The Social Reality of Virtual Worlds

METAPHYSICS

RESEARCH

ubiquity press

ROBERT FRASER

ABSTRACT

What is the ontological status of virtual worlds? The two prominent positions in the recent debate are David Chalmers's *virtual digitalism* and Neil McDonnell and Nathan Wildman's *virtual fictionalism*. In this paper, I argue that there are good reasons to be dissatisfied with both. To overcome their limitations, I propose a novel position, *virtual socialism*. Drawing on the 'two-dimensional' approach to social ontology articulated by Brian Epstein, I suggest that virtual objects are social objects grounded in the states of a computer, but 'anchored' by a variety of social and non-social factors. Virtual socialism, I suggest, makes the best sense of the messy relationship virtual reality bears to digital reality, as well as the fact that virtual reality can sometimes be inconsistent.

CORRESPONDING AUTHOR:
Robert Fraser

Swansea University, UK

r.f.knowles@swansea.ac.uk

KEYWORDS:
virtual reality; realism; fictionalism; social ontology

TO CITE THIS ARTICLE:
Fraser, R. 2024. The Social Reality of Virtual Worlds. *Metaphysics*, 7(1), pp. 85–98. DOI: https://doi.org/10.5334/met.156

DOI: 10.5334/met.156

1. INTRODUCTION

The ontology of virtual worlds is gathering interest. Not only is it an instructive arena in which to explore core metaphysical issues; it also has great social importance. Our judgements about the ontological status of virtual objects and events guide us in determining whether to lend them the same moral or legal status as their real-life counterparts. This is evident in how law-enforcement organisations have dealt with recent cases of virtual theft (see Wolswijk 2012) and virtual sexual assault (see Vallance 2024). The ontology of virtual worlds is therefore a pressing concern.

Two positions have emerged in recent debate. David Chalmers (2017, 2019) defends *virtual digitalism*, on which virtual objects are identical to digital objects, and enjoy the same kind of independent existence as physical objects. Neil McDonnell and Nathan Wildman (2019) defend *virtual fictionalism*, according to which virtual objects are fictional objects we imagine to exist when engaging with videogames and VR systems as representational artworks.

In what follows, I argue that we have reason to be dissatisfied with both, since they fail to account for the messy relationship between virtual and digital reality (§§2–3). In their place, I develop a novel position, *virtual socialism* (§4). Drawing on the 'two-dimensional' approach to social ontology articulated by Brian Epstein (2015, 2016), I argue that virtual objects are grounded in the states of computers, but 'anchored' by a variety of social and non-social factors. This, I argue, makes the best sense of the messy relationship between virtual and digital reality, and the fact that virtual reality can sometimes be inconsistent.

2. VIRTUAL DIGITALISM

To understand virtual digitalism, we must first understand what digital objects are. For Chalmers (2019), the simplest kind of digital object is a *bit*, a particular physical realisation of a 1 or a 0 on a computer. This will typically be a transistor on a circuit board having either a high or a low voltage. More complex digital objects—*data structures*—are arrays of bits. Data structures are the primary means by which information is stored on a computer.

Data structures exist. They have particular (albeit scattered) locations and causal powers. By changing the values of a particular data structure, we can make predictable and systematic changes in the functions a computer performs. A case for their existence can also be made by appealing to the explanatory role they play in computer science (cf. McDonnell & Wildman 2019: 388, fn. 19).

Once we accept data structures, there seems little reason to deny that 'high-level' data structures exist. Consider, for example, triples of numerical values, encoding the name, rank, and serial number of certain objects. Chalmers suggests that

> in general it will be built into the design of a system that these high-level structures are implemented (after compilation and/or interpretation) in certain low-level structures involving bits, which are themselves physically realised when a program is executed. (Chalmers 2019: 462–467)

Thus, we can say that these high-level data structures exist and are grounded in low-level data structures.

We can even allow, in principle, that which particular low-level data structures realise a given high-level data structure can change over time. For example, in the course of executing a program, different parts of the computer hardware may realise the same data at different times. Here, Chalmers relies on an analogy with physical objects. Water waves and tornadoes are (at a given instant in time) made up of particular bits of matter. However, over time, their composition changes. Such things are multiply realised by different bodies of matter across time. The same could be said of certain data structures.

As a concession to Chalmers, I assume that high-level data structures exist in addition to low-level data structures. However, Chalmers posits further digital objects at an even higher level of abstraction. For example, in response to critics (2019: 465, 471–472), he suggests that some digital objects can be realised redundantly by several duplicate data structures at the same time, stored across distinct computers. Further, his position implies the existence of digital objects that can survive the wholesale discontinuous replacement of constitutive data structures. No analogy with physical objects is forthcoming here, nor are any uncontentious examples of such digital objects (i.e. that aren't also virtual objects) provided. As we shall see, this is a source of concern.

The central claim of virtual digitalism is that virtual objects *are* digital objects. For Chalmers, digital objects exist, so this is a form of realism about virtual objects. Chalmers's (2017) case for virtual digitalism centres around causal considerations, encapsulated in the following argument:

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. Therefore, virtual objects are digital objects. (Chalmers 2017: 318)

(Chalmers 2017: 318 also provides an argument based on perception, but it rests on similar causal considerations.) There is something compelling about this argument. Virtual objects appear to have causal powers, and those powers must in some sense belong to digital objects. However, in §4 I will show that these considerations are compatible with the rejection of virtual digitalism.

My case against virtual digitalism rests on two observations that undermine the analogy with physical objects. First, our individuation of virtual objects is often driven by our expectations, and insensitive to the arrangement of underlying data structures (modulo the preservation of certain appearances). Second, there are situations where virtual objects have inconsistent virtual properties. These observations suggest that virtual objects are conditioned by our thoughts and practices to a greater extent than physical objects. To substantiate my first observation, I provide two examples.

