



The philosophy of the metaverse

Melvin Chen¹

Published online: 1 August 2023
© The Author(s) 2023

Abstract

How might we philosophize about the metaverse? It is traditionally held that the four main branches of philosophy are metaphysics, epistemology, axiology, and logic. In this article, I shall demonstrate how virtual walt-fictionalism, a particular version of virtual irrealism, is able to offer a straightforward, internally consistent, and powerful response about the metaphysics, epistemology, and axiology (ethics) of the metaverse. I will first characterize the metaverse in terms of a reality-virtuality (RV) continuum and distinguish between virtual realism and virtual irrealism, before elaborating on the explanatory power and potential of virtual walt-fictionalism relative to the metaphysics, epistemology, and axiology (ethics) of the metaverse.

Keywords Metaverse · Virtual realism · Virtual digitalism · Virtual irrealism · Virtual walt-fictionalism

The metaphysics of the metaverse

The constitution question

What is the metaverse? The metaverse is often associated with the term ‘extended reality’ (or XR). XR, in turn, is often used to refer to a series of immersive technologies that encompasses virtual reality (VR), augmented reality (AR), and mixed reality (MR). Physical reality is the world that we experience as our ancestors would have experienced it, without the benefit of any XR, VR, and AR technologies. A reality-virtuality (RV) continuum may help us to make sense of VR, AR, and MR (Milgram et al., 1995).

Relative to the RV continuum in Fig. 1, the real environment at the left of the continuum consists solely of real physical entities and the virtual environment at the right of the continuum consists solely of virtual entities. Many sorts of mixed reality may be located between the real (leftmost) and virtual (rightmost) extremities of the reality-virtuality continuum (Milgram & Kishino, 1994; Milgram et al., 1995, 1999). Furthermore, metaverse technologies are typically deployed in a mixed-reality environment in which we find both real physical entities and virtual entities.

When the metaphysical nature of the metaverse is under investigation, we might ask whether or how virtual entities can be constituted by digital entities. This question may be termed the constitution question.

The object-event distinction

Entities (virtual or digital) include both objects and events. It is helpful to understand the object-event distinction in terms of a distinction between something’s enduring and something’s perduring (Lewis, 1986). An entity x endures iff it persists by being wholly present at more than one time. Conversely, an entity x perdures iff it persists by having different temporal parts or stages at different times, although none of x is wholly present at more than one time. We typically say that objects exist, whereas events occur. Our linguistic distinction appear to track a metaphysical distinction between objects enduring and events perduring (Hacker, 1982b, a).

We mean by virtual objects objects that are contained in virtual worlds and that we perceive and interact with when using virtual reality. Paradigmatic virtual objects include avatars (virtual bodies). We mean by digital objects data structures or bits (0 s or 1 s) and bit strings (strings of 0 s or 1 s). It should immediately be noted that these data structures are abstract computational objects that may or may not be implemented.

✉ Melvin Chen
melvinchen@ntu.edu.sg

¹ School of Humanities, Nanyang Technological University,
48 Nanyang Avenue, Singapore 639818, Singapore




Fig. 1 Simplified representation of the RV continuum (Milgram et al., 1995, p. 283)

The identity-dependence distinction

As we shall use the concepts of identity and dependence to describe relations between entities (virtual or digital), it will be important to clarify what precisely might be meant by these concepts from the outset. Since the concept of identity is more straightforward than the concept of dependence, let us commence with the easier task of elucidating the former before proceeding with the more difficult task of making sense of the latter.

According to Leibniz's Law (hereafter: LL; also known as the principle of the indiscernibility of identicals), entities that are identical cannot differ in any respect. According to the principle of the reflexivity of identity (hereafter: RI), any entity is identical to itself. Formally:

$$(RI) x = x$$

$$(LL) x = y \rightarrow (\phi(x) \rightarrow \phi(y)) \text{ (Enderton, 1972; Deutsch \& Garbacz, 2022)}$$

According to the standard account of identity, the identity relation may be characterized as a conjunction of the postulates RI and LL (formally: $RI \wedge LL$). Symmetry (formally: $x = y \rightarrow y = x$) and transitivity (formally: $((x = y) \wedge (y = z)) \rightarrow x = z$) are characteristic properties of identity that may be deduced from the postulates RI and LL. Since any relation that is symmetric, reflexive, and transitive is an equivalence relation, the standard account of identity tells us that the identity relation is simply an equivalence relation that satisfies LL.

