Dialogue (2024), 1–21
doi:10.1017/S0012217324000155

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

The Constitution of Virtual Objects

Guillaume Bucchioni[1] and Alexandre Declos[2]

[1]Department of Philosophy, Aix-Marseille Université, France and [2]Institute of Philosophy, University of Neuchâtel, Neuchâtel, Switzerland
Corresponding author: Alexandre Declos; Email: alexandre.declos@gmail.com

Abstract
David Chalmers argues that virtual reality is a genuine kind of reality. In one of its readings, this "virtual realism" states that virtual entities ontologically depend on real digital entities. This article explores that suggestion and offers a novel account of the dependence of the virtual on the digital. Drawing on Lynne Rudder Baker's theory of constitution, we contend that virtual objects should be seen as *constituted* by digital objects, when these are placed in certain favourable circumstances. We explore the rationale and implications of this view, which we see as a promising form of virtual realism.

Résumé
David Chalmers maintient que la réalité virtuelle est une véritable réalité. Une version de ce « réalisme virtuel » affirme que les entités virtuelles dépendent ontologiquement d'entités numériques réelles. Nous explorons ici cette suggestion, en proposant un nouveau modèle pour décrire la dépendance du virtuel à l'égard du numérique. En nous appuyant sur la théorie de la constitution de Lynne Rudder Baker, nous défendons que les objets virtuels sont *constitués* par des objets numériques, lorsque ceux-ci se trouvent dans certaines circonstances favorables. Nous explorons les motivations et les implications de cette approche, que nous considérons comme une forme prometteuse de réalisme virtuel.

Keywords: virtual reality; constitution; realism; artifacts; Chalmers

1. Introduction

If you play *Skyrim VR*, you will encounter a plethora of *virtual objects*: virtual swords, virtual tables, virtual people, virtual trees, etc. You'll also take part in various *virtual events*: virtual fights, virtual walks, virtual conversations, and so on. Part of the recent philosophical literature on Virtual Reality (VR) has been focused on whether such objects/events are real. While David Chalmers argues that they are (Chalmers, 2017, 2019, 2022), others maintain that these entities are merely fictional (see e.g., Beisbart, 2019; Juul, 2019; McDonnell & Wildman, 2019).

Our concern in this article isn't the dispute between virtual realism and virtual fictionalism. Instead, our goal will be to flesh out a claim that was envisioned but

not fully developed by Chalmers, according to which virtual entities *ontologically depend* on digital entities. Specifically, we'll argue that virtual realists would do better to see virtual entities as *constituted* by digital objects. In order to develop this view, we shall draw on Lynne Rudder Baker's theory of constitution, which we adapt, with some departures, to the virtual realm. We'll explore the rationale and consequences of this constitution approach to the virtual, which we take to be a promising form of virtual realism. It should be stressed that our claim is only conditional: we won't defend virtual realism against fictionalist opponents here. Indeed, we only suggest that the constitution approach is an attractive, and as of yet underexplored, version of virtual realism.

The plan for this article is as follows. In Section 2, we discuss Chalmers' virtual realism and specify which version of this view we are interested in. In Section 3, we examine the notion of constitution, as theorized by Baker. Following that, we show how constitution theory can be applied to the virtual realm by considering in turn how digital entities (Section 4) and mental states (Section 5) are involved in the constitution of virtual objects. In Section 6, we highlight a number of interesting differences between digital and virtual objects. To conclude, we briefly review the virtues of our version of virtual realism (Section 7).

2. Virtual Realism(s)

Chalmers argues that "virtual reality" is aptly named. According to him, VR environments should indeed be considered as genuine realities, rather than mere representations or fictions (Chalmers, 2017, 2019, 2022). At the ontological level, this so-called "virtual realism" rests on the following claims:

> 1) Virtual objects really exist and are digital objects;
> 2) Events in virtual worlds are largely digital events that really take place. (Chalmers, 2017, p. 311)

It is important to distinguish two things in this twofold assertion. First, virtual objects and events are *real*: that's a metaphysical claim about the *status* of these entities. Second, virtual objects and events are *digital*: that's an ontological claim about the *nature* of these entities. Given our intent to leave aside the whole fictionalism/realism debate, we'll only focus on the ontological claim here. For simplicity, we shall also restrict ourselves to discussing virtual *objects*, although we think that what we'll say later on could equally apply to virtual *events*.

The fundamental ontological claim behind virtual realism is that virtual objects are digital objects. That's what Chalmers calls "digitalism." Digitalism can be understood in a number of ways, depending on what one means exactly by "digital object."[1] In a narrow and minimal sense, this notion can simply denote a *bit*, i.e., a 0 or 1 in a computational system. In a broader sense, a digital object is a "data structure": a computational object made of bits, which consists in a certain organization of data (e.g., a list, array, tree, linked list, table, …) and which is ultimately physically realized

[1] For a detailed discussion, see Beisbart (2019), Chalmers (2019), and Ludlow (2019).

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

as electrical impulses in integrated circuits. The latter characterization hints that data structures can be conceived either as abstract types, or alternatively, as the physical realization of abstract data in a piece of computer hardware. When Chalmers says that virtual entities are digital entities, he clarifies that what he has in mind is the notion of "physical data structure," understood as "a *realization* of the abstract data structures" (Chalmers, 2019, p. 460). In other words, a data structure can be seen as the concrete token of an abstract type.

The claim that virtual objects are digital objects is amenable to at least two different readings, depending on how data structures are factored in the equation:

> To a first approximation, [virtual objects] 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. (Chalmers, 2017, p. 317)

Chalmers is here considering two different views concerning the relation of virtual objects, understood as digital entities, and data structures. The first says that virtual objects and data structures are *identical*, while the second holds that virtual objects *depend* on data structures. Simplifying the picture, we can follow Neil McDonnell and Nathan Wildman (2019) in thinking that what Chalmers says here translates into two different versions of virtual realism. On the first version, virtual entities are taken to be strictly *identical* to digital entities, such as data structures. McDonnell and Wildman (2019) have labelled this view "Strong Virtual Digitalism" (SVD). On the second version, virtual entities and digital entities are taken to be distinct, but such that the former ontologically *depend* on the latter. A virtual object (e.g., a virtual sword) would depend on its associated data structure, much in the same way that a statue may be said to depend on the clay of which it is made. This is what McDonnell and Wildman call "Weak Virtual Digitalism" (WVD).

As Chalmers and others have noted, SVD is objectionable, as it fails to account for the *multiple realizability* of virtual objects. Indeed, SVD rigidly identifies any virtual object to a given data structure, although one and the same virtual object can be associated with a number of *different* data structures. This is what happens in cases of "porting," where a program written in a certain computer language is rewritten in another language fitted to another type of hardware. In such cases, the same virtual object — a virtual sword in *Skyrim VR*, say — corresponds to distinct data structures on different VR systems, say on *Oculus Quest 2* and *HTC Vive Pro 2*. In addition, online virtual environments can typically be accessed through a variety of VR devices, which differ in terms of software architectures. This suggests that different users can perceive the same virtual environment and objects, despite a difference of underlying digital entities in their respective hardware — this is what McDonnell

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

and Wildman (2019) call the "cross-play problem." For these reasons, among others, we think that SVD is not a promising option to defend virtual realism. This is also Chalmers' view (see Chalmers, 2019, p. 470).