When playing *Super Mario Bros*, the player controls an avatar that looks like a plumber in red overalls (Mario). When the image of Mario overlaps the image of a mushroom, the mushroom disappears and the image of Mario is replaced with an image of a taller plumber that otherwise looks alike. Suppose that small-Mario is encoded by one data structure, while tall-Mario is encoded by another. These data structures are wholly distinct, so that the described appearances are not affected by a single data structure undergoing a change, nor by one data structure becoming another, but by one data structure being 'swapped out' with another.

According to the standard fiction, when Mario collects a mushroom, he gets bigger. Thus small-Mario is identical to tall-Mario. This, I suggest, leads us to identify the virtual objects representing small-Mario and tall-Mario. So, according to virtual digitalism, the virtual object representing Mario (virtual Mario) is a high-level digital object grounded in two distinct data structures. However, there are other possible fictions we might have associated with the game. According to one, when Mario collects a mushroom, he is replaced by his taller twin brother of the same name. If that were the associated fiction, we would take virtual small-Mario to be distinct from virtual tall-Mario.

Thus, our individuation of virtual objects is insensitive to what is going on 'under the hood' of the computer, modulo the generation of certain appearances. We are led by the expectations we bring to our engagement. Indeed, as my second example demonstrates, such expectations even influence our deployment of ontological categories in carving up virtual worlds.

Games in the Half-Life series have impressive physics engines. In *Half Life: Alyx*, a VR entry in which you occupy the perspective of the protagonist Alyx, you can manipulate virtual objects in a way that is strikingly realistic. For example, you can pick up a bottle of vodka and move it around, watching as the liquid inside swills in realistic ways. Here, Chalmers will say that the vodka is a distinct virtual object.

However, it turns out that the effect is achieved by applying a shader to the bottle. That is, what looks like goings-on 'inside' the bottle is, in fact, a change to its surface appearance. This is surprising, but virtual digitalism is flexible enough to handle it. Chalmers will say that the virtual vodka is identical to the shader, a digital object.

But this reveals an inconsistency in our carving up of virtual worlds, since we are not tempted to treat the effects of other shaders similarly. For example, we would not take the shininess of a virtual bottle to be a virtual object; we would take it to be a virtual *property* of a virtual bottle. This discrepancy will be partly due to our understanding of the fiction, but the remarkable resemblance to portions of the physical world will also play a significant role.

My first observation reveals a disanalogy between virtual objects and physical objects. When it comes to physical objects, individuation is not led by our understanding of a fiction (however broadly construed), or the expectations we bring to the world, at least not to the degree demonstrated here.

My second observation is that situations can arise in which a virtual object has two inconsistent virtual properties at the same time. Suppose that two players A and B are playing an online first-person shooter together, controlling virtual avatars A* and B*, respectively. In such games, players occupy the first-person perspective of their avatars and engage in a battle royale. Suppose that, on each machine, A* and B* are each encoded by a single data structure.

Now, suppose that A and B are playing peer-to-peer, meaning their respective computers are running synchronised versions of the game, and there is some latency on the network. Then the following may occur. At time t_1, A* virtually shoots and kills B*. However, the latency in the network causes a delay in the synchronisation long enough to allow B* to go on to virtually shoot and kill A* at t_2. Eventually, the delay is rectified so that at t_3 both A* and B* are virtually dead. B* is virtually dead because A* virtually shot B* at t_1, and A* is virtually dead because B* virtually shot A* at t_2. So, B* is both virtually dead and not virtually dead at t_2.

The fact that this inconsistency concerns *virtual* properties does not make it innocuous. According to virtual digitalism, virtual properties are real properties of digital objects, and are ontologically on a par with properties of physical objects (Chalmers 2017: 320–326). In the situation described, Chalmers would say that B* is a high-level digital object realised by two distinct but parallel data structures, and thus its properties are determined by the properties of those data structures (see Chalmers 2019: 471–472). And, in the situation described, the properties of these data structures determine that B* has two mutually inconsistent properties at the same time.

The above line of reasoning presumes that each of the two data structures fully realises B*, so the properties of each taken separately are sufficient but not necessary to determine the properties of B*. Might Chalmers instead claim that B* is partially realised by each data structure, so that the properties of just one are necessary but not sufficient to determine the properties of B*? In that case, B* would have indeterminate properties—be neither virtually dead nor virtually alive—which is perhaps less disturbing.

This will not work. To see why, imagine that A disconnects from the game, but B keeps playing. Here, the relevant data structure on A's computer no longer determines the properties of B*. Intuitively, B* survives this event. But, if B* is only partially grounded in each data structure, then removing one will destroy B*, meaning B carries on playing with a distinct virtual object B** identical with the data structure on B's computer. This seems wrong.

Worse, if we countenance B**, then, while A and B are still playing together, there will be three virtual objects where we would only identify one. As well as B*, the high-level digital object grounded in two distinct data structures, there is B**, the data structure on B's computer, and B***, the data structure on A's computer. This proliferates virtual objects and suggests that our intuitive individuation of virtual objects is way off the mark. And this, in turn, casts doubt on our intuitive attributions of causal relations to virtual objects, undermining the central motivation for virtual digitalism.

One might object that both of my observations have parallels for ordinary physical objects. Physical tools and artifacts are individuated in part by our expectations, and the problems of identity and material constitution, like the Statue-Clay and the Ship of Theseus paradoxes, show that ordinary physical objects also generate paradox (see Gallois 2016 and Wasserman 2021). However, the analogy does not hold.

To begin with, the Mario example involves identity being preserved across a discontinuous and wholesale change in digital constituents. I cannot think of an example where a physical object survives an analogous change, even among tools and artifacts. For example, if all the parts of my car were replaced all in one go, it would not be the same car, even if its functions (and my associated expectations) were preserved.

