colors and shapes that things seem to have roughly reflect their actual colors and shapes, while in virtual reality they do not. We will see that the perceived colors and shapes at least match the *virtual* color and shape of a digital object, though, and that the perception here need not be illusory at all.

The fictionalist may try a counterargument along the following lines:

(1) My avatar is a dragon.

(2) No real object is a dragon.

———————————

(3) My avatar is not a real object.

In response, we need to distinguish physical dragons from virtual dragons. In the virtual world, there are no physical dragons, but there is a *virtual dragon*. In the real world, there are no physical dragons (giant creatures breathing real fire), but there are numerous virtual dragons (digital objects existing on computers in that world). Once the distinction is made, the conclusion does not follow from the premises. The virtual world, virtual dragons and all, is part of the real world, in virtue of existing on real computers.

Of course virtual objects do not *look* like digital objects, at least to the naive user. If one knows little about virtual reality, it may be surprising to discover that the objects that one is seeing and interacting with in VR are digital objects grounded in tiny chips on computer servers. In this respect, the claim that virtual objects are digital objects is a little like the claim that (apparently tiny) stars are enormous exploding balls of gas. It is also analogous to theoretical identifications such as the claim that water is H_2O or that lightning is electric discharge. One cannot tell that water is H_2O just by looking at it or thinking about it; one needs to know about the underlying processes. Likewise, one cannot know that virtual objects are digital objects just by looking at them or thinking about them; one needs to know about the underlying processes.

4 Virtual Properties and Virtual Events

I suspect that the real sticking point for many fictionalists involves *events* and *properties* in virtual worlds. In a virtual world, a virtual dragon flies through their air. In the real world, the corresponding digital object does not fly through the air, and indeed no real object flies through the air

as the virtual dragon does. If so, then either the virtual dragon is not real, or it is real but it does not really fly through the air. Either way, the event of the virtual dragon flying through the air is fictional. This conclusion seems to follow whether virtual objects are digital objects or not.

The same issue arises for *properties* in virtual worlds, such as colors and sizes. A virtual flower may be red, while the corresponding digital object is not red. Indeed, no object in the real world may have the precise shade of red that the virtual flower has. If so, then either the virtual flower is not real, or it is real but it does not really have the property of being red. Either way, the apparent redness of the virtual flower is fictional. In a similar way, when my avatar is apparently six feet tall, its having this property is fictional.

To address this concern, we have to get clear about properties and events in virtual worlds. In particular, just as we distinguished virtual objects from non-virtual objects, we have to distinguish virtual properties from non-virtual properties. A virtual flower is not red in the ordinary sense (non-virtually red), but it is virtually red. The corresponding digital object is also not red in the ordinary sense, but it is virtually red. My avatar is not six feet tall in the ordinary sense (non-virtually six feet tall), but it is virtually six feet tall. The corresponding digital object is also not six feet tall in the ordinary sense, but it is virtually six feet tall.

What is virtual redness? To answer this, we can step back and ask: what is redness? On an orthodox view, the property of redness is picked out in virtue of a certain sort of *effect*: in particular, the fact that red things normally cause red experiences. On one version of this view, redness is just the power to cause red experiences in normal circumstances. On another version, redness is the intrinsic property (a physical property of a surface, say) that causes red experiences in normal circumstances. There are some differences between these views, and more refined versions of each, but the differences will not matter much for our purposes. Views of this sort are sometimes called *functionalism* about color, because they understands colors in virtue of their functional (or causal) role.

Red roses are red, then, because they produce reddish experiences in the conditions that are normal for human perceivers. The digital object corresponding to a virtual red rose is not red,

[4]One refined view holds that redness is the physical property that normally brings about reddish experiences (or a disjunction of such properties). Another holds that redness is the higher-order property of having a physical property that normally brings about reddish experiences. These views arguably handle certain cases better, such as cases of systematic illusion in which a white object normally looks red. These views can also be seen as functionalist in a broad sense (the physical-property view is sometimes called realizer functionalism, while the other two are versions of role functionalism). One can straightforwardly generalize all these views to the virtual case just as I generalize the simple view below.

because it does not produce reddish experiences in normal conditions. Under normal conditions (that is, looking at the circuit with the naked eye) data structures are not really visible at all, but if they were visible there is no reason for them to produce reddish experiences. Now, the digital object *does* produce reddish experiences when it is accessed in a certain special way, namely through a virtual reality headset. Using a virtual reality headset is not (yet) a normal condition for ordinary human perception, so this is not enough to make the digital object count as red in the ordinary sense. But it is enough to make the object count as virtually red.

We can say that an object is *virtually red* when it produces reddish experiences in the conditions that are normal for virtual reality. Normal conditions for virtual reality currently involve access through an appropriate headset. The data structure corresponding to a virtual red rose really does cause reddish experiences when viewed in these conditions, so the data structure is virtually red. This allows us to say that the virtual rose is virtually red, even though it is not non-virtually red.

What is virtual redness? As before it might be construed either as the power to cause reddish experiences in normal VR conditions, or as the property that normally causes reddish experiences in those conditions. In any given VR environment, some digital property or properties will normally cause reddish experiences. In simple cases, these will involve certain values for an entry in a data structure. When a digital object has an entry with whose value is in the right range, the object is virtually red. In other cases the digital property will be more complicated, but the basic structure is the same. Virtual redness itself might be construed as a disjunction of all of these properties across different VR environments, or simply as the higher-order property of having some property that normally causes reddish experiences in the relevant environment.