WVD, by contrast, can account for the multiple realizability of virtual objects. Indeed, to say that a virtual object depends on a digital object effectively makes it a higher-order entity, which may be realized by a number of different digital bases — just as one and the same statue can depend on a variety of distinct physical arrangements. However, WVD is not free of trouble either. As Peter Ludlow (2019) remarked, the *social entities* we find in virtual worlds (e.g., clubs, teams, institutions, money, etc.) invalidate the claims of WVD. Consider, for instance, a virtual organization in a given VR environment, e.g., a criminal gang operating in *VR Chatroom*. It seems that the very same gang could migrate to a non-virtual environment and continue its illegal activities there. The gang might also survive the servers of *VR Chatroom* shutting down. If that's right, the nature of the virtual gang is not *uniquely* and *wholly* determined by underlying digital objects. This suggests, more generally, that virtual objects need not all be grounded uniquely in data structures. Some, like social objects, must also partly depend on human intentions or on some other mental factor.

WVD, understood as a narrow dependence of the virtual on the digital, is therefore unsatisfactory as well. Chalmers grants as much:

> One clear version of virtual digitalism might rely on [constitution] alone, saying that virtual objects are wholly constituted by, or grounded in data structures. In fact this is not my view. […] Virtual worlds involve not just data structures but also human users, and many virtual objects and virtual properties are grounded in part in the minds of the human users. […] So virtual objects need not be grounded in data structures alone. […]
>
> I suggest the following: virtual objects and properties are grounded in data structures and mental objects and properties. (Chalmers, 2019, p. 455)

On this revised version of WVD, virtual objects are no longer seen as depending uniquely on digital objects such as data structures. The relevant dependence relation now includes, in addition, some reference to the mental, i.e., to intentions, mental properties, and propositional attitudes. This gives us the following refined version of WVD:

WVD* Virtual objects are not identical to, but depend on, digital objects *and* mental objects/properties.

The view that we shall explore in this article is effectively a specification of WVD*, as it simply details what type of ontological dependence is involved here (viz., constitution). Our view, call it the "Constitution Theory of Virtual Objects" (CTVO), can be expressed as follows:

CTVO Virtual objects are *constituted* by digital objects, when certain favourable circumstances (which centrally involve mental states) occur.

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

We think that CTVO is quite in line with Chalmers' preferred version of virtual realism. This shouldn't come as a surprise. After all, he is the one who first brought up the notion of constitution as a potential way to elucidate the nature of virtual objects. Note, however, that we shall go further here than Chalmers himself ventured. Chalmers did not specify how the notion "constitution" should be understood exactly. Nor did he examine what is entailed by the claim that virtual objects are constituted by digital objects. In fact, Chalmers seems quite happy to leave the relation of ontological dependence of the virtual to the digital unspecified, since he also occasionally speaks of "grounding" rather than "constitution."

Our goal, in the rest of this article, will therefore be to flesh out the constitution view of the virtual, which Chalmers has merely sketched, and to show why it is a good option for those willing to endorse virtual realism. This, however, requires some explanation of what exactly the notion of constitution is supposed to be, and of the metaphysical role it can play. We turn to this now.

3. Constitution Theory

3.1. The General Idea

The notion of constitution can be approached through the infamous puzzle of a statue, "Goliath," and the lump of clay from which it is made, "Lump." At first glance, it is tempting to say that Goliath and Lump are just one and the same thing. After all, Goliath and Lump are co-located: they occupy a same region of space. Goliath is located exactly wherever Lump is, and vice versa. They also seem to share the same properties. For instance, if Lump weighs 20 pounds, has a humanoid shape, or was exquisitely carved by an Italian sculptor, the same holds for Goliath.

This identity claim, however, is problematic. Indeed, Goliath and Lump differ in *historical* and *modal* properties. Lump may have existed at a time when Goliath did not, namely, before the sculptor got to work. Moreover, Lump would survive being squashed into a flat disc, whereas Goliath would not. This difference of properties, by Leibniz's Law, entails that Lump and Goliath are numerically distinct entities. Therefore, it seems that this case forces one to admit the possibility of *spatially coincident objects* — i.e., of numerically distinct objects co-existing in the same place at the same time.

While some find such spatially coincident entities repugnant, a vast number of metaphysicians simply bite the bullet and accept the conclusion of the previous argument.[2] They will say that Lump and Goliath co-exist in a same place at the same time, but that they nevertheless are distinct entities, insofar as they differ in terms of certain (historical, modal) properties. This is often called the "constitution view." The point is also frequently expressed under the slogan that "constitution is not identity." Constitution differs from identity in that it is asymmetric (the lump constitutes the statue, but not vice versa), irreflexive, and meant to capture a certain kind of ontological *dependence*: the statue depends on the lump for its existence and properties, but not vice versa.

[2] See Wasserman (2021) for an overview.

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

It is important to stress that constitution is not simply a notion used to solve the odd metaphysical puzzle. Some, such as Baker (2002, 2007), see it as a ubiquitous phenomenon. Constitution would be the relation that is involved, for instance, between lumps of clay and statues, pieces of paper and bank notes, chunks of wood and tables, collections of atoms and molecules, DNA and genes, collections of cells and organisms, to mention but a few examples. Another crucial point is that constitution, on Baker's view, differs from *composition*, insofar as it is *not* a mereological relation: when x constitutes y, x is not a part of y.[3]

Thus, constitution can be seen as a relation that is halfway between identity and separateness. The constituted and constituting things are not identical to one another, but neither are they entirely separated nor independent from one another. Constitution, Baker says, may be thus understood as a form of "unity without identity." She defines it more precisely as follows:

> The fundamental idea of constitution is this: when a thing of one primary kind is in certain circumstances, a thing of another primary kind — a new thing, with new causal powers — comes to exist. (Baker, 2007, p. 32)

In this quote, we find the central tenets of Baker's account. Constitution, she proposes, is a relation holding between things of differing *primary kinds*, when certain *favourable circumstances* occur, and such that the constituted entity possesses *new causal powers* vis-à-vis the constituting entity. Let us examine these central notions in turn.

3.2. Primary Kinds

For Baker, constitution relates to things of different *primary kinds*. The notion of primary kind, which is Aristotelian in spirit, denotes a sortal property corresponding to what a thing is most fundamentally or essentially.[4] Baker explains this as follows:

> For any x, we can ask: What most fundamentally is x? The answer will be what I call x's "primary kind." Everything that exists is of exactly one primary kind — e.g., a horse or a passport or a cabbage. An object's primary kind goes hand in hand with its persistence conditions. (Baker, 2007, p. 33)

A thing's primary kind does two things.