Nor is there a clear example in physical reality where something's status as either an object or a property is determined by our expectations. Granted, physical reality can be carved up in different ways depending on our expectations. The category of culinary nuts is determined by our tastes and practices, for example. And where others recognise only a single marble block, a skilled sculptor may distinguish a beautiful statue from the surrounding marble. But such examples do not show that whether something is a physical property, on the one hand, or a physical object, on the other, is sensitive to our expectations. Yet the example of the vodka bottle shows that virtual reality can exhibit this more fundamental sensitivity.

Nor is there a clear example of a physical object that is fully and redundantly realised by two distinct arrangements of matter, such that the properties of each realiser are sufficient but not necessary to fix the properties of the realised. The example of virtual inconsistency I gave has no clear physical analogue.

My observations show that the relationship between virtual reality and digital reality is messier than the relationship between physical objects and their constituents. This suggests that virtual reality is conditioned by us to a greater degree than physical reality. To the extent that virtual digitalism implies otherwise, we should be dissatisfied with it.

3. VIRTUAL FICTIONALISM

Virtual fictionalism claims that virtual objects are fictional objects that we imagine exist when we engage with videogames and VR systems as representational artworks. On this view, virtual objects either do not exist at all, or 'exist' only in a weak sense, as the transient intentional objects of certain imaginative episodes (McDonnell & Wildman 2019: 392). Either way, virtual fictionalism is a form of anti-realism about virtual worlds.

McDonnell and Wildman (2019) draw on Kendal Walton's (1990) theory of fictions as props in games of make-believe. On this view, when we engage with a fiction, our imaginings are guided by objective features of the work (alongside the context of creation and critical appreciation) via *principles of generation*.

For example, when I read *Dracula*, I imagine certain things. To properly appreciate the work, what I imagine will be constrained by objective features of it, via certain principles of generation. For example, one principle is 'if there is a passage of prose following the words *Mina Harker's Journal*, imagine that it is an extract from Mina Harker's journal, and imagine that it is a faithful account of actual events.' In this way, the work serves as a prop in a specific game of make-believe.

For McDonnell and Wildman (2019), when we play videogames and use VR, we are engaging in games of make-believe constrained by objective features of the technology. In so doing, we imagine that we are interacting with virtual worlds; in reality, there are no such worlds.

Virtual fictionalism appears to fare better than virtual digitalism. If virtual objects are fictional, it is not surprising that we individuate them in accordance with the associated fiction. Further, the inconsistent state of affairs described in §2 can be treated like a glitch. It is an interesting question when and how glitches are generative of fictional truth (see Van de Mosselaer & Wildman 2021 and Fisher 2022). If the situation I described is not generative, there will be no contradiction. If it is, it will be fictional that the avatar is both virtually dead and virtually alive at the same time. But that is not troubling: many fictions are (intentionally or unintentionally) inconsistent. There is also something to be said for the unification promised by virtual fictionalism. Videogames and VR experiences are representational artworks, so it is attractive to fold our thinking about them into a general account of such things.

Unfortunately, however, virtual fictionalism raises problems of its own. First, it mischaracterises the phenomenology of attending to virtual worlds. Second, a sensible account of which objective features of the works constrain our imaginings is not forthcoming. I elaborate each in turn.

Virtual fictionalism says that we ascertain virtual goings-on by imagining them. Imagination is a distinctive psychological attitude that differs from other cognitive states, like belief, desire, and perception, by virtue of its cognitive role (see Nichols & Stitch 2000; 2003: Ch.2). This view has come to enjoy a broad consensus (see Carruthers 2006; Currie & Ravenscroft 2002: Ch.2; Friedman & Leslie 2007; Gendler 2006; Schroeder & Matheson 2006; Weinberg & Meskin 2006). The most influential accounts of fiction take imagining to be central to apprehending and appreciating fictional content (see Abell 2020; Currie 1990; Stock 2017; Walton 1990). Associated with this cognitive distinction, there is a phenomenological distinction. Believing or seeing that I have won the lottery feels very different to imagining that I have won the lottery.

There are detractors from the broad consensus that imagination is a distinctive cognitive attitude. For example, Peter Langland-Hassan 2020 reduces imagination to belief-based counterfactual reasoning. I am sceptical that our emotional engagement with fiction can be adequately explained on his view (cf. Currie 1990: 210), nor our capacity to be immersed in fiction in various ways (see Chasid 2017 and Kampa 2018). However, here I need only note that believing or seeing that *p* feels different to thinking about how things would be were *p* true. So, even if Langland-Hassan is right, there is still a phenomenological distinction to be drawn.

I suggest that the phenomenological distinction between believing/seeing and imagining explains the difference in how it feels to attend to a virtual world itself, and how it feels to attend to a fictional world represented by it.

Sometimes, when playing a videogame, the virtual world can intrude in ways that break immersion. For example, in games depicting outdoor areas, a common way of generating the appearance of open sky is to cover the map in a dome with a textured inner surface that looks like sky from within. From some vantage points, you may be able to see the dome. When you do, your imaginative engagement is interrupted. It feels like spotting how a magic trick is done—like seeing how things actually are, rather than how they are contrived to appear. For the fictionalist, however, cases like this are a matter of going from one imaginative episode to another.

The fictionalist may retort: when we notice a feature of game design like a sky-dome, we come to imagine something that conflicts with the fictional world imagined up to that point. In the world we have been imagining, only the sky is above us; in the world we are presently imagining, we are encased in a dome. This breaks our immersion because the fiction stops making sense. We realise that something has gone wrong, that we are not supposed to imagine what we have been led to imagine (see Van de Mosselaer & Wildman 2021 for discussion of fictional misfires).