(Of course if we are embedded in a permanent virtual reality, as in *The Matrix*, then virtual perception will be normal for us, and the virtual roses that normally cause our reddish experiences will be red in the ordinary sense.)

[4]What if the same data structure is used to generate reddish experiences for some users and grayish experiences for others, for example because of different headsets (color vs black and white)? If the first sort of headset is normal and the second is not, then one can reasonably say that the data structure is virtually red and not virtually gray, just as we say that apples are red even though colorblind people see them as gray. What if the same data structure is used in two quite different (and widely used) VR environments, normaly causing reddish experiences in one and normally causing greenshing experiences in the other? Then is is tempting to say that instances of the data structure in the first environment are red and those in the second environment are green. One model that can handle this: an object is virtually red if is in environment E and has property P such that P normally causes reddish experiences in E. There is a delicate issue about how finely to individuate environments here.

In other work, I have argued that something like this model applies to spatial properties too. That is, to a first approximation, an object is one meter tall when it normally causes experiences of being one meter tall. An object is square when it normally causes squarish experiences. And so on. This sort of view can be called *spatial functionalism*, because it understands space in terms of its causal role (though I am simplifying away from many important details here).

We can then say that virtually square objects are objects that produce squarish experiences under conditions that are normal for virtual reality. The digital object corresponding to a square table in a virtual world is probably not square in the ordinary sense, but it is virtually square. Likewise, the digital object corresponding to my avatar is not six feet tall in the ordinary sense, but it is virtually six feet tall. The virtual height of an avatar can be understood as the feature of the associated data structure (a value of a certain element, say, or a complex property that depends on many underlying elements) that typically brings about six-foot-tallish experiences.

One can also use spatial functionalism to understand virtual space in terms of the causal interactions between virtual objects. Here one inspiration is Brian Cantwell Smith's epigram "Distance is what there's no action at". The idea is that spatial relations (tend to) serve as a measure of causal interactions. We can combine this with "Distance is what there's no motion at", in effect imposing a constraint that space (tends to) serve as a locus of continuous motion. We can use these constraints to define a distance metric in terms of the dynamic evolution of and interaction between underlying objects. This applies equally to virtual objects. A virtual space is the space that virtual objects tend to vary continuously within, and tend to interact at short distances within. Two neighboring virtual objects will correspond to digital objects with much potential for causal interaction. Spatial functionalism in effect allow us to understand virtual space in terms of dynamic interactions in a digital world.

There is a lot more to say about virtual space, but this is enough for now. Virtual objects exist in their own virtual space, in virtue of their effects on each other and on our experiences. A digital object may exist simultaneously in non-virtual space (in a circuit board in a computer in a warehouse, for example) and in virtual space (outside on a virtual beach somewhere).

There are many virtual spaces. Every virtual world has its own virtual space. On my iPhone alone, there are dozens of virtual worlds, each with its own virtual space. The same data structure may occasionally be located in multiple virtual spaces (as well as in physical space), but it is more common to be located in just one. As usual, these virtual spaces are held together by their effects on users and their interactions with each other.

For any property X, there will be a corresponding virtual property virtual X. When a non-

virtual object has X, the corresponding virtual object will have virtual X. When X is picked out by a functional role, as for many properties, then virtual X will work as above. In other cases, virtual X may work quite differently.

For example, we saw earlier, there are some X for which a virtual X is an X:[5] for example, a virtual library is a library, and a virtual calculator is a calculator (at least if the virtual version is understood to simulate all the details, as opposed to merely a facade). In these cases, we can say that the digital object corresponding to a virtual library is not just a virtual library: it is really a library. Likewise, the digital object corresponding to a virtual calculator is really a calculator.[6]

With the distinction between non-virtual and virtual properties and events in hand, most arguments against the reality of virtual objects and events can be defused. The arguments presuppose that in virtual worlds, objects have certain *non-virtual* properties (such as flying in a certain way), which no real object has. But this premise is usually false: in virtual worlds, objects have only the corresponding virtual property (say, virtually flying in a certain way). Real digital objects really have those virtual properties, so there is no obstacle to the events in virtual worlds really occurring.

5 Is Perception of Virtual Reality Illusory?

Is perception in virtual reality illusory? If virtual objects are not real, then perception of them is a sort of hallucination, akin to perceiving a pink elephant. Even if virtual objects are real, however, perception of them might still be illusory, because we perceive virtual objects as having non-virtual properties that they do not really have.

A simple argument for virtual illusionism (as one might call it) goes as follows:

[6] In *The Conscious Mind* (1996) I discuss the conditions under which a simulated X is an X: it is when the property of being X is an organizational invariant, depending only on the abstract causal organization of the underlying system. Roughly the same goes for the question of when a virtual X is an X, although the virtual case adds to the simulation case the extra constraint that certain mental properties (the experiences of a user) are also present. If we assume that all mental properties from a non-virtual situation are duplicated in a corresponding virtual situation (perhaps because all minds are taking part as users of the virtual reality, or perhaps because they are brought about by the simulation?), then a virtual X will be an X when the property of being X can be analyzed in causal and mental terms. This is roughly the view that I defended in "The Matrix as Metaphysics" (2003). Philip Brey (2003; 2014) addresses the same question and answers that a virtual X is an X if and only if X is an *institutional* kind. I think that the "only if" claim is not quite right: virtual calculators are calculators and virtual love affairs are love affairs, where both are causal/mental kinds though neither are institutional kinds. But it is plausible that most institutional kinds are causal/mental kinds, so Brey's "if" claim is plausible.

