

The Virtual and the Real¹

David J. Chalmers
New York University

BIBLID [0873-626X (2017) 46; pp. 309–352]

DOI: 10.1515/disp-2017-0009

Abstract

I argue that virtual reality is a sort of genuine reality. In particular, I argue for virtual digitalism, on which virtual objects are real digital objects, and against virtual fictionalism, on which virtual objects are fictional objects. I also argue that perception in virtual reality need not be illusory, and that life in virtual worlds can have roughly the same sort of value as life in non-virtual worlds.

Keywords

Virtual reality, augmented reality, realism, fictionalism, structuralism.

How real is virtual reality? 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. In *Neuromancer*, William Gibson famously said that cyberspace (meaning virtual reality) is a “consensual hallucination”. It is common for people discussing virtual worlds to contrast virtual objects with real objects, as if virtual objects are not truly real.

I will defend the opposite view: virtual reality is a sort of genuine reality, virtual objects are real objects, and what goes on in virtual reality is truly real.

We can get at the issue via a number of questions. (1) Are virtual objects, such as the avatars and tools found in a typical virtual world, real or fictional? (2) Do virtual events, such as a trek through a virtual world, really take place? (3) When we perceive virtual worlds by having immersive experiences of a world surrounding us, are our experiences illusory? (4) Are experiences in a virtual world as valuable

¹ The *Petrus Hispanus Lectures 2016* were delivered by Professor David Chalmers at the University of Lisbon on June 8th and 9th 2016.

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:²

- (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.

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.

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 (which I will not presuppose any knowledge of here), I argued that if we are in a Matrix, most of our ordinary beliefs (e.g. that there are tables) are

² 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.

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 digital (or 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 digital objects.
- (2) Events in virtual worlds are largely digital events that really take place.
- (3) Experiences in virtual reality involve non-illusory perception of a digital world.
- (4) Virtual experiences of a digital world can be about as valuable as non-virtual experiences of a non-digital world.

We might call the combination of (1) and (2) *virtual digitalism*. 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 (1)–(4) even for the temporary and imperfect virtual realities that are possible with current VR technology.

1 Definitions

First, what is virtual reality? In general, the notion of “virtual X” is ambiguous between two readings. On a traditional reading, “virtual X” means something like “as if X but not X” (consider a virtual tie in an opinion poll, which is not exactly a tie, but functions as if it were a tie). On that reading, virtual reality would be an as-if reality that is not reality, and virtual realism would be ruled out by definition. On a more recent and now more common meaning, “virtual X” means something like “a computer-based version of X” (consider a virtual library, which is a computer-based version of a library). That reading is neutral on whether virtual X’s are X’s, and the answer may vary case by case. For example, it is plausible that a virtual kitten in this sense is not a kitten, but a virtual library in this sense is a library. Understanding the term “virtual reality” this way at least leaves it

open that virtual reality is a form of reality.³

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.⁴ Capturing this core, I will say that a virtual reality environment is an immersive, interactive, computer-generated environment. In effect, being computer-generated makes these environments virtual (as on the second definition above), and being immersive and interactive makes our experience of them at least akin to ordinary reality. 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 “presence”: that is, the sense of really being present at that perspective.⁵ Typically this involves inputs that yield 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 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 tools, handheld controllers,

³ The *Oxford English Dictionary* dates the traditional reading of “virtual” (“That is such in essence, potentiality, or effect, although not in form or actuality.”) to 1443 and the more recent reading (“not physically present as such but made by software to appear to be so from the point of view of a program or user”) to 1959. A similar ambiguity is familiar with the expression “artificial intelligence”, where “artificial” can be understood as “as-if” or as “synthetic”.

⁴ 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.

⁵ Slater (2003) suggests that the term “immersive” should be reserved for properties of the technology and the environment, while “presence” is used for the corresponding properties of a user’s subjective experience.

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.

The term "virtual reality" is often used in looser ways than this—sometimes so loose as to include almost any nonstandard means of 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 these conditions.

We can start by removing the three conditions one at a time. *Nonimmersive* VR includes computer-generated interactive 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 sometimes applied to environments satisfying 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 réalité 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 *mixed reality* involves immersive and interactive environments that are partly physical and partly computer-generated. The paradigm case of mixed reality is *augmented reality* where virtual objects are added to an ordinary physical environment. 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 people 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*, which are not strictly virtual realities, nevertheless involve virtual worlds. This usage suggests that 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 definitions 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*.⁶ 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 limited sort of value that engagement with fiction has.