This may account for certain jarring experiences, like seeing non-player characters behaving in inexplicable ways due to a bug; but it doesn't capture the sense that we are seeing things for how they are. Take, for example, playing a simple game like *Tetris*, where the player must fit differently shaped blocks ('tetrominoes') together as they fall with increasing speed. While you can imaginatively engage with *Tetris*, perhaps imagining that you are controlling physical blocks falling through a physical space, I doubt that most players are doing this. They are just seeing virtual blocks moving through a virtual space. This is reflected in the phenomenology: playing *Tetris* just doesn't feel like imagining anything. (Chalmers (2017: 320–321, 331) emphasises our ability to experience virtual worlds in this way.) More generally, attending to virtual worlds *as virtual worlds* just doesn't feel like make-believing.

The second problem for fictionalism is the absence of a sensible account of what governs our imaginings. In certain circumstances, there are determinate facts of the matter about a virtual world, even if we are not in a position to ascertain them. As we shall see, the kinds of props that are supposed to determine fictional truth cannot accommodate this.

The kind of situation I have in mind is this. Suppose that, while playing a videogame, I turn off my monitor and speakers but continue to use the controller. It seems clear to me that there are determinate facts of the matter concerning what is going on in the virtual world. After all, there are determinate facts of the matter concerning what is going on digitally, caused by my inputs to the controller, and these (along with broad features of context (see §2)) fix what goes on virtually. Specifically, there are facts about where in the virtual world my avatar is located, what actions my avatar is virtually performing, and what consequences those actions have in-game, even though I am currently not in a potion to ascertain them.

Virtual fictionalism says that the virtual world is generated by a game of make-believe, constrained by objective features of the game (the props). Which objective features are these? McDonnell and Wildman state that they will 'include digital elements like the particular images, sounds, and haptic feedback mechanisms (2019: 391).' The suggestion seems to be that the props are elements of sensory feedback.

Given this, the virtual fictionalist can say one of three things about the above situation. First, perhaps there are some prescriptions to imagine in the absence of props. It has long been recognised that we infer some fictional contents without their being explicitly represented (see Lewis 1978). For example, if I were to turn on the monitor again to find my avatar in a virtual room adjacent to the one it started in, we may be able to infer (given our knowledge of the game) that my avatar virtually walked (rather than teleported, say) to its current location. However, this will only allow us to infer *some* of what went on, and only *after* the sensory outputs have been restored. It does not capture the full range of specific facts concerning what went on during the blackout.

Second, perhaps we can infer some things on the basis of my inputs to the controller. This is better, but it still won't do. It will allow us to imagine some vague contents; but the facts about what is going on in the virtual world will far outstrip such guesswork in detail. Moreover, our guesswork may easily be off the mark. For example, suppose that, for the entirety of the blackout, I hold the button that makes my avatar virtually walk forward. I might imagine on this basis that my avatar continued to virtually walk forward during the blackout. However, suppose that, unbeknownst to me, right in front of my avatar was a virtual trapdoor, so that when the monitor is turned back on, I find my avatar virtually dead in a hole. Here, what was imagined on the basis of my inputs was not even close to what happened in the virtual world.

Finally, the virtual fictionalist could deny that there are determinate facts of the matter about the virtual world when the props are absent. In effect, the blackout renders the fiction incomplete and permits us to complete it in whichever way we choose (within certain constraints), after the sensory outputs have been restored. (Perhaps we are prescribed to complete it in some way, even if we are free to choose among a range of options. See Wildman & Woodward 2018.) I consider this to be a mark against virtual fictionalism, since it seems clear to me that there are determinate facts of the matter here, despite our lack of access to them.

Perhaps my intuitions here are not widely shared. Unfortunately, I lack a non-question-begging way of arguing that there are virtual facts of the matter in the absence of sensory feedback. All I can say is that, to the extent that one agrees with me that there are such facts, one will consider its inability to accommodate them a weakness of virtual fictionalism.

Might the virtual fictionalist instead appeal to data structures as the props that constrain our imaginings? This would solve the present concern, since the relevant facts concerning data structures are unaffected by the blackout. However, we require an account of how, in normal circumstances, the right goings-on among the data structures are identified by users and interpreted for fictional content. The only sensible account is that the props are perceived and interpreted via the appearances they generate through the output devices. For example, we might perceive certain goings-on among the data structures via their generation of the appearance of a tetromino, and on that basis come to imagine that we are seeing a tetromino.

DOI: 10.5334/met.156

But then why insist that the arrangement of data structures that produces the appearance only makes it *fictional* that one is seeing a tetromino? Why not accept that it makes it *true* that one is seeing a tetromino? Granted, when an arrangement of data structures produces the appearance of a dragon, it doesn't make it true that one is seeing a dragon. That would be absurd. If we were seeing a dragon, we would be able to see it without the aid of computer technology, and it would take up considerable physical space. But if we were seeing a virtual object like a tetromino, we would not be able to see it without the aid of a computer, and it wouldn't take up physical space. So why not say, simply, that one sees a tetromino via the appearances generated by its constitutive data structures? It is hard to see what we gain by characterising the tetromino as fictional, while we stand to lose a plausible account of the phenomenology of attending to it.

4. VIRTUAL SOCIALISM

Virtual digitalism imbues virtual reality with too much independence from us (§2), while virtual fictionalism imbues it with too little (§3). I suggest that we walk the line between these extremes by saying that virtual reality is socially constructed from digital reality.

Peter Ludlow (2019) argues that virtual reality is socially constructed, but my proposal differs in both ambition and detail. For Ludlow, absent a suitable social context, there is no objective fact of the matter regarding the computational state of a given physical system, from which it follows that virtual reality is 'social all the way down (Chalmers 2019: 457).' By contrast, my view treats digital reality as existing independently of social context. One could accept my view, and also accept that digital reality is socially constructed in some way, but the latter claim is not part of my view. In this sense, my view is less ambitious than Ludlow's.