1. We perceive virtual objects as having real (non-virtual) colors, locations, and shapes.

2. Virtual objects do not have real (non-virtual) colors, locations, and shapes.

3. If one perceives an object as having properties that it does not have, the perception is illusory.

---

4. Perception of virtual objects is illusory.

I have already argued for premise 2, and premise 3 can be regarded as a definition of "illusion". Premise 1 has some initial plausibility, but I will argue that it is false.

To make things more straightforward, we can start with the corresponding issues about belief rather than about perception. Are users' beliefs about virtual worlds false?

In response, it is natural to say that a naive user of virtual reality may have false beliefs. In particular, if someone is put in a virtual reality without knowing it is a virtual reality, they will probably come to believe that they are interacting with non-virtual objects in physical space. Even if they are told that it is a virtual reality, naive users may initially be unable to suspend the visceral belief that objects are present in a certain configuration in physical space, though on reflection they may judge that they are not.

For the sophisticated and experienced user of virtual reality, on the other hand, there is much less danger of acquiring false beliefs of this sort. Given that the user knows they are using VR, they will not form the belief they are interacting with non-virtual objects in physical space. They will know full well that they are interacting with virtual objects in virtual space.

What about perception? Naive users will plausibly suffer illusions. If they do not know they are in a virtual reality, they will certainly perceive the objects they are seeing to be around them in physical space. Even if we tell them they are in a virtual reality, it seems plausible that the illusion will persist. It is tempting to say that the illusion is hard-wired in so that it will persist for sophisicated users, but I am not so sure.

Here I think one can make a useful analogy with perception in mirrors. Does mirror perception, when one sees oneself or other objects in a mirror, involve an illusion? Here the relevant illusion is that the object looks to be on the far side of the glass, when in fact it is on the near side of the glass. A naive user who does not know they are using a mirror will plausibly experience an illusion of this sort. Even when they are told that they are told about mirrors, the illusion may persist. But it

is not so clear that the illusion persists for sophisticated and experienced user of mirrors.

Consider a car's rear-view mirror, as used by an experienced driver. When the driver looks in the mirror and sees cars that are actually behind her, does she have the visual illusion that the cars are in front of her, pointing toward her? This claim seems phenomenologically wrong to me. When I look in my rear-view mirror, the cars I see look to be behind me.

Now, someone who thinks that rear-view mirror is illusory will say that we *judge* that the cars are behind us, while nevertheless the cars *look* to be in front of us. Or perhaps they might allow that cars look to be behind us, but only in a sense where "look" is tied to judgment and other aspects of cognition—while at the level of visual perception, visual experience represents the cars as being in front of us. Once again, however, I think this gets the phenomenology of visual experience wrong.

There are certainly some cases in which mirrors yield illusions, even when the user knows that mirrors are present. One obvious case is a double mirror yielding an image of an infinite series of people— even though we know full well that just one person is present, it *looks* as if there are many people present. There are also cases in which objects seen in mirror clearly look to be on the other side of the mirror, despite knowing it is a mirror. One such case is the mirror box experiment in which one sees a reflection of one's left arm in the mirror, and it looks to be one's right arm on the far side of the mirror. These are cases where mirrors genuinely yield illusions. But the phenomenology of these cases is quite different from that of the rear-view mirror cases. In these cases, one has visual experience as of objects on the other side of the glass. In the rear-view mirror case, one has visual experience as of objects on the near side of the glass.

What are the key features of the rear-view mirror case that differentiates it from the various illusion cases and makes it a plausible case of non-illusion? One obvious factor is *knowledge*: we know a mirror is present. Another is *familiarity*: we are used to using mirrors, and we are especially used to using mirrors in this configuration. A crucial related factor is *action*: we have patterns of action that depend on a certain interpretation of what is seen in the mirror. For example, we may accelerate or turn depending on where we take the objects seen in the mirror to be. A fourth factor that may sometimes play a role is *naturalness*: the interpretation on which cars are on the other side of the mirror is extremely unnatural (it seems to require a narrow line of cars facing toward one amidst a an entirely different landscape, with an abriupt discontinuity between them), whereas the interpretation on which cars are on the near side of the mirror is much more natural.

The role of knowledge suggests that this is a case of *cognitive penetration*: that is, a case where cognition influences perception. Typical cases are cases where what one knows or believes

influences what one perceives. It is controversial whether there are any such cases (Macpherson, Scholl and Firestone, etc), but the mirror case is one of the more plausible examples. One can set up two parallel cases in which a subject sees a chair in a mirror, where in one case the subject believes a mirror is present and in the other the subject believes a window is present. The two subjects may have quite different visual experiences: the chairs appears to be on the near side of the glass for one subject, and on the far side for another. This suggests a direct dependence of perceptual appearance on belief.

We might call this sort of cognitive penetration *cognitive orientation*. In this case, background knowledge helps *orient* one to the perceived world, giving a global interpretation to what is perceived. To deny that this sort of cognitive orientation ever takes place with mirrors, an opponent will probably have to take a hard line and deny that objects seen in mirrors ever appear to be on the near side of the glass.