Many of these theorists have focused on the virtual worlds present in videogames, for which fictionalism is an especially natural thesis. For example, there are many videogames based on Tolkein's works and set in Middle 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

⁶ Varieties of virtual fictionalism are expounded by Juul 2005, Tavinor 2009, Bateman 2011, Velleman 2011, and Meskin and Robson 2012. To be fair, many of these theorists are making claims about videogame worlds rather than about virtual worlds more generally. 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.

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 this is unlikely to stay the case indefinitely. 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 socializing, 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 that 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. To a first approximation, they can be regarded as data structures, which are grounded in computational processes which are themselves grounded in physical processes on one or more computers. To a second approximation, one may want to invoke more subtle relations between virtual objects and data structures, just as theorists often invoke more subtle relations between high-level nonvirtual objects (e.g. a statue) and underlying physical entities (e.g. a lump of clay). For example, in some cases, multiple data structures will be associated with a single virtual object, in which case the virtual object will be a higher-level entity constituted by these data structures. I will focus on the simple data structure view here, but much of what I say should generalize to more complex views.

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 mo-

ment'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 our 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 photo or 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.

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 or that lightning is electric discharge. One cannot tell that water is 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.

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.

To flesh out this answer, we need to pay more attention to what it is for a real object to have virtual properties, such as being a virtual dragon. Questions of this sort are the subject of the next section.

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 the air. In the real world, the correspond-

ing digital object does not fly through the air. 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 understand colors in virtue of their functional (or causal) role.⁷

⁷ One refined view holds that redness is the physical property that normally

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, 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 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

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.

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.⁸

(Of course if we are actually 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.)

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 subtleties 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

⁸ What if different users use different headsets generating different color experiences? Here the issues parallel familiar issues about variation in nonvirtual vision. For example, if a data structure normally generates reddish experiences, but some users use a black-and-white headset for which the same structure generates grayish experiences, then one can reasonably say that the data structure is virtually red and not virtually gray, just as apples are red even though colorblind people see them as gray. What if the same data structure is used with two quite different (and widely used) VR headsets, normally causing reddish experiences in one and normally causing greenish experiences in the other? Our answer then may depend on which headset is normal for us: if the first headset is normal, we may say that the object is red, while if the second is normal, we may say the object is green. If both are equally normal, we may say that there is no absolute fact of the matter: the object is red relative to the first headset and green relative to the second.

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 allows 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 will suffice 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 by virtual entities' 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*. In cases such as those above, when *X* is picked out as what normally plays a certain functional role, then virtual *X* will typically be distinct from *X*. Virtual *X* will be a digital property that normally plays the role in virtual environments, while *X* will be a nondigital property that normally plays the role in non-virtual environments. In other cases, however, virtual *X* may work quite differently.

We saw earlier that there are some *X* for which a virtual *X* is an *X*: 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.

When exactly is a virtual X an X? In *The Conscious Mind* (1996), I answered a similar question—when is a simulated X an X?—by saying this holds when being an X is an *organizationally invariant* property: one that depends only on the abstract causal organization of the underlying system. Simulations are typically designed to replicate the abstract causal organization of an original system. A property such as being a calculator depends only on this organization, which is also present in a simulation, so a simulated calculator is a calculator. The same reasoning explains why a virtual calculator is a calculator.

One difference between ordinary virtual reality and simulations is that in VR, we typically assume that users having genuine experiences and mental states are present. If we assume that all relevant 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 situation in principle can replicate both the abstract causal organization and the mental properties of a nonvirtual situation. This suggests that a virtual X will be an X as long as X is a causal/mental invariant: one that depends only on the abstract causal organization and the mental properties of a situation. For example, it is plausible that being an action and being a philosopher can be analyzed in causal and mental terms. Correspondingly, I think that virtual actions are actions, and virtual philosophers are philosophers, at least given that all the relevant mental states are present.⁹

⁹ 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 (such as money), one that is constituted by collective social agreements in the right way. I think that the “only if” claim is not quite right: virtual calculators are calculators and virtual boredom is boredom, 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.

Sometimes when people ask the question “are virtual objects real?”, they are asking questions of the form “Are virtual X’s really X’s?”.¹⁰ To that question, the right answer is sometimes yes, and sometimes no, depending on the X. Virtual kittens are not really kittens, but virtual libraries are really libraries. But importantly, virtual kittens are still real objects. Virtual and nonvirtual kittens have a different underlying composition, but virtual kittens at least in principle can be just as rich and robust as nonvirtual kittens and play corresponding causal roles in virtual worlds.