First, it specifies what that thing is most fundamentally. What you are, most fundamentally, is arguably a person and not an animal, a philosopher, or a hunk of matter. If you're an animalist (see Olson, 2007), you'll say that your primary kind is *animal*. To take another example, the primary kind of the entity I am currently typing on is *computer*, rather than *object*, *artifact*, *black computer*. Examples of primary kinds include: *car, tree, stone, ID card, statue, quark,* and so on.

[3] See Baker (2007, Chapter 9).

[4] Baker assumes that everything belongs essentially to exactly one primary kind, and that primary kinds come with principles of individuation. Her view is reminiscent of Wiggins' "sortal essentialism," according to which a thing's principle of individuation is essential to it. For a discussion, see Mackie (2006, Chapter 8).

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

Second, a thing's primary kind determines its persistence conditions, in the following sense: if x has K for its primary kind, then x could not fail to be of kind K and continue to exist. Thus, *person* is a primary kind, since I arguably stop existing if I cease being a person. However, *teacher* is not a primary kind, since I can keep existing while ceasing to be a teacher. For Baker, everything essentially has one primary kind or another, each of them being associated with different persistence conditions.

3.3. Favourable Circumstances

The constitution relation, according to Baker, also crucially relies on the notion of "favourable circumstances." Favourable circumstances work as necessary conditions which "trigger" constitution. No general recipe can be given for such circumstances, which vary from case to case (Baker, 2007, p. 160). In the particular example of Lump and Goliath, they include the physical properties of the clay, the intention of the sculptor, and plausibly the social and cultural conventions that make artworks possible. Only in such circumstances does a lump of clay come to constitute a statue. By contrast, the circumstances favourable to the constitution of H_2O molecules by an oxygen atom and two hydrogen atoms are purely physical. In the case of social objects and artifacts (such as tables, flags, or money), the favourable circumstances will predominantly be intentional and cultural. At any rate, we see that for constitution to occur, the constituting object must find itself in certain special circumstances, without which the constituted object does not come into being.

3.4. Causal Powers

One final important point in Baker's account is that constitution brings into existence new things or new kinds of things, endowed with novel causal powers. Here's a suggestive example (Baker, 2007, p. 32): take a traffic sign made from a piece of metal. When this piece of metal is in certain peculiar circumstances — painted red with white marks spelling "S-T-O-P" on its surface, placed in a suitable location and in an environment that has certain laws and regulations, etc. — a new thing, namely a stop sign, comes into existence. The piece of metal has certain causal powers. For instance, it may reflect the 2 p.m. sun in such and such fashion, deflect a football thrown at it by children, etc. The constituted entity — the traffic sign — inherits these powers. Yet, the traffic sign has causal powers different from those of the piece of metal constituting it: most notably, that of directing drivers to stop at the location where it is placed. This is so even though the traffic sign does not exist separately from the constituting piece of metal. According to Baker, the constitution theorist can therefore accept a specific causation of constituted objects, whose causal powers do not reduce to that of their constituters.

3.5. Constitution and the Virtual

Now that we have examined Baker's theory of constitution in more detail, we can move on to our central point, which is to see how and to what extent it may apply to virtual objects. As stated initially, we suggest that virtual objects are constituted

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

by digital objects (data structures) when certain favourable circumstances occur. To paraphrase Baker, we get something like the following:

> When a digital object of a certain primary kind is in certain circumstances, a thing of another primary kind — a virtual object with new causal powers — comes to exist.

This view, which is the one we'll develop here, raises three questions. First, what are the primary kinds of digital and virtual objects? Second, what are the favourable circumstances to the coming into being of virtual entities, and how do mental phenomena factor in these circumstances? Third, what are the differences of causal powers between digital and virtual objects? These are the issues we'll address in the rest of this article.

4. The Digital and the Constitution of the Virtual

As we have seen, constitution relates to entities of fundamentally different sorts, i.e., which differ in primary kinds. To apply Baker's view to virtual objects, then, we must start by determining what the primary kinds of virtual and digital objects are.

4.1. The Primary Kinds of Virtual and Digital Objects

Let's consider virtual objects first. What are the primary kinds of these constituted entities? On the face of it, virtual worlds are populated by a plethora of different kinds of virtual objects: virtual mountains, virtual trees, virtual lakes, virtual cats, virtual people, virtual guns, virtual money, virtual shops, and so on. The primary kind properties of virtual objects, we suggest, can be determined through an analogy with the kinds of non-virtual objects that surround us. Ludwig, my cat, has the primary kind property *cat*. Notre-Dame has the primary kind property *cathedral*. In just the same manner, some virtual objects will have *virtual cat* and *virtual cathedral* for primary kind properties.

This suggestion, however, raises three questions:

First, what licenses the introduction of *sui generis* virtual kinds? For instance, why say that an object that I encounter in VR has *virtual cat* for its primary kind, rather than simply *cat*? The obvious answer is that this object, contrary to non-virtual cats, depends for its existence a digital software and hardware basis, and could not be perceived nor interacted with without the proper VR headset. It will also fail to satisfy whatever biological or evolutionary criteria we take to be the discriminating mark of cathood. Cats and virtual cats, that is, have different identity and persistence conditions. In addition, and given the current limitations of VR, virtual objects are ontologically deficient in that they do not possess all the properties, affordances, and capacities of their non-virtual counterparts (see e.g., Juul, 2019). Virtual kinds, for these reasons, can normally be distinguished from their non-virtual counterparts as being irreducibly different kinds of entities.

Second, what allows VR users to identify a virtual object x as being a virtual F, i.e., an object of a given virtual primary kind? This is a complex issue. Most of the time,

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

this will be a matter of the virtual object's function, affordances, and causal role within the VR environment. Thus, a virtual key normally has the property of being of that kind because it can be used to lock and unlock doors, chests, etc. Likewise, an on-screen object will typically be identified as a virtual ball because it can, "do the ball-like things relevant for what we are trying to do" (Juul, 2021). But, in other cases, the object may be recognized as a virtual *F* simply in virtue of its appearance and superficial resemblance to non-virtual *Fs*. Thus, an object could be identified as a virtual cat merely because it looks like a cat. Some virtual objects are thus virtual *Fs* because they are recognized by users as F-like entities.

Third, what determines that *F* is a *primary* virtual kind? Baker provides a test for determining whether a kind is primary: "An object x has K as its primary kind only if: x is of kind K every moment of its existence and could not fail to be of kind K and continue to exist" (Baker, 2007, p. 35). For instance, Ludwig has *cat* for primary kind, for he could not lose the property of being of that kind without ceasing to exist. *Needy black cat* is not a primary kind, however, for Ludwig could stop being needy or being black without ceasing to exist. We think that the same story holds for virtual objects. A virtual cat has *virtual cat* as its primary kind as it could not lose that property of being of that kind without ceasing to exist. However, *sleeping virtual cat* is not a primary kind, for a virtual cat might stop sleeping without ceasing to exist.