By contrast, dependence relations come in various stripes: realization, concretization, grounding, and supervenience (to name but a few possibilities). The realization relation denotes some dependence relation between higher-level and lower-level or more fundamental properties or activities. At a first pass, we mean by a realization relation the following: an object x 's having some higher-level property or activity ϕ is realized by x 's having a lower-level property or activity ψ (Wilson & Craver, 2007).¹ The realization relation is

¹ Wilson and Craver (2007) argue that the concept of realization serves several masters: it is regularly invoked by metaphysicians and philosophers of mind to discuss relationships between the mind and

asymmetric: x 's having ψ realizes x 's having ϕ but not vice versa. After all, ϕ could be realized by different objects and different properties. For example, let us consider a mousetrap. The mousetrap is a device that consists of several parts: a base, a bait holder, a trigger or lever, an impact bar, a latch mechanism, etc. These different parts and their associated activities (viz. ψ_1, ψ_2, ψ_3 , etc) realize the mousetrap's behaviour of catching mice (viz. ϕ -ing). At the same time, the different parts of a mousetrap and their associated activities are sufficient though not necessary for mice-catching, since the behaviour of catching mice may be multiply realized by other devices (e.g. balance cages, pit traps, etc) (Wilson & Craver, 2007). Therefore, we can distinguish between identity (symmetric in nature, as might be inferred from the postulates RI and LL) and realization (asymmetric in nature).

The concretization relation is another distinct species of dependence relations that receives comprehensive treatment in the Information Artifact Ontology (hereafter: IAO) (Ceusters & Smith, 2015).² According to the IAO, an information content entity (hereafter: ICE) is an entity that is generically dependent on some material entity and that stands in a relation of aboutness to some entity.³ The ICE is therefore an entity that is about something in reality and it can migrate or be transmitted (e.g. through copying) from one entity to another. Entity x generically depends on y iff x exists, y exists, and for some universal Y , y is an instance of Y and (necessarily) if x exists then some Y exists.

A generically dependent entity is concretized in each case. Just as a DNA sequence is concretized in a specific pattern of nucleotides in a molecule and a sentence is concretized in a pattern of ink marks on a piece of paper, a data structure is concretized in the implementation of bit strings (strings of 0s or 1s) that can facilitate the efficient storage, manipulation, and retrieval of data. IAO covers both objects (endurants or continuants) and events (perdurants or occurments): we may deal equally with artifacts (e.g. databases, scientific publications) and the processes through which ICEs are created, understood, and communicated. Furthermore, a concretization relation is rooted in time and context: the sentence 'Obama was never President of the USA', written on a piece of paper in 2007, was true when it was written but false when read by an observer today.

Footnote 1 (continued)

the brain and appears equally central to the explanatory and investigative practices of cognitive scientists.

² IAO is an ontology of information entities, originally driven by work on the Ontology for Biomedical Investigations (hereafter: OBI).

³ 'Aboutness' corresponds to what is referred to by means of the terms 'reference' or 'denotation' and includes both linguistic reference and relations of cognitive or intentional directedness (Ceusters & Smith, 2015).

Table 1 The identity relation and the various dependence relations

Relation	Identity	Dependence			
		Realization	Concretization	Grounding	Supervenience
Symmetry	Symmetric	Asymmetric	Asymmetric	Asymmetric	Non-symmetric
Reflexivity	Reflexive	Irreflexive	Irreflexive	Irreflexive	Reflexive
Transitivity	Transitive	Transitive	Transitive	Transitive	Transitive
Nature of relata	Entities	Properties or activities	Information entities	Facts or entities	Sets of properties

Grounding is a dependence relation in which the grounding entities are prior to or more fundamental than the grounded entities. For example, the existence of a set is grounded in its members: the members of a set are prior to the set itself and each set obtains in virtue of its members. Furthermore, any entity x is fundamental or ontologically independent iff nothing grounds x (Schaffer, 2009, p. 373). The grounding relation is transitive: if x grounds y and y grounds z , then x also grounds z . Grounding is irreflexive, since nothing can ground itself. Grounding is asymmetric, since nothing can ground something that also grounds it.