The phenomenon of cognitive orientation naturally extends to video. If one's car uses a rear camera instead of a rear-view mirror, for example, after a while one will perceive objects seen on the screen as behind the car. Something similar goes for side cameras. One could extend the phenomenon to cameras on remote cars or on robot bodies, where on expert use one will perceive objects as standing in a certain relation to the remote car. We can also extend to different scales. If a camera is attached to a tiny robot, such as the shrunken submarine in the movie *Fantasic Voyage*, then an expert user will not undergo the illusion that the objects are much larger than they are: instead they will correctly perceive small objects as small.

From here it is not an enormous step to virtual reality. The virtual reality case is in many ways parallel to the mirror case. A naive user who does not know they are using virtual reality will undergo the illusion that certain objects are present in physical space in front of them. After they learn they are using virtual reality, they will not be fooled into believing that the objects are present, but the perceptual illusion will persist. After some time, however, a sophisticated user will become familiar with VR, and they will act in ways that turn on interpreting themselves to be in VR. For example, they may learn to use the distinctive input controllers of VR. They will learn just how far they need to reach or to step to get to a certain virtual location. They will learn to exploit distinctive affordances in current VR: for example, the ability to walk right through many virtual objects. All this will give them a sort of cognitive orientation to VR, not unlike our cognitive orientation to mirrors.

I think it is plausible that after this period of cognitive orientation, a sophisticated user of VR may perceive virtual objects *as virtual*. They will not perceive the objects as present in physical

space, any more than we perceive objects as being on the far side of the mirror. Instead, they perceive the objects as being in virtual space.

Just as visual experience alters for an experienced user of mirrors, I think visual experience may alter for experienced users of VR. When the sophisticated user of mirrors knows they are looking into a mirror, they have a distinctive mirror phenomenology. When the sophisticated user of VR knows they are looking at virtual objects, they have a distinctive *phenomenology of virtuality*.

This is particularly clear in the case where virtual objects are associated with distinctive affordances for action: the ability to pick them up in certain distinctive ways or to walk right through them, for example. A number of philosophers (e.g. Siegel) have argued that affordances are something that we can visually perceive and that are reflected in the character of our visual experience. The affordance to walk through a virtual object might be reflected in a perceived insubstantiality of that object. This is one aspect of the phenomenology of virtuality.

When a sophisticated user has the phenomenology of virtuality, it is plausible that they *perceive* the objects they are interacting with as being virtual objects in virtual space. The interpretation of them as being physical objects in physical space has been left behind, just as the interpretation of mirror objects as being on the far side of the glass has been left behind.

All this suggests that for a sophisticated user of VR, their perceptions of a virtual world need not be illusory. They need not misperceive virtual objects as being in physical space. Instead, they will correctly perceive those objects as being in virtual space.

Of course illusions will still be possible for sophisticated users of VR, just as they are possible for sophisticated users of nonvirtual reality. One way this can happen is when virtual objects cause experiences in some abnormal way. For example, someone might tamper with my headset so that I perceive a treasure as being close to me, when in fact (in the virtual world of the servers) it is a long way away. There can also be invisibility shields in VR that make us see nothing in front of us when something is there, and so on. But none of this leads to the systematic illusoriness posited by the virtual irrealist.

Many hard questions remain. What should we say about augmented or mixed reality? Here a lot depends on whether the user can differentiate virtual and non-virtual objects, and on how much the two sorts of objects interact. If there is differentiation and little interaction, the user may perceive non-virtual objects as being in physical space and virtual objects as being in virtual space (with an associated phenomenology of virtuality). If there is differentiation and much interaction, or if there is no differentiation, the user may perceive all objects as being in an expanded

physical/virtual space. If the physical space is dominant and just occasional undifferentiated virtual objects are inserted, the user may perceive these as being in physical space. That will be an illusion, but it is a special case.

What about the perception of colors? I am inclined to say that the sophisticated user may see objects as having virtual colors, though perhaps this is not as straightforward as the case of perceiving virtual space.

What about proprioception—perception of one's body? This is a particularly hard case for the non-illusion view, as the experience of one's virtual body is at least extremely closely connected to experience of one's physical body, and the latter experience presumably represents a body in physical space. Still, there are cases where one's physical body and one's virtual body have quite different properties (perhaps one has a much longer reach than the other), and one can choose to attend to either the physical and the virtual body, with different resulting experiences. One can make a case that at least for a sophisticated user, attention to the physical body will represent it as being in physical space, while attention to the virtual body will represent it as being in virtual space.

What should we say about language use in virtual reality? When a users sees a virtual table and says "That is a table", are they saying falsely that it is a table, or truly that it is a virtual table? I think once again there is a difference between naive and sophisticated users. The naive user maybe be falsely saying that it is a table. The sophisticated user will certainly be intending to convey that it is a virtual table, and it is reasonably plausible that their utterance of "table" should be interpreted as meaning "virtual table". As they go back between virtual and non-virtual contexts, the meanings will tend to switch quickly and easily. There are various linguistic mechanisms by which this could happen, and I do not mean here to choose between them. But just as cognitive orientation affects what we perceive and what we believe, I think it can plausibly affect what we say and what we mean.

6 Digital Worlds and Fictional Worlds

Let us return for a moment to the question of whether virtual worlds are digital worlds or fictional worlds. Earlier, I allowed that at least some virtual realities involve fictional worlds: a Lord of the Rings videogame, for example. How do I reconcile that with my claim that virtual worlds are digital worlds?