Virtual objects, on our view, can therefore be of many different virtual primary kinds. Now, what about the *digital* entities from which virtual objects are constituted? Here too, an analogy with the physical world may be helpful. Consider an ordinary wooden table. According to Baker, there is a constitution chain going on here. The table is constituted by an aggregate of natural polymers, which is constituted by an aggregate of molecules, itself constituted by an aggregate of atoms, and so on until the bottom physical level, if any. For simplicity, suppose that any ordinary physical object is ultimately constituted by aggregates of atoms. If that is so, and given that constitution is transitive, whatever constitutes ordinary objects has ultimately *aggregate of atoms* for its primary kind.

Now, we think that things work quite similarly in the case of virtual objects. Take a virtual table in a given virtual world. The virtual table is constituted by a data structure, which is a digital object corresponding to a complex algorithmic structure. For simplicity, let us assume that data structures are ultimately constituted of bits (we are thus leaving aside the physical realization of these bits, which are certain material states and processes in circuits and transistors). According to Baker, "If we descend down any chain of constitution relations, sooner or later we will come to aggregates as constituters" (Baker, 2007, p. 35). We suggest that this is also the case here. A data structure is ultimately constituted by aggregates of bits, i.e., by series of 0s and 1s. Aggregate of bits and data structures should be distinguished. The former are unstructured mereological sums, while the latter are structured entities corresponding to sets of algorithms. Data structures are multiply realizable, to the difference of aggregates of bits. In virtue of their organization, they also have distinct causal powers. If that's so, aggregates of bits constitute data structures, and data structures constitute virtual objects in turn. Since constitution is transitive, the primary kind of the ultimate constituters of virtual objects is *aggregate of bits. Aggregate of bits* is

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

a primary kind, for a given aggregate of bits could not stop being of that kind without ceasing to exist.

4.2. The Constitution of Virtual Objects

Now that we have specified the primary kinds of virtual and digital objects, we may be more specific about what it means to say that a virtual object is constituted by underlying digital objects. Let us illustrate through the example of a virtual table in a given virtual environment. The virtual table, we claim, is ultimately constituted by an aggregate of bits. This means that:

1) The constituting and constituted entities have different primary kinds. The first has *aggregate of bits* for primary kind, while the second's primary kind is *virtual table*.
2) The aggregate of bits is in circumstances favourable to the coming into being of the virtual table. These circumstances are complex. They include physical processes in the hardware, computational processes at the software level but also, as we shall see, intentional or mental factors.
3) When the virtual table comes into being, a novel entity appears, with causal powers different from those of the underlying digital object. This is something that we'll also argue for later on.

On our view, the virtual table is thus distinct from the aggregate of bits constituting it. The aggregate is not a part of the table, but is its constituter. The aggregate and the virtual table form a unity without identity.

4.3. Digital and Virtual Objects Do Not Coincide

It is important, at this stage, to mention a significant difference between the cases of constitution discussed by Baker and our application of that theory to the realm of virtual objects. On Baker's view, constituted and constituting objects *spatially coincide*. The lump of clay, for example, occupies exactly the same region of space as does the statue. The same holds for a piece of paper and a bank note, an aggregate of atoms and a molecule, etc. However, this observation does not go for virtual and digital objects. When I play *Vader Immortal*, the virtual lightsaber waved by my avatar does not spatially coincide with the data structure and the aggregate of bits that constitutes the saber. The latter are located in the hardware, say in the circuits of my VR headset, whereas the virtual lightsaber is located in the virtual space of the virtual world.

The constitut*ed* virtual objects, then, do not share their spatial location with the constitut*ing* digital objects. Isn't this a problem, since other cases of constitution involve such coincidence? We do not think so. Despite the common assumption to the contrary, spatial coincidence does not seem to be a necessary requirement of constitution. Consider social entities such as NGOs and other kinds of organizations. Although Baker remains silent on this, it is possible to hold that such entities are constituted at least in part by their members. Yet, we cannot say that NGOs spatially

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

coincide with their constituters, even taken collectively. Indeed, it would be false to say that such entities are located at the sum of the spatial regions occupied by their members. If that's so, the usual contention that the unity of constitution consists in material (spatial) coincidence of the constituted and constituting object ought to be rejected. This point is made by Frank Hindriks, who embraces a constitution theory of organizations, while noting that "the condition of spatial coincidence is problematic for organizations […]. [T]he spatial location of a constituted object can be distinct from that of the constituting object" (Hindriks, 2013, p. 419). As such, virtual objects would not be the only kind of constituted entities that fail to spatially coincide with their constituters. One reply to the previous concern is thus that virtual objects wouldn't be exceptional in that respect.

Now, some might still balk at this extension of Baker's view. Constitution without coincidence, they could say, is no constitution at all. One reason to think so is that with coincidence, also goes the "inheritance of properties," which is an important feature of Baker's theory. On her view, constituter and constituted share most of their properties, including their sortal properties. Lump and Goliath, for instance, are both statues. The difference is that Goliath is so *essentially*, while Lump is so merely *accidentally*, in virtue of its constituting a statue. Lump, Baker says, has the primary kind *statue* "derivatively" (Baker, 2007, p. 38), whereas Goliath has it "non-derivatively." This "inheritance" of properties is what seems to underlie the notion of a "unity without identity" at the heart of constitution theory. Without coincidence, that aspect of the theory is gone. A virtual gun, for instance, does not share its properties with the underlying data structure. So, does our view still qualify as a version of constitution theory?

We grant that our account departs from Baker's in two significant respects. However, we think that the problematic features (i.e., coincidence and property inheritance) are not essential for making sense of constitution. Also, recall that our goal here is to draw on Baker's theory to shed light on the relation of digital and virtual entities, not necessarily to endorse all aspects of her view. Some may persist in thinking that the corresponding relation shouldn't be called "constitution." That is fine. Terminological issues aside, what we are after here is a relation of ontological dependence that is irreflexive, asymmetric, and transitive; that holds between entities of different (primary) kinds, that occurs given the presence of certain favourable circumstances, and that allows for emergence and multiple realizability. Even if Baker's theory weren't strictly applicable to virtual objects, it satisfies these desiderata, and thus gives us a good dialectical starting point.

5. Virtual Objects as Artifacts

Until now, we have focused on how digital entities are involved in the constitution of virtual objects. This, however, is only part of the story. As we have already seen, Chalmers admits that virtual objects are partly dependent on the mental, while Baker contends that mental factors play a part in the constitution of many entities, including artifacts. In line with this, we claim that mental states (intentions, beliefs, etc.) feature in the circumstances that are necessary for the coming into being or existence of virtual objects. To show why this is so, we'll now consider the artifactual

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

nature of virtual objects and explain why this makes them partly dependent on the mental.