We overlap in what we say about the virtual counterparts to ordinary physical objects. We both claim that being a virtual object of a certain kind, or having a certain virtual property, or being a virtual object at all, is partly socially determined (see Ludlow 2019: §4). However, Ludlow is quiet about *how* virtual objects are socially constructed. What little he says suggests that a virtual fact obtains by virtue of collective acceptance that it obtains (2019: 365–366). The view I develop draws on a more sophisticated account of how the social world is constructed, where the obtaining of a virtual fact need not require people to bear any particular psychological attitude towards it.

I draw on Brian Epstein's (2015, 2016) *anchor-grounding model* of the social world. On this model, to understand a social fact, we can ask after its grounds (the bits and pieces that make it up) and its anchors (that which explains why those bits and pieces make it up). This is a 'two-dimensional' approach to social ontology (see Bouwer 2022: §1.3 for discussion).

To illustrate, consider the fact that the thing in my pocket is a one-pound coin. To investigate its grounds, we can ask what *determines* it. A plausible answer is that the coin has certain physical qualities, $q_1 \dots q_n$, and is issued by the Royal Mint. If that's right, then the following grounding principle is in place:

> If an object x has physical qualities $q_1 \dots q_n$ and is issued by the Royal Mint, then this fact grounds the fact that x is a pound coin.

We can also ask what it *depends* on. Suppose that all pound coins must be minted by the Royal Mint, but that they must satisfy only one of a disjunction of physical qualities, $q_1 \dots q_n$ or $r_1 \dots r_n$ or $s_1 \dots s_n$. Then the following grounding principle is also in place:

> If an object x is a pound coin, then this fact is grounded in the fact that x has $q_1 \dots q_n$ or $r_1 \dots r_n$ or $s_1 \dots s_n$ and is issued by the Royal Mint.

To investigate the anchors, we ask why *that* is what it takes for the fact to obtain, or why the relevant grounding principles are in place. Here, we want to know why having physical qualities $q_1 \dots q_n$ and being issued by the Royal Mint makes the thing in my pocket a pound. Regarding the latter condition, we might point to the fact that all coins issued by the Royal Mint are authorized by Royal Proclamation, in accordance with the Coinage Act of 1971. Depending on one's views about laws, one may take this to ultimately bottom out with facts about conventions, attitudes, or behavioural regularities. Regarding the former condition, we might point out that the qualities q_1 … q_n ensure that pound coins have certain properties that are desirable given their function, such as being long-lasting and difficult to counterfeit.

Whether or not the above is an accurate or complete account, it illustrates two important features of the model. First, grounds and anchors are heterogeneous. In our example, they include social facts alongside physical facts. Second, while the anchors and grounds in the example may involve the attitudes and behaviours of individuals (upstream to the cited social facts), neither require anyone to have any particular attitude towards my coin. My coin can satisfy the above grounding conditions without anyone recognising that it does. If it fell behind my sofa and was forgotten about completely, it would remain a pound coin. Even if it outlasted humanity, it would remain a pound coin.

Nevertheless, the above example involves a portion of the social world that is to some degree constructed deliberately. As artifacts, I take virtual worlds to be similar. However, other parts of the social world are not like this (see Khalidi 2015). For example, an economy can enter a recession without anyone intending it, being aware of it, or even having the concept of a recession. The anchor-grounding model accommodates such examples (see e.g. the discussion of un-American food in Epstein 2019: §7), but I won't discuss them further here.

Let us apply the anchor-grounding model to virtual worlds. I begin with general facts concerning their existence and constitution, and propose that the following grounding principle is in place:

> If D is a plurality of data structures forming an ongoing, interactive computer-based causal process generating a complete and interconnected space, then this fact grounds the fact that a particular virtual world W exists and is constituted by D.

A few clarificatory points are in order. First, virtual worlds are constituted by (the elements of) certain processes. When it comes to videogames, we tend to call these processes 'playings' or 'playthroughs'. The virtual world I engage with when I play *Super Mario Bros* is distinct from the virtual world my friend engages with when playing the same game on her computer. What happens in one doesn't affect what happens in the other.

Second, the notion of space I invoke is the intuitive one that applies to both virtual and physical space. Following Chalmers (2022: 430–432), we can elaborate on this by invoking a functionalist notion of space, where a space is just that which mediates motion and interaction in the right way, and which causes spatial perception. Certain elements of a playing will have spatial values determining where and how they appear and move onscreen, causing spatial-like perceptions on the basis of player interactions. Moreover, the game will regulate the interaction of these elements so that they interact more in certain ways as their spatial values converge, and less as their spatial values diverge. In this way, a playing can generate a space.

Finally, constitution could be taken to be a primitive. However, I understand it to be a relation between a single constituted entity and a distinct constitutive entity or plurality that obtains just in case the constituents coincide with the constituted entity, and facts about them partially ground the fact that the constituted entity has the makeup it has (see Epstein 2016: Ch 10). For example, a statue is constituted by a lump of clay just in case the lump coincides with the statue, and the existence of the lump at least partially grounds the fact that the statue has the parts it has. In the present case, talk of parts is strained, but we can talk about local physical realisations, i.e. bits. Thus, a plurality of data structures constitutes a virtual entity just in case the plurality coincides with the virtual entity, and its existence partially grounds the fact that the virtual entity has the physical realisations it has.