5 Is perception of virtual reality illusory?

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

Correspondingly, an opponent might accept everything I have said so far, while holding that virtual worlds are nevertheless illusory. The reason is that we undergo illusions when we perceive virtual objects, because we perceive them as nonvirtual. When I perceive a red cube in VR, it appears to me as if I am seeing a nonvirtual object that is nonvirtually red and nonvirtually cubical. But I am not. So my experience is illusory.¹¹

We might call this view *virtual illusionism*. We can formulate a simple argument for virtual illusionism as follows:

- (1) We perceive virtual objects as having the ordinary (non-virtual) colors, locations, and shapes that a corresponding non-virtual object has.

¹⁰ In comments on this paper, Cheryl Chen endorsed Austin’s suggestion that when we ask whether something is real, we always need a “trouser-word” X to ask whether it is a real X. I think we can ask simply whether something is real (to ask whether Santa Claus is real we need not ask whether he is a real man), but if a trouser-word is desired, “real object” or “real entity” will do.

¹¹ A relative of this line is taken by Slater (2009), who says that the sense of presence in a virtual environment always involves a “place illusion”: we perceive ourselves as being located in a “real place” when we are not.

- (2) Virtual objects do not have the ordinary (non-virtual) colors, locations, and shapes that a corresponding nonvirtual object has.
- (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? A perceptual illusion is a case where an object looks a certain way, when it is not that way. For example in the Müller-Lyer illusion, one line looks longer than the other, even though it is not. Perceptual illusions like this often persist even when the subject is not deceived at the level of belief. Even after we know the lines have the same length, we continue to perceive one has longer than the other.

Might there be perceptual illusions like this in virtual reality? Naive users can certainly 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 at least for some time. It is tempting to say that the illusion is hard-wired and will persist indefinitely, even for sophisticated users. I think this temptation should be resisted, however.

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 anything about mirrors will plausibly experience an illusion of this sort. Even an experienced user can experience such an illusion when they do not know that a mirror is present, as for example when one walks into a restaurant with mirrors on the walls and initially has an experience of a bigger space. But does the illusion persist in every case of mirror vision, including cases involving experienced users who know that a mirror is present?

The view we might call *mirror illusionism* says that mirror vision always or at least typically involves perceiving objects as being on the far side of the glass, and thereby always or at least typically involves a spatial illusion.¹² The view we might call *mirror veridicalism* says that mirror vision (at least for experienced users who know that a mirror is present) typically involves perceiving objects as being on the near side of the glass, and so need not involve a spatial illusion. I think that mirror veridicalism is the correct view, as reflection on cases of mirror use by experienced users brings out.

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, do the cars look to be in front of the driver, pointing toward her? Or do they appear to be behind the driver, following behind her vehicle? My own intuition, and those of most people I have asked about this case, is clear. When I look in my rear-view

¹² Maarten Steenhagen (2017) independently formulates mirror illusionism as what he calls "specular illusionism", and argues against it using arguments distinct from those given here. Roberto Casati (2012) and Clare Mac Cumhaill (2011) discuss some related questions about mirror perception.

mirror, the cars I see look to be behind me.

Now, someone who thinks that rear-view mirrors are 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, and indeed from ordinary mirror perception. In the mirror box case, for example, one has visual experience as of an arm on the other side of the glass—the arms visually *seems* to be on the other side of the glass, even though one knows it is not. With a rear-view mirror, and with typical mirror perception, one has visual experience as of objects on the near side of the glass, and the objects visually seem to be on this 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-dependence*: 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 abrupt 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, but the mirror case is one of the more plausible examples.¹³ 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 chair 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 *Fantastic Voyage*, then an expert

¹³ See Macpherson 2012 and Firestone and Scholl 2016. Firestone and Scholl are most concerned to argue against the cognitive penetrability of relatively early vision, a module whose products may differ from the (further downstream) contents of visual experience, so their view is consistent with effects of mirror knowledge on perceptual experience, as long as these effects are relatively late. Correspondingly, my own view is consistent with there being early levels of representation in which objects seen in mirror are represented as being on the far side of the glass.

user will not undergo the illusion that the objects are much larger than they are: instead they will correctly perceive small objects as small.

What applies to video also applies to (nonvirtual, camera-based) video images viewed through an immersive headset. With appropriate knowledge and expectations about how the headset works, one could experience objects in the center of one's visual field as being behind one, or as being in front of a remote robot, or as being extremely 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, the perceptual illusion may persist for a period, but they will not be fooled into believing that the objects are present. After some time, 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. And this perception will be correct.

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 2014) 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. I think that the phenomenology of virtuality goes beyond the perception of these affordances, though, and can be present even in a virtual world that presents relatively normal affordances.