The answer is that in these cases, there is *both* a digital world and a fictional world. In a Second

World War videogame, for example, one sees and interacts with many real digital objects with real virtual colors, located in a real virtual space. At the same time, there is an associated fiction that all this is taking place in Europe in the 1940s, which it is not. This fiction is not absolutely necessary in order to play the game: one *could* treat the game as involving simply virtual objects in virtual space. But most users will deploy the fiction to interpret what is going on in the game, giving a further level of meaning to the game. When one "sees Hitler" in the game, I would say that one actually sees a digital object, but one sees it as Hitler. In effect, there is a digital world (with virtual space) that one interacts with, and a fictional world (with physical space) that one represents.

The digital world has a certain priority over the fictional world, however. When one brings a fictional interpretation to bear, a pre-existing digital world is being interpreted as having a certain fictional content, just as pre-existing physical objects might be interpreted as having fictional content in a non-virtual role-playing game. Furthermore, every VR environment involves a digital world, while only some of them involve an associated fictional world. This strongly suggests that digital worlds are better candidates than fictional worlds to be the basic sort of virtual worlds.

Around here it is useful to distinguish two sorts of fictional content in a virtual world. *Specific* fictional content involves specific physical spatial locations (e.g. Germany), times (e.g. 1945), and individuals (e.g. Hitler). Many videogames involve specific fictional content, and it can also play a role in other uses of VR: say, training and navigation uses of VR where a city such as New York is depicted. However, specific fictional content is quite optional in virtual worlds. Many parts of virtual worlds such as Second Life seem to lack specific fictional content entirely.

A more serious challenge is posed by *generic* fictional content: the representation of objects as occupying physical space and as having shapes, sizes and relative positions, along with other primary and secondary qualities such as colors and perhaps masses and sounds. While specific fictional content is found in only some virtual worlds, it is arguable that generic fictional content can be found in all or almost all VR environments, or at least in those that involve immersive experiences of a three-dimensional environment. Any three-dimensional VR environment *can* be interpreted or imagined as involving objects in physical space, and it will typically be natural to interpret it in this way. Given that in real physical space, there are no objects arranged in this way, it seems that this interpretation of a virtual world must involve fictional content.

At this point, I think one should agree that every virtual reality environment *can* be associated with both a digital world (with virtual space) and a fictional world (with physical space). However, the digital world is always present. The fictional world involving physical space is optional. The invocation of a fictional world depends entirely on the interpretation of the user, and in many cases

that interpretation will not be present at all

In some VR environments, the fictional world involving physical space will be highly salient for most users. These include videogames with specific fictional content (set somewhere on Earth, say) as well as flight simulators and other training programs where simulation of the physical world plays a crucial role in preparing for it. In these cases, while it may be possible in principle to engage with the virtual reality without engaging with the fiction (by taking the attitude that one is in a virtual space, but a physical space), this may be unnatural for most users.

For other VR environments, the fictional world will not be at all salient. For an extreme case, the game of Pong *can* be interpreted as representing a game of tennis in physical space, but few users will interpret it this way. For an intermediate case: for users of Second Life, a fictional interpretation of this world as a physical space may well be set aside in favor of a correct interpretation of the world as a virtual space. Environments that involve unusual forms of embodiment and unusual laws of physics maybe be especially apt for being interpreted as virtual rather than as physical.

Of course naive users of virtual reality are more likely to deploy a fictional interpretation. On one extreme, users who are confronted with a familiar-looking reality that they do not know is virtual will take themselves to be in a physical space, and that space will usually be fictional. But as we have already seen, as users become more experienced, they become cognitively oriented to VR, and the interpretation involving physical space may fall away entirely. For sophisticated users, there need be no sense that they are moving through physical space, interacting with physical objects. Instead, they will take it that they are moving through virtual space, interacting with virtual objects.

So while it is true that any VR environment can be associated with both a digital world and a fictional world, it is also true that every use of a VR environment involves a digital world, while only some involve fictional worlds. So if we take virtual worlds to be something that are associated with every use of a VR environment, and that have a uniform nature, then we should take them to be digital worlds rather than fictional worlds.

What if the conclusions of the previous section are wrong, and every user of virtual reality perceives objects as being in a surrounding physical space? Then at least at the level of perception, every virtual reality will be associated with a world in which physical space is configured the way things look to be, and this world will usually be fictional. We could at this point be dualists about virtual worlds, saying that there are two kinds of virtual worlds: digital worlds and fictional worlds. It is digital worlds that users really interact with, but it is fictional worlds that they perceptually

represent. Of course there would remain some substantive issues: for example, digitalists and fictionalists might still argue substantively about which world is the one that we primarily think about and talk about.

There also remains the abovementioned priority for the digital world that we really interact with. A useful analogy is with philosophical views on which all perception of the physical world involves some sort of illusion. For example, many people hold that physical objects appear to be colored, but they are not really colored (colors exist only in the mind). In that case, we could say that people perceptually represent a fictional world with colors, though they inhabit and interact with a nonfictional world that lacks color. We *could* then be dualists about worlds, but we would certainly say that the real, nonfictional world that we interact with has a certain primacy as the world which we inhabit. By analogy, I think that even if virtual reality involves illusions of a fictional world, the real digital world that we interact with when using virtual reality has a certain primacy as being the reality which we inhabit.