Last but not least, the supervenience relation is yet another species of dependence relations. A set of properties Φ is said to supervene on another set of properties Ψ iff some difference in Ψ is necessary for any difference in Φ to be possible. In other words, no two entities x and y can differ with respect to Φ -properties without differing with respect to their Ψ -properties. Suppose that you are in a happy mental state. Your mental state has the property of happiness that arises due to specific properties of your brain (viz. neuronal firing patterns, release of neurotransmitter, etc). We say that the mental properties supervene on the physical properties: you cannot change the mental properties without first changing the relevant physical properties. Furthermore, if any of the relevant physical properties were altered, it is likely that a different mental state (e.g. anger, frustration, boredom) would be experienced by you.

The supervenience relation is reflexive, since for any set of properties Φ , there cannot be a Φ -difference without a Φ -difference. Supervenience is also transitive: if Φ -properties supervene on Ψ -properties and Ψ -properties supervene on Λ -properties, then Φ -properties supervene on Λ -properties. However, supervenience is non-symmetric. In certain instances, it is asymmetric: while mental properties may supervene on physical properties, physical properties do not supervene on mental properties. In other instances, it is symmetric: every reflexive case of supervenience is symmetric in a trivial sense. Unlike the identity relation (symmetric in nature), supervenience is non-symmetric. Unlike the realization relation (asymmetric in nature), supervenience is non-symmetric: it sometimes holds symmetrically.

Table 1 represents both the similarities and differences between the identity relation and the various dependence relations:

Virtual digitalism

Strong virtual digitalists maintain a constitution-as-identity relation between virtual objects and digital objects and virtual events and digital events. Strong virtual digitalism is equivalent to the first-approximation view in Chalmers (2019). In other words, virtual objects are simply identical to data structures that have been implemented and virtual events are simply identical to the physical processes associated with the implementation of these data structures. The relation between virtual and digital entities is symmetric, reflexive, transitive, and in accordance with LL.

Weak virtual digitalists, on the other hand, deny that the constitution-as-identity relation holds between virtual and digital entities. Weak virtual digitalism is equivalent to the second-approximation view in Chalmers (2019). Instead, they might say that virtual entities are at least in some sense dependent on digital entities. Constitution is not identity, since virtual and digital entities differ in important respects. If an appeal is made to the realization relation, then weak virtual digitalists might say that the higher-level properties or activities of virtual entities are realized by the lower-level or more fundamental properties or activities of digital entities. Conversely, she might say that data structures are generically dependent objects (continuants) that are concretized in specifically dependent objects (continuants) upon implementation and further that these data structures permit the concretization of virtual entities (a species of information entities). Alternatively, it might be held that virtual entities are grounded in data structures, bits (0 s or 1 s), and bit strings (strings of 0 s and 1 s). Last but not least, it could be asserted that the properties of virtual entities supervene on the properties of the digital entities that constitute them.

Virtual realism versus virtual irrealism

Virtual digitalism (strong or weak) defends both the reality of virtual entities and the existence of certain relations (viz. identity or dependence) between the virtual and the digital.

If virtual digitalism (strong or weak) is correct, then we will have good grounds to accept the reality of virtual entities: virtual objects genuinely exist in the sense that real physical objects exist and virtual events genuinely occur or take place. The view according to which virtual entities are real is known as virtual realism. Virtual realism implies a certain level of artifice relative to the RV continuum, since virtual entities in the virtual environment are as real as physical entities in the real environment. Perhaps we shall be better off speaking solely in terms of real virtual entities and real physical entities.

Virtual entities can be perceived, interacted with, and bring about certain effects within the virtual environment (e.g. video game, VR simulation, computer-generated world, etc). It is therefore convenient for us to describe virtual entities as having some form of existence within the virtual environment. Virtual realism appears compatible with this intuition. At the same time, virtual realism incurs a distinct set of ontological commitments and introduces the worry that the philosophy of the metaverse may result in a needlessly bloated ontology. According to Ockham's Razor, we ought to refrain from multiplying entities beyond necessity.⁴ Unless we have compelling reasons to posit the reality of entities, Ockham's Razor suggests that we have parsimonious grounds to resist the conclusions of virtual digitalism.