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. Furthermore, the sense of “presence” need not involve suffering an illusion that one is in a nonexistent physical location. Rather, it may involve the correct perception that one is in a virtual location.

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 about the perception of colors in virtual reality? 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, as in the so-called "very long arm illusion" of Kilteni et al 2012), and one can choose to attend to either the physical or the virtual body, with different resulting experiences. One can make a case that at least for a sophisticated user, proprioceptive attention to the physical body will represent it as being in physical space, while proprioceptive attention to the virtual body will represent it as being in virtual space.

VR researchers (e.g. Blanke and Metzinger 2009; Maselli and Slater 2013) often talk of "bodily ownership illusions" and "full-body illusions", where users are given the illusion of bodily ownership with respect to a virtual body. If I am right, this phenomenon need not always be an illusion (though certainly some experimental conditions involve illusions). It is possible to have a sort of bodily ownership of a virtual body (especially when one has appropriate perception of and control over that body), and sophisticated users who experience a virtual body as their virtual body may be correct. Likewise, the so-called "virtual arm illusion" (Slater et al 2008), where users experience a virtual arm as theirs, need not always be an illusion. Even the "very long arm illusion" may sometimes involve non-illusory perception of a virtual long arm.

Finally, what should we say about language use in virtual reality? When a user 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 may 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 virtual environment (including *Second Life* and the like) *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 not 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. A residual issue concerns which world is the one that we primarily think about and talk about: when I form beliefs about an avatar, or talk about it, am I thinking or talking about a digital object or a fictional object? I would argue that the digital object has primacy here, at least for experienced users: at least in thought (if not in perception), these users are not inclined to regard virtual worlds as physical worlds, even fictional ones. Rather, they form true beliefs about digital virtual worlds.

Even if we accept this dualism about virtual worlds, the digital world will as before have a certain primacy as the one 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?

A virtual fictionalist may hold that life in VR has only the limited sort of value that one gets from engaging with fiction. Engaging with VR is analogous to engaging with a good book or a good movie, which may have considerable value, but only a limited sort of value compared to the full range of value available in nonvirtual life. By contrast, a virtual realist like me holds that life in VR can in principle have most or all of the value deriving from nonvirtual life.

I have already argued against virtual fictionalism, and in doing so have rebutted that sort of objection to the value of VR. Here I will address some other objections to the value of VR.

I will focus initially on Robert Nozick’s parable of the Experience Machine, which 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 life's 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 won't know that you're there; you'll think it's 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: 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 human-made world and rules out contact with a deeper reality.

I think that whether or not these are good reasons not to plug into the Experience Machine, they are not good reasons not to use virtual reality. If we understand virtual reality properly, the first two objections do not apply to it, and while the third objection may apply, it does not have much force.

To start with the third objection: it is true that virtual environments are usually human-made. But if this is an objection to living in virtual reality, it is also an objection to living in a modern city such as New York. But billions of people lead meaningful lives in human-made environments such as cities. It is certainly reasonable to value naturalness in an environment, but this seems an optional value and for most people, not the sort of thing that makes the difference between a meaningful life and a meaningless one.

The first two objections are potentially more serious, but they

do not apply to virtual reality. In virtual reality environments, users make real choices, they really do things, and they are genuine sorts of people. Even in limited existing environments such as *Second Life*, a user can genuinely write a novel, or make a friend, or read a book (to use Nozick's examples). They can choose whether to (virtually) attend a concert or to build a house. They can be honest or dishonest, and shy or courageous. In principle, a subject living in a long-term virtual reality could make their own life there.

Nozick's first two objections seem to stem from two features of the Experience Machine as he conceives of it. First, it is illusory, so that what seems to happen in the Experience Machine does not really happen. Second, it is preprogrammed, so that one's life experiences are programmed in advance rather than depending on one's choices along the way. Whether or not these features are true of the Experience Machine, however, I think they are not true of virtual reality.

I have already discussed illusoriness, and I have argued that virtual reality need not be illusory. At least for sophisticated users of VR, what seems to happen in VR by and large really happens. One may seem to have a conversation, and one really does have a conversation. One may seem to enter a virtual house, and one really does enter a virtual house. One may seem to be virtually flying, and one really is virtually flying. I would add that virtual actions are plausibly real actions (albeit with a virtual body), so that when one performs virtual actions, one really is doing something.

A related worry for the Experience Machine is ignorance. In the Experience Machine, users presumably often will not know that they are using an Experience Machine, so there is much more scope for them to have false beliefs as well as perceptual illusions. But none of this is an objection to use of standard virtual realities, where users know perfectly well that they are inhabiting virtual worlds.