Of course, beyond a certain point, once we agree on all the properties of digital worlds and fictional worlds, arguing over which of these worlds is a "virtual world" is something of a verbal dispute. But however we use the labels, our understanding of virtual reality is improved once we recognize the centrality of real digital worlds in VR.

7 The Value of Virtual Worlds

Are experiences in virtual reality less valuable than experiences outside it? If I climb a virtual mountain, is that less of an accomplishment than climbing a non-virtual mountain? If I win a chess game in VR, does that count for less? If I build a business in Second Life, is that less meaningful than building a business in the non-virtual world? If I fall in love in VR, is the relationship less significant?

Robert Nozick's famous parable of the Experience Machine is often taken to argue that life in virtual reality is much less valuable than life in nonvirtual reality. Nozick introduces the idea as follows:

Suppose there was an experience machine that would give you any experience you desired. Super-duper neuropsychologists could stimulate your brain so that you would think and feel you were writing a great novel, or making a friend, or reading an interesting book. All the time you would be floating in a tank, with electrodes attached to your brain. Should you plug into this machine for life, preprogramming your life experiences? If you are worried about missing out

on desirable experiences, we can suppose that business enterprises have researched thoroughly the lives of many others. You can pick and choose from their large library or smorgasbord of such experiences, selecting your lifes experiences for, say, the next two years. After two years have passed, you would have ten minutes or ten hours out of the tank, to select the experiences of your next two years. Of course, while in the tank you wont know that youre there; youll think its actually happening. Others can also plug in to have the experiences they want, so there is no need to stay unplugged to serve them. Would you plug in? (Nozick 1974, p. 44-45)

Nozick goes on to argue that one should not plug in, for three reasons. First, we want to *do* things, and not just have the experience of doing them. Second, we want to be a certain sort of person, and in the experience machine we are not really any sort of person. Third, the experience machine limits us to a man-made world and rules out contact with a deeper reality.

To respond, it is useful to enumerate various respects in which one might think that the Experience Machine and/or virtual reality are worse than non-virtual reality. I will argue that few of these respects have much force.

Preprogramming. Nozick stipulates that the Experience Machine is entirely preprogrammed: what happens is determined by users and/or programmers in advance. It is arguably this preprogramming that is responsible for much of people's negative reaction to the Experience Machine. A typical reaction is that because it was preprogrammed that one would win the championship (say), one did not really achieve anything in doing so. The preprogramming also seems at least partially responsible for Nozick's worry that one is not doing anything in the Experience machine, and that one is not any sort of person in the Experience Machine.

The Experience Machine is often described as a sort of VR, but preprogramming makes it quite unlike standard VR. Because the experiences in the Experience Machine are entirely preprogrammed, the machine is not interactive: one's actions make no difference to what happends. It is perhaps closer to a passive VR, such as an immersive movie.

It is actually quite obscure how the Experience Machine could be preprogrammed while still maintaining the constraint that the experiences in an experience machine are the same as those outside the machine. Certainly the experience of an immersive movie, even if filmed from an agent's point of view, is typically quite unlike the experience of that agent: there might be similar perceptual experiences, but very little experience of agency and acting. Perhaps Nozick has in mind that one's brain is directly manipulated, and not just a virtual environment, so that one has the full experience of agency; or perhaps the idea is that the brain is analyzed in advance and a virtual environment is constructed in which the brain is guaranteed to do certain things. It is not at

all clear that either of these things are possible, however.

In any case, ordinary interactive VR is not preprogrammed. What one does in the VR makes a big difference to how it evolves. Many videogames have some degree of preprogramming (a series of levels to get through, say, and programmed levels of difficulty), but many other virtual realities (Second Life, say) are largely open ended. No outcomes are guaranteed in Second Life or in World of Warcraft; one has to act in the right way to make them happen. If one attains a certain level in World of Warcraft, it is a real achievement. If one makes a friend in Second Life, it is a real achievement. So even if preprogramming undermines the value of the Experience Machine, it does not undermine the value of VR.

Ignorance. Another way in which the Experience Machine differs from VR is that ordinary users of VR know they are using VR, but Nozick stipulates that his subjects do not know they are using the machine. As a result, these subjects are likely to have false beliefs that they are in non-virtual reality. Certainly ignorance, with the resulting false belief, is a plausible reason to reject the Experience Machine. But it does not apply to a clear-eyed choice to use VR.

Illusions. Some will argue that even knowledgeable use of VR is subject to a related worry: illusions. Nozick says that in the Experience Machine, we seem to do things without really doing them. Some aspect of that worry may be generated by preprogramming and by ignorance, but many will take it to apply to ordinary VR. In using VR, one may seem to be flying, but one is not really flying. I have already argued, though, that knowledgeable users of VR need not be subject to these illusions: they seem to be virtually flying, and they really are virtually flying. If I am right about this, illusion is not a reason to reject life in VR.

Relationships. A common reason to reject the experience machine, and sometimes VR, is that other people will not be present there. The worry may be that no other people are present at all, and at best there will be some simulated people who are not really people. Alternatively, may be that certain specific people who one cares about (family, friends, and partners) will not be present; at best there will be simulated versions of them who are not really them. Neither worries applies to VR in general, however. Familiarly, many VRs have multiple users. This automatically overcomes the first worry: other people are genuinely present. It may leave the second worry, but this worry can be overcome if one's loved ones enter the VR along with one. Any worries here seem broadly analogous to those that apply to moving to another country. They do not bring out a source of disvalue that is distinctive to the virtual.