5.1. Artifacts and Proper Function

The non-virtual world is full of artifacts, that is, objects that are intentionally produced by human beings in order to serve some given purpose. However, it is also populated by a myriad of natural beings like animals, mountains, and trees. These entities, unless we are living in a simulation, are not artifacts. The same does not go for virtual worlds, however. Virtual worlds are themselves artifacts, and they are only comprised of artifacts.[5] There are no natural beings or natural kinds in virtual worlds.

The artifactual nature of virtual entities, we think, has a crucial consequence: it makes them partly dependent upon the mind. This conclusion owes to a specific feature of artifactual primary kinds. Here's what Baker has to say on that matter:

> What distinguishes artifactual primary kinds from other primary kinds is that artifactual primary kinds entail proper functions, where a proper function is a purpose or use intended by a producer. […] The nature of an artifact lies in its proper function — what it was designed to do, the purpose for which it was produced. An artifact's proper function is an *intended* function. Since artifacts have intended functions essentially, they are [intention-dependent] objects: they could not exist in a world without beings with propositional attitudes. (Baker, 2007, pp. 51–52)

The proper function of an artifact, according to Baker, is always an "intended function," i.e., one that has been intentionally determined and assigned by the agent who conceived or produced the artifact.[6] As a result, artifacts are "intention-dependent objects": they depend on human intentionality. On this view, artifacts could not exist without beings with propositional attitudes (beliefs, desires, intentions). In a world without beliefs, desires, or intentions, there would be no cars, churches, credit cards, or tables, but only particles arranged in certain ways. Propositional attitudes, that is to say, are centrally involved in the "favourable circumstances" necessary for the constitution of artifacts.

[5] A reviewer asked what we should say about automatically generated virtual objects. For instance, someone could create a virtual forest by using a pre-made template, without designing any of the particular virtual trees composing the forest. In such a case, we think that the virtual trees can still be considered as artifacts, insofar as they are derivatively produced by human intentionality. More importantly, the virtual trees would still have an intended proper function in this scenario. The fact that something is generated automatically, indeed, does not mean that it is created *unintentionally*. Although more would have to be said regarding the case of automatic generation, we do not think that it threatens the view we develop below.

[6] Baker is not alone in thinking that the identity of an artifact crucially depends, whether directly or not, on the intention of its creator (see e.g., Evnine, 2016; Thomasson, 2003, 2007). A less popular view gives precedence to users' intentions over those of the artifact's creator (Preston, 2013). Our discussion of virtual artifacts will be framed in terms of Baker's author-intention-based account, though we reckon that it is vulnerable to certain issues (for a discussion, see Koslicki, 2018, Chapter 8).

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

This account of artifactual kinds, we think, can help us understand the nature of virtual objects. Given that all virtual objects are artifacts, they too are intention-dependent objects that partly depend on the mental for their existence. Put otherwise, the constitution of virtual objects requires something else than the mere presence of digital objects. Mental states or propositional attitudes must feature in the favourable circumstances required for the constitution of virtual objects.

To illustrate, take a virtual screwdriver in a VR environment. This virtual object is constituted by an aggregate of bits. But what makes this aggregate constitute a virtual screwdriver is also a collection of propositional attitudes that determine its proper function. The virtual screwdriver must have been intended to be usable to screw and unscrew things in the virtual world. This particular functional role is no less essential to the existence of the virtual screwdriver than the underlying digital object, viz., the aggregate of bits. This also means that, in a world without intentions, propositional attitudes, or mental states, there could be no virtual screwdrivers.[7] The virtual screwdriver, that is, requires for its existence a background of mental states, which are part of the "virtual-screwdriver-favourable" circumstances.

The same story, of course, holds for all other virtual objects, since all of them qualify as artifacts involving one function or another. For that matter, we agree with Ludlow (2019, p. 365) when he contends that even non-social virtual objects such as virtual sushi partly depend on the mental for their identity, as they require intentions or "social consensus" to be identified as such. Chalmers is also sympathetic to a view of this sort:

> … part of what it is to be a table is to be treated as a table, or to be disposed to be treated as a table, or to be designed as a table, or something like that. […] On this picture, both being a table and being a virtual table will be partly grounded in the mental. (Chalmers, 2019, p. 459)

5.2. How Virtual Objects Depend on the Mental

As just made clear, we claim that virtual objects partly depend on the mental, in the sense that mental factors are centrally involved in the circumstances necessary to the constitution of virtual objects. That said, there are different ways in which something may depend on the mental. Drawing on Amie L. Thomasson (2005), we may differentiate *direct* and *indirect* dependence on the one hand, and *rigid* and *generic* dependence on the other.

Consider the first distinction. *X directly* depends on the mind iff *x* depends uniquely on intentional acts, without the mediation of another mind-external entity. For instance, mental images as they occur in dreams or reveries depend directly on the mental. By contrast, *x mediately* depends on the mind iff its dependence on intentional acts is mediated by another mind-external entity. Fictional characters, at least on Thomasson's (1999) view, are dependent entities in this sense, insofar as their existence requires, in addition to mental entities, various mind-external "supporting" entities (e.g., physical copies of a novel).

[7] In fact, we think that digital entities would not exist either in such a world, insofar as they can also be seen as artifacts that partly depend on mental states or intentions for their existence.

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

Now, let's turn to the other distinction. We can say that *x rigidly* depends on the mind iff *x* depends for its existence on a specific intentional act; and that x *generically* depends on the mind iff *x* depends on some intentional act of a certain sort. An example of the former case might be private mental images. The cold beer that I imagine after a long and warm day rigidly (and directly) depends on my intentional act, as it wouldn't exist without my imagining specifically. In cases of generic dependence, by contrast, the dependent object will tend to be an intersubjective entity whose properties result from a collection of intentional acts and public practices. In that case, the dependent object does not depend on any specific act of consciousness for its ongoing existence. Plausible candidates for this type of dependence include artifacts, along with social entities such as money. While these generically depend on subjectivity for their ongoing existence, they remain independent of any given particular act of consciousness.

We think that virtual objects indirectly and generically depend on the mental. I can't alter a virtual object at will, because it does not depend only or wholly of me and my conscious acts. Virtual objects only indirectly depend on the mind, as their dependence to intentional acts is mediated by entities that are external to the mind (i.e., digital objects). In addition, virtual objects only generically depend on the mind, as they depend on collective intentions and propositional attitudes rather than on individual and particular intentions or mental states. As such, our suggestion that virtual objects depend on the mental does not make them mere subjective projections or purely intentional entities.

5.3. Function and Appearance

We now have clarified what we mean when we say that virtual objects are constituted by digital objects placed in certain favourable circumstances. Given the role mental states play in such circumstances, regardless of the considered virtual object, we may consider that all virtual objects partly depend on the mental. Given the artifactual nature of virtual objects, this shouldn't come as a surprise. The core argument, once again, is this. First, all artifacts partly depend on the mental, as they are essentially tied to their intended proper functions. Second, all virtual objects are artifacts. Therefore, all virtual objects partly depend on the mental.