Let us turn to facts about particular virtual objects and properties. Such things are identified on the basis of their distinctive causal roles. For example, I identify virtual Mario by the following: it appears as a mustachioed man in red overalls; it moves in certain ways when I input certain commands into the controller; it grows when it touches a virtual mushroom; and so on. Similarly, during a playing of *Half-Life: Alyx*, I identify the shininess of a virtual bottle by its causal role. This suggests the following grounding principles:

> If W is a virtual world constituted by D, and D′ is the smallest part of D that discharges the causal role distinctive of virtual object V, then these facts ground the fact that V exists in W and is constituted by D′.
>
> If W is a virtual world constituted by D, and D′ is the smallest part of D that discharges the causal role distinctive of virtual object V, and D″ is the smallest part of D′ that discharges the causal role distinctive of the instantiation of virtual property P by V, then these facts ground the fact that V has P in W and P's instantiation by V is constituted by D″.

So far, I have only provided grounding principles for how virtual facts are determined. There will also be principles specifying what virtual facts depend on, such as the following:

> If virtual object V exists in virtual world W, then this fact is fully grounded by the fact that there is a unique smallest part of the plurality of data structures that constitutes W that discharges the causal role distinctive of V.

Note that, while it is facts about particular data structures that determine facts about virtual worlds, facts about virtual worlds depend only on there being some plurality of data structures or other that discharges the relevant causal role.

The above is only a sketch: a more thorough statement of grounding principles may include a temporal index, for example. However, it should be clear enough how the grounding inquiry should proceed from here. For now, let us ask after the anchors of these grounding principles.

Consider first the principles specifying what virtual facts depend on. A virtual fact depends on some plurality of data structures or other discharging a certain causal role. The identity, number and arrangement of data structures does not matter, so long as the causal role is discharged.

We are therefore asking why the discharging of *this* causal role is required for the relevant virtual fact to obtain. In some cases, the obtaining of a virtual fact will require that data structures discharge a causal role that is sufficiently similar to the real-life equivalent. For example, in a life-like VR experience, to constitute a virtual apple, a plurality of data structures will have to produce appearances that closely resemble real apples. In other cases, however, appearing as a reddish blur may be sufficient for a plurality of data structures to constitute a virtual apple, so long as it is mutually understood that a reddish blur represents an apple.

Thus, among the anchors of the grounding principles that specify what virtual facts depend on are features of the context in which the virtual world is designed and enjoyed. In cases where a virtual world aspires to be life-like, facts about resemblance with real life will come into play. In cases where a story is being weaved, facts about the shared understanding of the fiction will come into play. In both cases, facts about authorial intentions and their uptake are relevant, since what the videogame and VR experience aspires to must be, to some extent, mutually understood. Further, in both cases, biological and psychological facts about how people are disposed to interpret sensory information will play a role as well.

Turning now to the principles that specify determining facts, the question here is why *this* (rather than *that*) particular arrangement of data structures makes it the case that the virtual fact obtains. Here, among the anchors will be facts about the developers. Developers will ensure that the relevant causal roles are discharged in some way, but the specific way they choose will be a function of the technology they use, their skill and creativity in using it, the time and resources available to them, and so on. Such factors will explain why virtual facts are grounded in the particular ways they actually are.

We can now state virtual socialism more conspicuously. Facts about virtual worlds (and the objects, properties, and events that populate them) are grounded in facts about pluralities of data structures. These grounding conditions are anchored by a range of different facts (physical, psychological, biological, and social) concerning the context in which virtual worlds are designed and enjoyed. We have seen that details will change from case to case, and I have only sketched anchors and grounds for a handful of virtual facts. But I hope to have done enough to show how the details might be filled in, and which considerations will be pertinent to the task. To close, I will defend virtual socialism by demonstrating how it avoids the difficulties faced by its two rivals.

We saw that virtual digitalism draws an analogy between physical objects and virtual objects. As such, it fails to account for the extent to which our individuation of virtual objects is led by our expectations, and the ways in which virtual worlds can generate inconsistency. Virtual socialism has no such difficulty.

Regarding individuation, virtual socialism assigns a role to our expectations. Recall the example involving Mario growing. The context in which *Super Mario Bros* games are developed and enjoyed is such that there is a shared understanding of the narrative elements of the game. This fixes the causal role distinctive of virtual Mario, and anchors grounding principles relevant to Mario's growing.

The fact that Mario grows on a particular occasion depends on some arrangement of data structures or other effecting the small-Mario image being replaced with the tall-Mario image. These grounding conditions are insensitive to the details regarding the number and arrangement of data structures involved. Thus, the obtaining of the virtual fact is insensitive to what is going on 'under the hood', modulo the preservation of the right appearances. This explains why virtual objects can survive a wholesale and discontinuous replacement of constitutive material. Moreover, we can see that, were the anchors sufficiently different, such that the mutually understood fiction said that Mario is replaced by his taller brother when he eats a mushroom, different grounding principles would be in place that connect the same digital facts and appearances with a different virtual fact, namely the fact that virtual small-Mario is replaced by virtual large-Mario. Thus, virtual socialism predicts that our individuation of virtual objects will be led by our expectations, including those generated by the associated fiction.

Our shared understanding of narrative elements will also play a role in our identification of the virtual vodka as a virtual object in *Half-Life: Alyx*. The gritty fictional surroundings we are to imagine means we are prone to interpret certain appearances in certain ways. However, unlike *Super Mario Bros*, *Half-Life: Alyx* aspires to give players a relatively life-like experience. Thus, resemblance to real life objects and properties plays a role here. The appearances generated by the relevant shader closely resemble vodka sloshing around in a bottle—closely enough to ensure that the shader constitutes a virtual body of vodka.

Regarding the shininess of the bottle, our understanding of the narrative, along with the fact that the relevant shader produces appearances that closely resemble the shininess of a bottle, anchors grounding principles to the effect that this shader constitutes a property instance of an existing virtual object, rather than a further virtual object. In this way, the context in which a virtual world is produced and enjoyed can anchor grounding conditions that exhibit a striking lack of parity in terms of which kinds of digital facts determine which kinds of virtual facts.