Interference. Some worry that too much use of VR will interfere with one's non-virtual life. Perhaps it may distract one, so one ignores responsibilities and duties. Perhaps time in VR will

lead one to neglect physical exercise, so one's non-virtual health and one's non-virtual capacities will diminish. Perhaps it may shape one to a VR in a way that is unsuitable for non-virtual reality, as when for example one becomes used to violence in a VR. These are reasonable worries, but most of them apply equally to non-virtual realities, where it is common for one activity to interfere with others (a new relationship may distract one from responsibilities, a desk job may be bad for physical health, a job involving violence may desensitize one). So this problem does not seem distinctive to the virtual domain.

Disembodiment. A common worry is that VR is disembodied. One does not really have a body, and the body is the source of much value in life. Now, I think it is not entirely correct that one lacks a body in VR: one has a virtual body (an avatar), and this virtual body may in principle play many of the same roles as non-virtual bodies. Of course in the present day virtual bodies are much more limited than non-virtual bodies. In some sophisticated virtual realities one can control their movements in considerable detail (for example by moving and tracking one's physical body), and one may also perceive the world from a perspective that depends on one's virtual body. But many bodily functions (food and sex, for example) are either not present or present in an extremely limited way, and most forms of bodily perception are missing from one's relation to one's avatar. Of course one's physical body can supply some of these things, but then one's avatar is not really playing the role of a body, and one is relying on physical rather than virtual reality. Still, there is at least a degree of embodiment in one's avatar, and it is easy to imagine that as the technology becomes more sophisticated, virtual bodies will be able to do everything that physical bodies do (as they do in movies such as *The Matrix*. So while disembodiment is certainly a source of disvalue in current VR, it is probably not an essential and permanent problem that applies to VR across the board.

Quality. A related worry is that VR is of lower-quality than physical reality in various respects. For example, it has lower visual resolution, it has less fine-grained detail, there are fewer modalities of perception, and so on. The obvious reply is as before: while this is a source of disvalue for current VR, it is probably not a permanent or essential problem. It is likely that eventually there will be VR that is largely indistinguishable from ordinary physical reality, in which case it will be just as high-quality. In the long run, it could be that physical reality is low-quality compared to VR.

Artificiality. Nozick objects that the experience machine limits us to human-made reality. This worry applies equally to VR, in which environments are usually human-made, or sometimes computer-made. Either way, the environments are not natural. There is something to this worry,

but it is limited. Some people may value a natural environment, but others will not. Many choose to live in cities in which the environment is largely artificial. It is hard to see that this worry makes VR much worse off than living in a non-virtual, artifically constructed city.

Transience. One may hold that ordinary physical reality has a long past and (we hope) a long future, while VR is typically transient. Many virtual worlds are in effect created at the moment one enters them. Even those that have been around for some time typically have a limited history, perhaps a few years. Furthermore, virtual worlds typically have a limited future, lasting a few years before users move on to new worlds. Many hold that considerable value in the non-virtual world derives from history and from the future. This value is missing in transient virtual worlds. I think there is something to this. As with artificiality, matters such as history are of more value to some people than to others. Further, it may well be that many virtual worlds have an extended future, for example being upgraded with new technology. The lack of history is less straightforward to overcome. Of course one could quickly simulate a long history in a virtual world, but this will be limited in various respects; for example, this will not substitute for a subject's own history in a non-virtual world. Still, it is not obvious that spending time in a world with little history is much worse than spending time in a city or a country with little history.

Birth and death. There is arguably no real birth in VR, and no real death. There might be simulated birth, but no one is really born. There might be simulated death, but no one really dies. When an avatar is destroyed, one can typically "reincarnate" in another avatar. Even if one cannot, one's life will continue, albeit elsewhere. Perhaps a VR device could be arranged to ensure non-virtual death under certain circumstances, and perhaps even non-virtual birth (when two people have virtual sex, their non-virtual genes might be used to create a non-virtual baby that is then reattached into the virtual world), but this in effect is to piggyback virtual birth and death on nonvirtual birth and death. Alternatively, if there are genuine artificial minds in a virtual world, these may undergo birth and death within the world. But if we think there is something special value attaching to our own death (or to human birth), this may have to happen non-virtually. Still, there may be analogs to birth and death, when people enter and leave a VR for example, and lives without birth and death may still have considerable value.

Overall: Some of the worries here (pre-programming, ignorance) do not really apply to ordinary virtual reality. Other worries (illusions, relationships, embodiment, quality) apply to some virtual realities, but do not apply to others. Some (e.g. interference) seem to apply equally in nonvirtual lives. Three of the worries (artificiality, transience, lack of birth and death) are harder to avoid, but at least these sources of disvalue are somewhat limited, and all of them have analogs

in nonvirtual lives.

One can illustrate the situation by a hypothetical nonvirtual reality that we might call terraform reality, or TR. In the future, technology gives us the capacity to terraform new planets very quickly into existence, and to build them up in whatever shape we like. These planets become very popular: they have much more space than Earth, there is an enormous variety of lifestyles, and they allow many more possibilities. Many societies are set up on these planets, and new planets and societies are introduced all the time. People can acquire new bodies in terraform reality, and most choose to do so.