As for preprogramming, Nozick's description stipulates that the Experience Machine is entirely preprogrammed: what happens is determined by users and/or programmers in advance. However, most virtual reality is not preprogrammed. Our definition of core VR requires that it is interactive: a user's actions make a difference to what happens. In some videogames the range of actions is admittedly limited, but even here there is still some room for free choice. In an open-ended world such as *Second Life*, what happens depends

almost entirely on the free choice of users. Perhaps there are some constraints on actions, but there are also constraints in the non-virtual world. One's life in a virtual world certainly need not be preprogrammed.

It is arguably 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, and perhaps one was not genuinely doing anything at all. But even in *Second Life* or *World of Warcraft*, no outcomes are guaranteed; 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 this objection to the Experience Machine has little purchase on VR.

The Experience Machine is often described as a sort of VR, but as described it is not interactive: one's actions make no difference to what happens. It is perhaps closer to a passive VR, such as an immersive movie. Because of this, it is 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 is possible, however.

So we have seen that Nozick's three objections to the Experience Machine are not strong objections to living in VR: people can do things in VR, they can be genuine sorts of people in VR, and while virtual environments are artificial, this does not stop people from living a meaningful life there.¹⁴

¹⁴ Cogburn and Silcox (2014) also give a detailed critique of Nozick's three objections to the Experience Machine as applied to VR. Their responses are

Are there other reasons to think that life in VR must lack the sort of value that life in nonvirtual reality has? A few other potential worries include the following.

Relationships. Some reject the experience machine on the grounds that other people, or perhaps specific people who they love, will not really be present there. Does this worry apply to VR? Certainly there are some virtual realities in which other people are not present: perhaps at most there will be a few “non-player characters” that are not genuine people at the current level of technology. But many VRs have multiple users who are people, and it is perfectly possible to enter a VR with specific loved ones. So there is not a principled objection to VR here.

Interference. Some worry that VR will interfere with one’s non-virtual life. Perhaps it will distract one from responsibilities and duties; perhaps one will neglect one’s nonvirtual health; perhaps violence in VR will make one more violent outside 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 in VR one lacks a body, which is the source of much value in life. One can have a virtual body in VR, but at least for now these are much more limited than nonvirtual bodies, and lack many of their functions. Eating, drinking, exercising, and having sex, for example, are either impossible or at least extremely limited in current VR. One’s physical body can supply some of these things, but then 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

somewhat different to mine, but they also emphasize the interactiveness of VR in contrast to the passiveness of the Experience Machine. Cogburn and Silcox oppose “brain-in-a-vat-ism” about VR, which holds that life in VR has only the limited value and epistemic status of the life of a brain in a vat. By my lights (Chalmers 2003) they are too pessimistic about brains in vats: at least some brains in vats can have good lives and plenty of knowledge.

disembodiment is certainly a source of disvalue in current VR, it is probably not an essential and permanent problem.

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. However, this is likely to be another short-term problem. Eventually there may be VR that is largely indistinguishable from ordinary physical reality, and in the long-term VR could be much higher quality in various respects.

Transience. It is arguable that much value in ordinary reality derives from its long past and (we hope) its long future. Virtual worlds typically lack a long past and a long future. Many virtual worlds are in effect created at the moment one enters them and disappear when one leaves them. Even older virtual worlds are typically only a few years old, and most have a limited future. Perhaps some virtual worlds may come to have indefinite futures (with occasional upgrades in technology), and perhaps they could even have a long history if that history is simulated fast enough, but there will be some limits here. This simulated history will not be a substitute for a subject's own history in a non-virtual world. Still, as in the case of artificiality, the worry here is limited. Plenty of people live meaningful lives in environments with little relevant 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 attached to 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 some 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 (quality, disembodiment) are worries for current VR technology but not principled worries in the long term. They do not apply to what we can call *rich virtual reality*: a future level of VR with the complexity of ordinary physical environments, where these temporary technological limitations have been overcome. Other worries (relationships, interference) seem to apply equally to moves between different parts of a nonvirtual world. Three of the worries (artificiality, transience, lack of birth and death) are harder to avoid. These bring out sources of value (naturalness, history, birth and death) that are hard to replicate in virtual worlds, at least without going to a full-scale Matrix-style simulations. But the disvalue that results from their absence is somewhat limited, and in each case it has 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, but it is at least roughly as valuable overall. On the plus side, it may 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 weighty 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 can be about as valuable as life in terraform reality. Existing limited forms of VR may have somewhat limited value, but life in rich VR (VR with roughly the complexity of ordinary reality, in which short-term technological limitations have been overcome) will be about as valuable as life in a corresponding terraform reality. Virtual reality has if anything even more upside, because much more is possible in VR. As for downside, perhaps one

difference is that TR allows birth and death more straightforwardly than VR, though perhaps rich VR in the long-term future may overcome this problem. (Alternatively, we could stipulate a form of TR where the brain stays on earth, thereby making birth and death as difficult as in VR.)