Before moving on, we'd like to consider a potential objection. We said, following Baker, that virtual kinds are essentially tied to intended proper functions. For instance, if a virtual object has been designed to be a movable barrier that allows entry or exit from an enclosure, then the object is a *virtual door*. However, some could object that some on-screen objects which *prima facie* qualify as virtual doors do not fulfill this particular role or function. Consider, for instance, the case, familiar to gamers, of a virtual door that cannot be opened by the user or interacted with by any means. In that case, we think that the relevant virtual object might actually not be a *virtual door*, strictly speaking. It is rather a *virtual door image*, i.e., a decorative virtual object whose proper function is to affect our perception in a door-ish way.[8] Virtual worlds, of course, are full of such purely decorative virtual objects.

[8] See Aarseth (2005), Grabarczyk and Pokropski (2016), Juul (2021).

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

Generally speaking, we suggest that these decorative virtual Fs are better thought of as virtual F-*images*, whose intended proper function is to produce F-ish experiences for users in normal VR conditions.

We have mentioned cases where virtual objects look like Fs, without having the usual intended proper function associated with Fs. But, conversely, some virtual objects may have the intended and proper function of Fs without having the typical appearance associated with Fs. For instance, a virtual sword might be reprogrammed to launch rockets, effectively making it a virtual rocket launcher, which happens to have the visual appearance of a sword. In such a case, and supposing that the virtual object cannot be used as a sword any longer, the newly acquired function takes over the graphical skin to determine the object's primary kind. The virtual sword, that is to say, has been replaced by a virtual rocket launcher. To consider another example, imagine that the developers of a VR app re-skin a virtual key to make it appear as a virtual pony to users. Assume that this involves absolutely no change of function on the part of the object: the "pony-key" can still be used to open virtual doors or virtual chests exactly as it did before. Here too, we suggest that function supersede appearance: the considered virtual object would still be a virtual key, even though it may look like a virtual pony.

The previous discussion hints that function and appearance may and do often come apart in virtual worlds. Some virtual Fs will look nothing like non-virtual Fs. This is why we can identify some objects as virtual libraries, even though they do not visually resemble non-virtual libraries in any shape or form. Some virtual Gs will look like non-virtual Fs without being virtual Fs. This is why some things are mere *virtual libraries images*, insofar as they visually resemble non-virtual libraries but fail to have the proper function or role usually associated with non-virtual libraries.

6. Digital and Virtual Entities

Now that we have detailed our constitution account of the virtual, we will see that three crucial differences between digital and virtual objects result from this account.

6.1. A Difference in Persistence Conditions

The first notable difference brought to the fore by CTVO regards the persistence conditions of virtual and digital objects. Let us start with digital objects. Data structures, as we have seen, are ultimately constituted by aggregates of bits. The important point is that aggregates of bits, just like other types of aggregates, comply with a principle known as "Mereological Essentialism" (ME). According to this principle, any part of an object is essential to existence. Applied to the digital realm, this view means that any part of an aggregate of bits is essential to the aggregate: add or remove any part (i.e., any 0 or 1), and you destroy the corresponding digital object.

Each part of an aggregate of bits, then, is essential to it. This, crucially, does *not* entail that the same data structure cannot be constituted by different arrays of bits. Indeed, ME is a claim about objects and their *parts*; not about object and their *constituters*. Recall that constitution, on Baker's view, is not a mereological relation.

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

The lump of clay is not a *part* of the statue, no more so than the aggregate of bits that constitutes a data structure.

Now, and to the difference of aggregates of bits, virtual objects do not comply with ME. Intuitively, we think that virtual objects can undergo (at least some) changes without ceasing to exist. For instance, I may add a door to my virtual house or change the outfit of my avatar without thereby destroying these virtual objects. Put otherwise, a virtual object can gain or lose virtual parts without ceasing to be the *same* virtual object. The persistence of virtual objects, for that matter, isn't solely determined by the identity of their parts. This point was also stressed by Chalmers: "a virtual statue is not exactly the same as a structure of bits. Bits may change and the statue may remain. The statue may be destroyed and the bits may remain" (Chalmers, 2022, p. 195).

Virtual objects can therefore survive changes of parts. But then, what else is involved in the persistence conditions of virtual objects? The reply here will depend on the specifics of one's theory of artifacts. On most views, such as Baker's, proper function plays a key part in the identity of artifacts. A natural thought would be to say that a virtual object persists just as long as it can perform its intended proper function. A virtual screwdriver, for instance, would persist as long as it can do what it was designed to do, viz., screwing and unscrewing things in the virtual environment. Some alterations brought to that object will be inconsequential to its function, and will therefore be identity-preserving. For instance, one may change the shape or colour of the screwdriver's handle without destroying it as a virtual object. It will be the *same* virtual screwdriver. If the screwdriver was reduced to (virtual) bits or dust and couldn't perform its function any longer, however, that virtual object would be destroyed. On this view, the proper function of a virtual object is therefore a necessary condition for its persistence. (It need not be a sufficient one, though, for other properties might be deemed necessary to the persistence of the relevant virtual object. For instance, it is arguable that the virtual character Darth Vader in *Vader Immortal* would not survive a drastic alteration of its skin or front-end features, even if the rest of its in-game behaviour was left untouched.)

The previous proposal about the persistence of virtual objects, while intuitive, is nonetheless quite objectionable. A disassembled phone sent for repair has stopped working, but it is still a phone, intuitively. So, why couldn't virtual objects survive, similarly, the loss of their proper function, at least in some cases? Couldn't a virtual screwdriver be malfunctioning, for example, because of a bug, and yet still be a virtual screwdriver? Some proponents of intention-based accounts of artifacts are happy to say that a broken artifact can still be a member of its artifactual kind. Simon Evnine, for instance, contends that "there can be broken Ks […] that are still Ks" (Evnine, 2016, p. 126). Baker argues for something similar, when writing that "an artifact may survive some malfunctions (the brakes can be fixed) but not others (the gas tank exploded and blew up the car)" (Baker, 2007, p. 57). If we allow, along these lines, that some cases of malfunction preserve the identity of the artifact, a buggy virtual screwdriver might still be a virtual screwdriver. But this, of course, requires an account of the conditions in which malfunction is persistence-compatible. Various candidates come to mind here. Perhaps the virtual screwdriver can survive the loss of its function if it *was* working at some point. Or, as long as it could be repaired.

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

To fully spell out the persistence conditions for virtual objects would require no less than a comprehensive theory of malfunction. We can't address that complex issue here, and so postpone it to further work on virtual objects. Fortunately, the lesson to take home from constitution theory is independent of these matters. The important point is that the persistence of virtual objects requires neither the diachronic identity of their parts nor the sameness of underlying digital objects. Thus, our view accounts for the different persistence conditions of digital objects (aggregates of bits) and virtual objects. This difference is explained by a difference in primary kind.