Is terraform reality as valuable as ordinary Earth reality? On the face of it, it has pros and cons, with its value somewhere in the same league. On the plus side, it maybe be more pleasant than earth with many more possibilities. On the minus side, terraformed planets are artificial and they lack much history, so that life on these planets may seem less weight than life on Earth. Still, it would seem perfectly reasonable for many people to choose to spend considerable time in a terraform reality, or even to move there long-term.

I think that life in virtual reality is about as valuable as life in terraform reality. Virtual reality has if anything even more upside, because much more is possible in VR, while having similar downsides. Of course there are better and worse terraformed worlds, and there are better and worse virtual worlds, but they roughly correspond to each other. Perhaps one difference is that TR allows birth and death more straightforwardly than VR. To even out this difference, we can stipulate that people "travel" to TR by leaving their brain on Earth and connecting it (via an instantaneous wormhole connection) to a new body of their choice on TR. Then birth and death will be equally complicated on TR, but though we can suppose various ways they might work analogous to the possibilities for VR.

This yields an argument:

1. Life in virtual reality is roughly as valuable as life in a corresponding terraform reality.

2. Life in terraform reality is roughly as valuable as ordinary non-virtual life.

———————————-

3. Life in virtual reality is roughly as valuable as ordinary non-virtual life.

I think that this conclusion is correct. Life in VR may be in some respects better than a corresponding life outside VR, and in some respects worse. But overall, they are about as good as

each other. Since in many cases we may have the option to enter a virtual world without the option to enter a corresponding nonvirtual world, it may well be that in many cases it is rational to enter a virtual world.

8 Philosophical Underpinnings

What is the underlying philosophical view that leads to this virtual realism? Some philosophers will be led there by idealism, saying roughly that reality is all the mind. If so, then if a virtual object looks and sounds and feels real, then it is automatically real. I am not an idealist, however: I think there is a great deal of non-mental reality outside the mind.

Instead, my philosophical view is a sort of *structuralism*.[7] Physical reality can be characterized by its causal structure: the patterns of interaction between physical objects, and their effects on our experience. Exactly the same goes for virtual reality. Digital objects in general are characterized by their patterns of interaction, which is ultimately a matter of causal structure. Furthermore, the same patterns of causal structure that are present in physical reality can be present in virtual reality. Non-virtual reality and virtual reality are just two different implementations of closely related structures. There may be some differences, but not enough to make on real and valuable whether the other is not.

One does not need to be an structuralist or an idealist to hold hold that virtual reality is real. One simply needs the plausible claim that digital processes on a computer are real, and that virtual reality consists in digital processes on a computer. These claims can be accepted by people with all sorts of metaphysical commitments. I think that at least the first two tenets of virtual realism can be accepted by people with little sympathy for structuralism or idealism.

Structuralism, though, allows not just that digital reality and ordinary physical reality are both real, but that they are on a par. Both can embody the same sorts of structure, and for the structuralist, structure is really what matters. For example, if perception and thought most fundamentally represent the structure of the world, and that structure can be present in virtual reality, it is then

[7] For much more on structuralism and its role in analyzing these matters, see chapter 8 and excursus 17 in *Constructing the World*. Philosophers of science traditionally distinguish two forms of structuralism: ontological structural realism, which holds that all reality is structure, from epistemological structural realism, which holds that all we can know of reality is structure. I think that both are relevant here, with the epistemological version (and also a conceptual version discussed in the *CTW*) playing more crucial role. I should also note that my own structuralism is restricted to physical reality and does not extend to consciousness; but it is the status of virtual reality compared to physical reality that is most central here.

no surprise that perception and thought in virtual reality need not be illusory. And if it is structure (perhaps along with the mind) that gives things value, it is no surprise that virtual reality (along with the mind) can be valuable.

9 Conclusion

One can summarize the position I have argued for by saying: virtual reality is not a second-class reality. Or at least, virtual reality need not be a second-class reality. In the short term, of course, virtual realities may be inferior to physical realities in all sorts of respects (while perhaps being superior in others). But even in the short term, virtual reality may be real, non-illusory, and valuable. In the long term, and in principle, virtual reality may well be on a par with physical reality.

References

Aarseth, E. 2007. Doors and perception: fiction vs simulation in games.

Artaud, A. 1938. *The Theatre and its Double*.

Brey, P. 2003. The social ontology of virtual environments. *American Journal of Ethics and Sociology* 62: 269-82.

Brey, P. 2014. The physical and social reality of virtual worlds. In (M. Grimshaw, ed) *the Oxford Handbook of Virtuality*.

Chalmers, D.J. 1996. *The Conscious Mind*. Oxford University Press.

Chalmers, D.J. 2003. The matrix as metaphysics. Online at thematrix.com. Published in print in (C. Grau, ed) *Philosophers Explore the Matrix* (Oxford University Press, 2005).

Cogburn, J. & Silcox, M. 2008. *Philosophy through Video Games*. Routledge.

Heim, M. 1993. *The Metaphysics of Virtual Reality*.

Heim, M. 2000. *Virtual Realism*. Oxford University Press.

Juul, J. 2005. *Half-Real: Videogames between real rules and fictional worlds*. MIT Press.

Meskin, A. & Robson, J. Fiction and Fictional Worlds in Videogames. In (J. Sageng, ed) *The Philosophy of Computer Games*. Springer.

Nozick, R. 1974. *Anarchy, State, and Utopia*.

Tavinor, J. 2009. *The Art of Videogames*. Blackwell.

Velleman, D. 2013. Virtual selves. In *Foundations for Moral Relativism*. Open Book Publishers.