Virtual socialism also accommodates cases of virtual inconsistency. Sometimes, a collection of grounding principles can, in certain circumstances, generate inconsistent social facts from consistent grounds (Brouwer 2022). The following provides a simple illustration. (The example is Priest's (2006 [1987]: 184–185); Brouwer (2022: §3) puts it in anchoring/grounding terms.) Imagine a society where two laws are passed. One states that women do not have the right to vote; another states that anyone with a certain amount of property has the right to vote. Imagine further that the notion of women owning the required amount of property is unimaginable to the lawmakers.

The passing of the first law anchors a principle that means someone's being a woman determines their lack of right to vote. The passing of the second law anchors another principle that means someone's owning the specified amount of property determines their right to vote. Now, suppose that the imagined society undergoes some social progress, and eventually a woman comes to

DOI: 10.5334/met.156

own the amount of property specified by the second law. In this situation, legally, this woman has the right to vote and does not have the right to vote. Thus, we have a contradiction in social reality, while the facts that ground it are consistent, since being a woman is not inconsistent with owning any amount of property.

In the above situation, the inconsistency arose in part because of a lack of imagination on the part of the lawmakers. I suggest that something similar is going on with the case where the avatar B* is both virtually dead and not virtually dead at the same time. Here, rather than a lack of imagination, the developers fail to make physical reality cooperate fully with their designs. The developers anchored a collection of grounding principles that can, in situations where there is significant delay on the network, generate inconsistent virtual facts from consistent digital grounds.

Not only does virtual socialism illuminate how contradictions arise in virtual reality; it renders them tolerable. That there are inconsistent social facts seems beyond reproach (Brouwer 2022), so, if we take virtual realities to be social, inconsistent virtual realities are not a further source of inconsistency. Moreover, even though I assume that social reality is real, it is not fundamental, and, for all we have seen, its ultimate grounds are consistent. Thus, virtual socialism reassures us that virtual inconsistency does not run deep.

As well as overcoming its difficulties, virtual socialism is compatible with the causal considerations presented in favour of virtual digitalism. Chalmers's case rests on the fact that virtual objects appear to have causal powers, and that those causal powers surely belong to digital objects. According to virtual socialism, if a given virtual object exists in a virtual world, it has a distinctive causal role in terms of which it is individuated, and that causal role is discharged by a plurality of data structures. It follows that, each time a virtual object causally affects or is affected by something, that causal relationship is realised by certain data structures. Thus, we can agree with Chalmers that virtual objects have causal powers that in some sense belong to digital objects, and still deny that virtual objects *are* digital objects.

There is more to be said about causation in virtual reality (see Wheeler 2022 for recent discussion). Some might take what I have said to imply that there is no causation at the virtual level, only the appearance of causation produced by causation at the digital level. This will depend on one's general views about causation. On the view I favour, one thing causes another just in case, under a suitable intervention on the former, the latter changes. We intervene on virtual objects all the time, and these interventions result in systematic changes in other virtual objects. This can be so, even if each causal relationship is realised by data structures. While I recognise that all this is far from uncontroversial, I lack the space to elaborate and defend it fully here. (See Woodward 2015 for how one might proceed.)

Virtual socialism also fares better than virtual fictionalism. First, it is compatible with the Waltonian account of our appreciation of videogames and VR experiences. According to virtual socialism, virtual reality is real and objective, relative to certain anchors being in place. That is, given the broader context in which videogames are produced and enjoyed, the grounding conditions for virtual facts are fully determinate and not dependent on our ability to recognise their satisfaction on any given occasion. In this way, the virtual reality constituted by a given playing is prior to our imaginative engagement with it and is thus able to appropriately constrain our imaginings. So, we can say that, when we properly appreciate a videogame or VR experience as a work of fiction, we engage in a game of make-believe that is appropriately constrained by observable features of the virtual world constituted by our playing.

Second, virtual socialism avoids the problems raised for virtual fictionalism. It captures the phenomenological contrast between attending to a virtual world and attending to a fictional world. On virtual socialism, when we attend to a virtual tetromino, we perceive a real virtual object; and when we attend to the fictional physical block it perhaps represents, we imagine we are seeing a physical block. Further, since virtual socialism identifies the props that constrain our imaginative engagement with virtual objects and properties, it accommodates the fact that there are determinate facts of the matter about virtual worlds, even when we cannot ascertain them.

So long as the right anchors and grounds are in place, things can go on in virtual worlds without our knowing about them.

ACKNOWLEDGEMENTS

Thanks to Leonid Tarasov and Andrew Kirton for formative discussions and for the vodka-bottle example. Thanks to members of the Philosophy Research Group at Swansea University for their valuable feedback on an early draft. Thanks to the two anonymous reviewers at Metaphysics for their constructive comments. And thanks to the Paul Schmitt Foundation for their financial support of the publication of this article.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR AFFILIATIONS

Robert Fraser orcid.org/0000-0002-1475-1863
Swansea University, UK

REFERENCES

Abell, C. 2020. *Fiction: A Philosophical Analysis*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/oso/9780198831525.001.0001

Brouwer, T. 2022. 'Social Inconsistency'. *Ergo,* 9. DOI: https://doi.org/10.3998/ergo.2258.

Carruthers, P. 2006. 'Why Pretend?' In Nichols, S. (ed.) *The Architecture of the Imagination: New Essays on Pretense, Possibility, and Fiction*. (pp. 89–109). Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199275731.003.0006

Chalmers, D. 2017. 'The Virtual and the Real'. *Disputatio,* 9: 309–352. DOI: https://doi.org/10.1515/disp-2017-0009

Chalmers, D. 2019. 'The Virtual as the Digital'. *Disputatio,* 11: 453–486. DOI: https://doi.org/10.2478/disp-2019-0022

Chalmers, D. 2022. *Reality+. Virtual Worlds and the Problems of Philosophy*. London: Allen Lane.