This approach also has the virtue of making sense of the much discussed multiple realizability of virtual objects. Just as one and the same table or flag may be constituted by different aggregates of matter at different times; the same virtual entity may be constituted by different digital entities — and aggregates of bits — at different times. And, just as the flag of a certain country can be constituted by various different material substrates, a virtual object can be constituted by a number of distinct data structures, as evidenced by cases of porting or cross-play (McDonnell & Wildman, 2019).

6.2. A Difference in Nature

A second important aspect of the constitution view is that it brings out a difference of nature between virtual and digital objects. As initially stressed, digital objects may be considered in two different ways. On the one hand, they can be seen as abstract types. This would make them non-spatio-temporal entities, just like propositions, numbers, or mathematical entities. On the other hand, and as Chalmers prefers to say, digital objects can be seen as the concrete instantiations of abstract types. On that view, each digital object is ultimately physically realized as electrical impulsions in the integrated circuits of servers or computer systems.

Now, what about virtual objects? Are these entities abstract or concrete? The answer will vary depending on the form of virtual realism advocated. Those who consider that virtual entities just *are* digital entities (e.g., data structures) will probably see them are concrete physical entities, since digital entities can ultimately be mapped onto certain physical processes in the hardware. This can't be our view, however, since we reject that identity claim in favour of a dependence claim. That said, it seems to us that CTVO leaves at least three options open regarding the nature of virtual objects.

The first option would be to say that virtual objects are abstract, but in a peculiar way. On a traditional or so-called "Platonist" account, abstract entities are seen as eternal and mind-independent, and hence uncreated. Sets, numbers, or propositions are frequently seen as abstract in this sense. However, various authors have argued that there are such things as non-eternal, mind-dependent, and hence created abstract entities, or in short, "abstract artifacts." This type of view has been used to account for fictional characters (Thomasson, 1999), repeatable artworks (Mag Uidhir, 2012), and software (Irmak, 2012). This kind of view might be adapted, *mutadis mutandis*, to virtual objects.

A second possibility would be to claim that virtual objects are *sui generis* entities, in that they are neither concrete nor abstract. This line of thought might be motivated

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

as follows. Virtual objects are not abstract entities in the traditional (Platonic) sense, since they exist within a simulated space-time in which they possess causal powers. Yet, they are not concrete physical objects either, as they do not exist in physical space and do not have any causal powers there (a virtual fire will cause no physical damage). This view would stress that the virtual space inhabited by virtual objects is different from the physical space in which their constituters (i.e., digital objects) are found. If that is right, virtual objects would have their own, distinctive nature.[9]

The third option would be to say that virtual objects are a peculiar type of concrete entities. On a traditional view, to be concrete is to exist in space-time. This is entirely compatible with the suggestion that there are immaterial concreta. Possible candidates would be ghosts (if there are or could be any) or, less exotically, some social entities such as institutions. Virtual objects too could belong to that category. The fact that we can't physically stumble on virtual stones doesn't necessarily exclude virtual objects from the realm of concreta. This only shows that these entities aren't *material*, not that they are non-concrete. To defend this view, we'd have to say that virtual objects exist in space-time. This isn't untenable. After all, we do perceive these entities when we put on our VR headsets. They may not be made of atoms and quarks, but they have a beginning and an end in time. So, they may be taken to exist in physical space-time. Virtual objects also have a virtual spatio-temporal location, whatever the relation of that virtual space-time to physical space-time may be.

No matter which option we choose, CTVO highlights a difference in nature between virtual and digital objects. Whereas digital objects can be seen as concrete and material entities, virtual objects are either abstract entities of a peculiar type, or non-concrete and non-abstract, or concrete but immaterial. Note that each of previous alternatives can explain why virtual objects do not coincide spatially with their constituters (viz., digital objects). The explanation, in the second case, is that digital and virtual objects occupy different spaces, whose locations cannot overlap. In the first and third case, coincidence is proscribed because virtual objects aren't material. Also note that all of these views are compatible with Chalmers' contention that the virtual is real. The claims that x is an abstract artifact, that x is neither concrete nor abstract, or that x is concrete but immaterial, do not entail anything regarding the reality or fictionality of x.

6.3. A Difference in Causal Powers

The third important aspect of the constitution view is that it helps to make sense of why virtual and digital objects differ in terms of causal powers. Where Chalmers (2017, pp. 318–319) seems happy to say that virtual causation is ultimately a matter of physical causal interactions at the hardware level, Baker's constitution theory affords us an extra resource here, insofar as it credits constituted objects with irreducible causal powers.

Seen as concrete realizations, digital objects are physical entities that possess physical causal powers. Such physical causation is the one at work within the hardware. However, it must not be conflated with the one existing within virtual words, when virtual objects causally interact with one another. It is true that such virtual

[9] For a related view, see Aarseth (2005).

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

causality is ultimately underlain by physical causal interactions within the integrated circuits of servers or computer systems. Still, the two causal levels seem irreducible.

To see why, imagine that I play a game of tennis table in a virtual world. When my virtual racket hits the virtual ball, the latter is propelled toward the right-hand side of the virtual tennis table. It is true that there are digital and ultimately physical causal processes underlying such a virtual event. Yet, these physical processes within the hardware simply are not the same as the event taking place within the virtual world. This is so simply because no event such as the trajectory of the ball from A to B occurs within the physical world. In the physical world where digital objects are found, there are no virtual racket, table, and ball, but only electrical impulses happening in some piece of hardware. While virtual causality is ultimately underlain by physical causal interactions within the hardware, there is no necessary isomorphism between the two causal levels.[10] This is one reason to think that virtual objects have their own causality, which is irreducible to that of the objects that constitute them.

Here, the analogy with what Baker (2007, pp. 97–119) calls "commonsense causation" proves enlightening. According to her, ordinary macroscopic objects have distinctive causal powers, irreducible to those of their constituters. As she writes:

> [With constitution], new things of new kinds, with new kinds of causal powers, come into being. An organism — but not the aggregate of cells that constitutes it — can eat its prey. A flag — but not the aggregate of pieces of cloth — may cause a veteran to cry. (Baker, 2004, p. 100)

This commonsense causation, for Baker, is linked to the primary kinds of objects, insofar as "[a]n object's primary-kind property […] confers on the object causal powers that cannot be manifested at lower levels" (Baker, 2007, p. 112). It is the object whose primary kind is *flag* that causes the veteran's tears; not the piece of cloth. We submit that the same thing holds for the virtual racket and the virtual ball. It is the object whose primary kind is *virtual racket* that causes the ball's propulsion, not the aggregate of bits. And, likewise, it is that particular virtual event (the virtual ball moving from A to B) that causes the excitement of the player, and not the underlying digital event (the electrical processes in the integrated circuits).

Constitution theory, equipped with the notion of primary kind, can therefore help make sense of causal interactions within virtual worlds. On this version of virtual realism, virtual objects have their own causal powers, irreducible to that of their constituters. Just as Baker suggests that there is a commonsense causation irreducible to that of the constituters of ordinary objects, we suggest that there is a virtual causation irreducible to that of the constituters of virtual objects.