All this can be put as an argument:

- (1) Life in rich VR 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 rich VR is roughly as valuable as ordinary non-virtual life.

I think that this conclusion is correct. Life in a rich 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 Other realities

It is natural to ask how much of what I have said generalizes to other “realities” that are like virtual realities at least in some respects: mixed realities, dreams, delusions, fictions, and other cases,

Mixed reality. A mixed reality is an environment that is partly computer-generated and partly not.¹⁵ The best-known case here is that of augmented reality, where VR technology is used to add virtual objects to standard perception of a physical world. There are also cases of so-called “augmented virtuality”, where physical objects are added to our perception of a virtual world, and other sorts of mixing.

I would say that virtual objects in mixed realities have the same

¹⁵ The term “mixed reality” was introduced by Milgram et al (1994), who also discuss many sorts of mixed reality along the “reality-virtuality continuum”.

ontological status as those in pure virtual realities: they are digital objects and are perfectly real. When we interact with a mixed reality, we interact with both nonvirtual (physical) objects and virtual (digital) objects, and these will often interact with each other. Things get more complicated when we think about how we perceive and represent mixed realities, and what properties we attribute to objects within them: for example, our perception of their location in space. Much depends on how the mixed reality is set up.

Two relevant aspects of the set-up are whether virtual objects are distinguishable from nonvirtual objects, and whether and how the two sorts of object interact. If virtual objects are distinguishable (as is now usually the case), users will plausibly perceive them as virtual, with an associated phenomenology of virtuality. If virtual objects interact with each other but not with nonvirtual objects, the user may perceive non-virtual objects as being in physical space and virtual objects as being in a separate virtual space. If the objects are distinguishable but have rich interactions, the users will plausibly perceive them as inhabiting a single space. Perhaps this might initially be a physical space (at least for an augmented reality where most objects perceived are physical), with some associated illusions for virtual objects: they seem to be in physical space, but they are not. However, a sophisticated user might eventually come to represent the objects as being located in a disjunctive “mixed space”, or perhaps to represent virtual objects as being “virtually located” in physical space, where virtual location is a relation that virtual objects really bear to the space (a virtual book might be virtually located on my desk, even though it is not physically on my desk). If so, perception of virtual objects need not involve illusions.

What about mixed realities where the virtual objects are indistinguishable from nonvirtual objects? I suspect that these mixed realities will be rare for the foreseeable future, partly because of technological limitations, but also partly because there will be strong pressure from users to be able to distinguish the virtual from the nonvirtual. Many users may come to rely on the phenomenology of virtuality. But for an environment where the two are indistinguishable and interact, then as above it is plausible that users will represent them as being in a single space. As before, naive users may have the illusory experience that the virtual objects are in physical space.

Sophisticated users of these realities may come to represent the objects they perceive as being located in a mixed space, or perhaps as being quasi-located in a physical space, where quasi-location corresponds to location for physical objects and to virtual location for virtual objects. If so, it need not involve illusion.

Another issue arising from mixed reality concerns “mixed” objects—objects that are partly virtual and partly based on real objects. Even in standard VR it is common to have images of virtual handheld controllers that track the location and orientation of non-virtual controllers with the same shape. In augmented reality, a physical object can be “transformed” into a virtual object based on that object. In these cases, is one seeing a nonvirtual object, a virtual object, or both? Does one see it as inhabiting a virtual space, a non-virtual space, or both? The matter is complex and depends on the case, but I am inclined to think that in many cases, both a virtual and a nonvirtual object is present, and one can reasonably be said to be seeing both of them.

Dreams. Ordinary dreams are far more fragmented and far less stable than ordinary virtual or nonvirtual realities. At least in principle, however, there could presumably be rich and stable lucid dreams governed by regular laws. A dream like this would be structurally analogous to a virtual reality, except that a part of one’s brain is playing the role of the computer in implementing the reality. Events in the dream world could at least in principle be identified with brain events. We could run a version of the earlier argument from perception: dream events are the objects of perception in dreams, but brain events cause our experiences and therefore are the objects of perception in dreams, so dream events are brain events. This would yield a sort of “dream realism”, at least for these rich and stable dreams.