Chasid, A. 2017. 'Imaginative Content, Design-Assumptions and Immersion'. *Review of Philosophy and Psychology,* 8: 259–272. DOI: https://doi.org/10.1007/s13164-016-0315-2

Currie, G. 1990. *The Nature of Fiction*. Cambridge: Cambridge University Press.

Currie, G and **Ravenscroft, I.** 2002. *Recreative Minds: Imagination in Philosophy and Psychology*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780198238089.001.0001

Epstein, B. 2015. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199381104.001.0001

Epstein, B. 2016. 'A Framework for Social Ontology'. *Philosophy of the Social Sciences,* 46: 147–67. DOI: https://doi.org/10.1177/0048393115613494

Epstein, B. 2019. 'Anchoring versus Grounding: Reply to Schaffer'. *Philosophy and Phenomenological Research,* 99: 768–781. DOI: https://doi.org/10.1111/phpr.12644

Fisher, A. 2022. 'Truth in Interactive Fiction'. *Synthese,* 200: 1–18. DOI: https://doi.org/10.1007/s11229-022-03935-0

Friedman, O and **Leslie, AM.** 2007. 'The Conceptual Underpinning of Pretense: Pretending is Not 'Behaving-As-If''. *Cognition,* 105: 103–124. DOI: https://doi.org/10.1016/j.cognition.2006.09.007

Gallois, A. 2016. 'Identity Over Time'. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* URL = <https://plato.stanford.edu/archives/win2016/entries/identity-time/>.

Gendler, T. 2006. 'Imaginative Contagion'. *Metaphilosophy,* 37: 183–203. DOI: https://doi.org/10.1111/j.1467-9973.2006.00430.x

Kampa, S. 2018. 'Imaginative Transportation'. *Australasian Journal of Philosophy,* 96: 683–96. DOI: https://doi.org/10.1080/00048402.2017.1393832

Khalidi, MA. 2015. Three Kinds of Social Kinds. *Philosophy and Phenomenological Research,* 90: 96–112. DOI: https://doi.org/10.1111/phpr.12020

Langland-Hassan, P. 2020. 'Explaining Imagination'. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/oso/9780198815068.001.0001

Lewis, DK. 1978. 'Truth in fiction'. *American Philosophical Quarterly,* 15: 37–46.

Ludlow, P. 2019. 'The Social Furniture of Virtual Worlds'. *Disputatio,* 11: 345–369. DOI: https://doi.org/10.2478/disp-2019-0009

McDonnell, N and **Wildman, N.** 2019. 'Virtual Reality: Digital of Fictional?' *Disputatio,* 11: 371–397. DOI: https://doi.org/10.2478/disp-2019-0004

Nichols, S and **Stich, S.** 2000. 'A Cognitive Theory of Pretense'. *Cognition,* 74: 115–147. DOI: https://doi.org/10.1016/S0010-0277(99)00070-0

Nichols, S and **Stich, S.** 2003. *Mindreading*. New York: Oxford University Press. DOI: https://doi.org/10.1093/0198236107.001.0001

Priest, G. 2006. [1987] *In Contradiction*. 2nd ed. Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199263301.001.0001

Schroeder, T and **Matheson, C.** 2006. 'Imagination and Emotion'. In Nichols, S. (ed.) *The Architecture of the Imagination*. (pp. 19–40). Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199275731.003.0002

Stock, K. 2017. *Only Imagine: Fiction, Interpretation, and Imagination*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/oso/9780198798347.001.0001

Vallance, C. 2024. 'Police Investigate Virtual Sex Assault on Girl's Avatar'. *BBC News* 2 January.

Van de Mosselaer, N and **Wildman, N.** 2021. 'Glitches as Fictional (Mis)communication'. In Barker, T. & Korolkova, M. (eds.) *Miscommunications: Errors, Mistakes, Media*. (pp. 300–315). London: Bloomsbury Academic. DOI: https://doi.org/10.5040/9781501363825.0029

Walton, K. 1990. *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Cambridge, MA: Harvard University Press. DOI: https://doi.org/10.2307/2108134

Wasserman, R. 2021. 'Material Constitution'. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/fall2021/entries/material-constitution/>.

Weinberg, J and **Meskin, A.** 2006. 'Puzzling Over the Imagination: Philosophical Problems, Architectural Solutions'. In S. Nichols (ed.) *The Architecture of Imagination* (pp. 175–204). Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199275731.003.0010

Wheeler, B. 2022. 'Causation in a Virtual World: A Mechanistic Approach'. *Philosophy and Technology,* 35: 1–26. DOI: https://doi.org/10.1007/s13347-022-00502-w

Wildman, N and **Woodward, R.** 2018. 'Interactivity, Fictionality, and Incompleteness'. In Robson, J. & Tavinor, G. (eds.) *The Aesthetics of Videogames*. (pp. 112–127). New York: Routledge. DOI: https://doi.org/10.4324/9781315210377-8

Wolswijk, H. 2012. 'Theft: Taking a Virtual Object in RuneScape: Judgment of 31 January 2012, case no. 10/00101 J'. *Journal of Criminal Law,* 76: 459–62. DOI: https://doi.org/10.1350/1740-5580-76.6.459

Woodward, J. 2015. 'Interventionism and Causal Exclusion'. *Philosophy and Phenomenological Research,* 91: 303–347. DOI: https://doi.org/10.1111/phpr.12095

TO CITE THIS ARTICLE:
Fraser, R. 2024. The Social Reality of Virtual Worlds. *Metaphysics*, 7(1), pp. 85–98. DOI: https://doi.org/10.5334/met.156

Submitted: 15 March 2024
Accepted: 04 October 2024
Published: 25 November 2024

Metaphysics is a peer-reviewed open access journal published by Ubiquity Press.