7. Conclusion

We have explored the features and implications of the constitution approach to the virtual. To conclude, we would like to review five merits of CTVO, which make it, in our eyes, an attractive and promising version of virtual realism.

[10] See McDonnell and Wildman (2019, p. 377) for a parallel point.

https://doi.org/10.1017/S0012217324000155 Published online by Cambridge University Press

First, CTVO is *prima facie* plausible. Just as a statue and the clay from which it is made, an aggregate of bits and the associated virtual object differ in terms of properties and persistence conditions. CTVO does not take this difference as a mysterious and brute fact, but explains it as stemming from a difference of primary kinds.

Second, CTVO accounts for the manifest differences in causal powers between digital objects and virtual objects. This is because constitution brings about novel and irreducible causal powers.

Third, CTVO responds to Claus Beisbart's objection to Chalmers that

> claims about grounding or the constitution of objects do not answer the question of what the objects are. To know what certain objects are we should at least be given the category they belong to as well as some of their essential features or at least the types of features they must essentially have. (Beisbart, 2019, p. 320)

CTVO specifies what virtual objects are and how they differ from digital objects. It specifies the various kinds to which virtual objects may belong. It also has the potential of explaining which features are essential to these entities (this issue, as we noted, will depend on one's metaphysics of artifacts and malfunction).

Fourth, CTVO resolves the problem of multiple realizability and thus escape the worries faced by SVD. Again, this is so because constitution is no mereological relation. The constituting objects may change without affecting the identity of the constituted object.

Fifth, CTVO explains why and how virtual objects partly depend on the mind, just like other types of artifacts and social entities. This, importantly, does not threaten their reality nor make them mere projections. It is therefore a promising view for virtual realists.

Acknowledgements. Ancestors of this paper were presented at the "Virtual and Augmented Realities" conference in Collège de France on June 5, 2023 and at the University of Neuchâtel on November 1, 2023. We want to thank five anonymous reviewers for helpful comments. Our gratitude also goes to Jill Flohil for her invaluable help with the revision of this manuscript. One of the authors of this research (Alexandre Declos) was supported by the Swiss National Foundation project n°100015_212493: "An Ontology of Production, Products, and By-Products."

Competing interests. The authors declare none.

References

Aarseth, E. (2005). Doors and perception: Fiction vs. simulation in games. *Proceedings of the Digital Arts and Culture Conference 2005* (Copenhagen, Denmark). https://doi.org/10.7202/1005528ar

Baker, L. R. (2002). On making things up: Constitution and its critics. *Philosophical Topics: Identity and Individuation*, *30*(1), 31–52. https://doi.org/10.5840/philtopics200230112

Baker, L. R. (2004). The ontology of artifacts. *Philosophical Explorations*, *7*(2), 99–112. https://doi.org/10.1080/1386979041000169446 2

Baker, L. R. (2007). *The metaphysics of everyday life*. Cambridge University Press. https://doi.org/10.1017/cbo9780511487545

Beisbart, C. (2019). Virtual realism: Really realism or only virtually so? A comment on D. J. Chalmers's *Petrus Hispanus* Lectures. *Disputatio*, *11*(55), 297–331. https://doi.org/10.2478/disp-2019-0008

Chalmers, D. J. (2017). The virtual and the real. *Disputatio*, *9*(46), 309–352. https://doi.org/10.1515/disp-2017-0009

Chalmers, D. J. (2019). The virtual as the digital. *Disputatio*, *11*(55), 453–486. https://doi.org/10.2478/disp-2019-0022

Chalmers, D. J. (2022). *Reality+: Virtual worlds and the problems of philosophy*. Allen Lane.

Evnine, S. (2016). *Making objects and events: A hylomorphic theory of artifacts, actions, and organisms*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198779674.001.0001

Grabarczyk, P., & Pokropski, M. (2016). Perception of affordances and experience of presence in virtual reality. *Avant: Trends in Interdisciplinary Studies*, *7*(2), 27–44. https://doi.org/10.26913/70202016.0112.0002

Hindriks, F. (2013). The location problem in social ontology. *Synthese*, *190*(3), 413–437. https://doi.org/10.1007/s11229-011-0036-0

Irmak, N. (2012). Software is an abstract artifact. *Grazer Philosophische Studien*, *86*(1), 55–72. https://doi.org/10.1163/9789401209182_005

Juul, J. (2019). Virtual reality: Fictional all the way down (and that's OK). *Disputatio*, *11*(55), 333–343. https://doi.org/10.2478/disp-2019-0010

Juul, J. (2021). The game of video game objects: A minimal theory of when we see pixels as objects rather than pictures. https://www.jesperjuul.net/text/gameofobjects/essay.html

Koslicki, K. (2018). *Form, matter, substance*. Oxford University Press. https://doi.org/10.1093/oso/9780198823803.001.0001

Ludlow, P. (2019). The social furniture of virtual worlds. *Disputatio*, *11*(55), 345–369. https://doi.org/10.2478/disp-2019-0009

Mackie, P. (2006). *How things might have been: Individuals, kinds, and essential properties*. Oxford University Press. https://doi.org/10.1093/0199272204.001.0001

Mag Uidhir, C. (Ed.) (2012). *Art and abstract objects*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199691494.001.0001

McDonnell N., & Wildman, N. (2019). Virtual reality: Digital or fictional? *Disputatio*, *11*(55), 371–397. https://doi.org/10.2478/disp-2019-0004

Olson, E. (2007). *What are we? A study in personal ontology*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195176421.001.0001

Preston, B. (2013). *A philosophy of material culture: Action, function, and mind*. Routledge. https://doi.org/10.4324/9780203069844

Thomasson, A. L. (1999). *Fiction and metaphysics*. Cambridge University Press. https://doi.org/10.1017/CBO9780511527463

Thomasson, A. L. (2003). Realism and human kinds. *Philosophy and Phenomenological Research*, *67*(3), 580–609. https://doi.org/10.1111/j.1933-1592.2003.tb00309

Thomasson, A. L. (2005). Ingarden and the ontology of cultural objects. In A. Chrudzimski (Ed.), *Existence, culture, persons: The ontology of Roman Ingarden* (pp. 115–136). Ontos. https://doi.org/10.1515/9783110325621.115

Thomasson, A. L. (2007). Artifacts and human concepts. In E. Margolis & S. Laurence (Eds.), *Creations of the mind: Theories of artifacts and their representation* (pp. 52–73). Oxford University Press. https://doi.org/10.1093/oso/9780199250981.003.0004

Wasserman, R. (2021). Material constitution. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2021 ed.). Stanford University. https://plato.stanford.edu/archives/fall2021/entries/material-constitution

Cite this article: Bucchioni, G., Declos, A. (2024). The Constitution of Virtual Objects. *Dialogue* 1–21. https://doi.org/10.1017/S0012217324000155