This dream realism is perhaps more counterintuitive than virtual realism, and might be resisted by pointing to key differences between dreams and virtual reality. One difference is that unlike virtual worlds, dream worlds depend on our own minds. Mind-independence is often invoked as a criterion for “reality”, so this might yield one sense in which stable dream worlds are less real than virtual worlds. Furthermore, one could argue that in ordinary dreams we represent objects as mind-independent, and mind-independence is required to be an object of perception. If so, brain events may not

qualify as dream events after all. In addition, mind-dependence may entail that experience in dream worlds is less valuable than nonvirtual experience, since we value engaging with a world outside our minds.

Are dreams illusions? In ordinary dreams, we do not know we are dreaming. As in the case of unknowing use of VR, it is plausible that we represent dream events as located in ordinary physical space around us. Those events are usually not taking place there, so ordinary dream experience will be either illusory, if dream events are really brain events, or hallucinatory, if dream events are not brain events. Either way, dream events of the sort that we think are taking place (people walking and talking in the physical world, for example) will typically be happening only in a fictional world. In lucid dreams, however, we know that we are dreaming. As with VR, a sophisticated lucid dreamer might come to represent dream events as taking place in a self-generated dream world, so that their thought and perception need not be illusory.

Delusions and hallucinations. What about the worlds represented in cases of mental illness, such as the delusions and hallucinations associated with schizophrenia? These are at least somewhat analogous to the dream case, with the difference that ordinary physical reality plays a much larger role. Some delusions involve just mild alterations to ordinary physical reality. Even in extreme hallucinations, perceived physical reality typically plays a role. So these cases perhaps stand to dreams as augmented and mixed realities stand to virtual realities, with some perceptual experiences caused by the external world and some by internal brain processes. In cases where subjects do not know they are suffering from delusions or hallucinations, they will plausibly undergo illusions and false beliefs, representing things as going on in the external physical world where they are not. For sophisticated subjects who can distinguish their “delusion” experiences from other experiences, things are not so clear. In some cases they may come to experience relevant events as happening in an alternative virtual space. As in the case of dreams, this space will be mind-dependent, and it may be more fragmented and less stable than ordinary physical and virtual realities, but as in the case of dreams there may be a case for identifying perceived events with brain events all the same.

Fictions. What about events that take place in the worlds described in novels. Do these really take place, perhaps in the head of the author or the reader? I think not. The head of the reader is much too limited to properly ground the events of a fictional world. The head of an author may be richer, involving a detailed model of the fictional world, and certain components of these models may be causally responsible for a reader's experience of fictional events. But there remain strong obstacles to identifying fictional events with brain events. First, the brain events will typically not stand in the causal and counterfactual relations that the fictional events stand in. In the fictional world it may be true that if John had not gotten angry, no one would have been killed, but corresponding claims about brain events may be false; the novelist wanted to write a murder mystery, so if John had not gotten angry, someone would have gotten killed anyway. To have something closer to the virtual reality case, with the right cause and counterfactuals, we would need a novelist who writes a novel by setting up a simulation and some rules and having it all unfold in their head. But this would be an extremely unusual case.

Interactive novels and text adventures. Another difference between novels and virtual realities is that novels are typically not interactive, and even interactive novels do not have anything like the rich degree of interaction typical of virtual and ordinary reality. That said, a sufficiently interactive novel would in effect be something like a text adventure game, such as Colossal Cave Adventure, in which users receive text descriptions of their location in a virtual world and issue text commands to see what comes next. I am inclined to think that someone playing Colossal Cave is genuinely interacting with real virtual events in a digitally realized virtual world (as well as representing an associated fictional physical world). This virtual world is simpler than the worlds presented in videogames and is presented via text rather than through perception, but otherwise is similar in kind. Something similar goes for the virtual worlds realized in games such as Dungeons and Dragons, which are traditionally in the notebooks, props, and memories of participants.

9 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 in the mind, so that if we have rich enough perceptions as of a world around us, that world is real. 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.¹⁶

Instead, my philosophical view is a sort of *structuralism*.¹⁷ 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. To take an extreme case, a virtual simulation of the entire physical world will replicate the causal structure of the physical world, with causal relations between physical entities mirrored by causal relations between corresponding digital entities. Non-virtual reality and virtual reality are just two different implementations of closely related structures. There may be some differences, but these differences are not enough to make one real and valuable while the other is not.

¹⁶ I do have some sympathy with versions of panpsychism on which all reality is mental. This is a sort of idealism, but a different sort from the one discussed in the text where reality is grounded in perception of it: on this view there will still be a great deal of reality outside *my* mind.

¹⁷ For much more on structuralism and its role in analyzing these matters, see chapter 8 and excursus 17 in *Constructing the World*, and my recent article “Structuralism as a Response to Skepticism”. Philosophers of science traditionally distinguish two forms of structuralism: ontic structural realism, which holds that all reality is structure, and epistemic structural realism, which holds that all we can know of reality is structure. The epistemic version (and especially a conceptual version holding that our relevant concepts of physical reality are structural) is most relevant to the current discussion, which can stay neutral on the truth of ontological structuralism. 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.

One does not need to be a structuralist or an idealist to 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 in important respects 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 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.

10 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. It may be a second-level reality, in that it is contained within physical reality and realized by processes in the physical world, but this need not make it less real or less valuable. In the short term, of course, virtual realities may be inferior to physical realities in all sorts of respects (while perhaps beginning to be superior in other respects). 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.¹⁸

David J. Chalmers
Department of Philosophy, New York University

¹⁸ Thanks to audiences at Arizona, the Australasian Association of Philosophy, Brooklyn Library, Glasgow, Harvard, Hay-on-Wye, La Guardia College, NYU, Skidmore, Stanford, and Sun Valley, and especially at the Petrus Hispanus Lectures at the University of Lisbon. For comments on drafts of this paper, thanks to Cheryl Chen, Grace Helton, Frank Lantz, Eric Schwitzgebel, and David Yates.

5 Washington Place
New York NY 10003, USA
chalmers@nyu.edu

References

- Aarseth, E. 2007. Doors and perception: fiction vs simulation in games. *Intermedialities* 9:35–44.
- Artaud, A. 1938. *The Theatre and its Double*. Gallimard.
- Bateman, C. 2011. *Imaginary Worlds*. Zero Books.
- Blanke, O. and Metzinger, T. 2008. Full-body illusions and minimal phenomenal selfhood. *Trends in Cognitive Sciences* 13: 7–13.
- 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 *The Oxford Handbook of Virtuality*, ed. by M. Grimshaw.
- Casati, R. 2012. Illusions and epistemic innocence. In *Perceptual Illusion: Philosophical and Psychological Essays*, ed. by C. Calabi. Palgrave MacMillan.
- 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 *Philosophers Explore the Matrix*, ed. by C. Grau. Oxford University Press, 2005.
- Cogburn, J. and Silcox, M. 2008. *Philosophy through Video Games*. Routledge.
- Cogburn, J. and Silcox, M. 2014. Against brain-in-a-vatism: On the value of virtual reality. *Philosophy and Technology* 27:561–79.
- Firestone, C. and Scholl, B.J. 2016. Cognition does not affect perception: Evaluating the evidence for ‘top-down’ effects. *Behavioral and Brain Sciences* 39:1–72.
- Heim, M. 1993. *The Metaphysics of Virtual Reality*. Oxford University Press.
- Heim, M. 2000. *Virtual Realism*. Oxford University Press.
- Juul, J. 2005. *Half-Real: Videogames Between Real Rules and Fictional Worlds*. MIT Press.
- Mac Cumhaill, Clare. 2011. Specular space. *Proceedings of the Aristotelian Society* 111: 487–95.
- Macpherson, F. 2012. Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research* 84:24–62.
- Maselli, A. and Slater, M. 2013. The building blocks of the full body ownership illusion. *Frontiers of Human Neuroscience* 7: 83.
- Meskin, A. and Robson, J. 2012. Fiction and fictional worlds in videogames. In *The Philosophy of Computer Games*, ed. by J. Sageng. Springer.
- Milgram, P., Takemura, H., Utsumi, A., and Kishino, F. 1994. Augmented Reality: A class of displays on the reality-virtuality continuum. *SPIE Conference on Telemanipulator and Telepresence Technologies* 2351:282–92.

- Nozick, R. 1974. *Anarchy, State, and Utopia*. Basic Books.
- Siegel, S. 2014. Affordances and the content of perception. In *Does Perception Have Content?*, ed. by B. Brogaard. Oxford University Press.
- Slater, M. 2003. A note on presence terminology. *Presence Connect* 3(3).
- Slater, M, Perez-Marcos, D., Ehrsson, H. and Sanchez-Vives M.V. 2008. Towards a digital body: the virtual arm illusion. *Frontiers of Human Neuroscience* 2:6.
- Slater, M. 2009. Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical Transactions of the Royal Society of London B* 364: 3549–57.
- Steenhagen, M. 2017. False reflections. *Philosophical Studies* 5: 1227–42.
- Tavinor, J. 2009. *The Art of Videogames*. Blackwell.
- Velleman, D. 2013. Virtual selves. In *Foundations for Moral Relativism*. Open Book Publishers.