Lest we forget, there is also the competing intuition that while virtual entities can interact with other virtual entities, possess properties, and bring about a significant impact within virtual environments, they do not exist as physical entities do in the real world. Virtual realism has trouble dealing with this competing intuition. It is precisely because of the tension between the two competing intuitions that the term 'virtual reality' strikes us as neither a contradiction nor a mere tautology (i.e. saying the same thing twice) in rhetorical terms. In addition, there is a lack of consensus within this camp about the relations between entities (virtual or digital): are these relations identity relations or dependence relations? If these relations are dependence relations, are they realization, concretization, grounding, or supervenience relations?

Cross-play and distributed realization

Two scenarios should suffice to provide an illustrative sense of the problems that confront strong virtual digitalism. After all, once the reality of virtual entities is admitted, nothing prevents further questions being asked about the nature of the identity relation that holds between the virtual and the digital, how virtual objects get properly re-identified as one

and the same across time, how we distinguish between token instantiations of the same type, etc. Our first scenario concerns cross-play, while our second scenario concerns distributed realization. As we shall see, attempts to address these further questions tend to lead strong virtual digitalism on the royal road to metaphysical hairsplitting.

The cross-play phenomenon arises whenever users on multiple and distinct systems are able to interact within a single virtual environment. In an example from McDonnell and Wildman (2019), a VR application named *Frisbee* supports cross-play: multiple users can play virtually and throw a virtual frisbee around. A peer-to-peer network (without any mediating server to control the game) can be used to allow a user on a 32-bit HTC Vive and another user on a 64-bit Oculus Rift to engage with a single virtual object: the virtual frisbee. Given the radical differences between HTC Vive and Oculus Rift, we may equally conclude that there is no single digital object with which both users are engaged. Since virtual entities such as the virtual frisbee may be constituted in the absence of any distinct digital entities, the cross-play phenomenon appears to imply that virtual realism is false.

The strong virtual digitalist might counter that we have two virtual objects whose properties are being synchronized in real time via the peer-to-peer network, even though it might look like the two users are interacting with a single virtual object. We have one application on the HTC Vive, another application on the Oculus Rift, two suites of virtual objects that are specific to each application, and the peer-to-peer network that supports the synchronization of disparate objects across these devices. The proponent of strong virtual digitalism will have no issues with declaring that we have two distinct digital objects and two distinct virtual objects in the cross-play scenario.

It could be further urged that if one device (e.g. the Vive) stops working whereas the other (e.g. the Rift) keeps going, there is no reason why the Rift user would suddenly become unable to continue interacting with the virtual frisbee. However, if we accept an identity relation between virtual and digital entities, what is it in virtue of which we maintain that one and the same device-specific virtual frisbee (either the Vive-specific one or the Rift-specific one) endures, even as the number of nodes in the network, computing power, and amount of resources available for sharing vary from t_n to t_{n+1} within the peer-to-peer network? The strong virtual digitalist owes us an account of the conditions under which a virtual entity (e.g. the Vive-specific virtual frisbee, the Rift-specific virtual frisbee) at t_n is properly reidentified at t_{n+1} .⁵ Alternatively, we might say that the strong virtual digitalist owes us

⁴ For more information about the history of Ockham's Razor, see Brampton (1964), Maurer (1978).

⁵ In the metaphysics of personal identity, this is known as the reidentification question (Schechtman, 1996).

a criterion of diachronic numerical identity (i.e. a criterion of what makes a virtual entity one and the same as itself at different times). Absent this criterion and the claim that we have two virtual objects (viz. the Vive-specific virtual frisbee and the Rift-specific virtual frisbee) in the cross-play example does not seem any more persuasive than the claim that we have one virtual object (viz. a single virtual frisbee across the Vive and Rift devices) in the same example.

Our second scenario concerns distributed realization: a single virtual entity may be realized by, grounded in, concretized in, or dependent on multiple bodies of data across different machines. In other words, there is a single digital object whose parts are physically distributed across multiple machines (Beisbart, 2019; Ludlow, 2019).⁶ This phenomenon of distributed realization motivates the idea that one and the same virtual object may somehow be dependent on different digital objects (e.g. concretized data structures). The strong virtual digitalist might counter that although a virtual object is realized by, grounded in, concretized in, or dependent on multiple bodies of data across different machines, each token instantiation of this virtual object will still be identical to a particular digital object. In other words, while it might appear that the same type of virtual object is realized by, grounded in, concretized in, or dependent on different digital objects, the strong virtual digitalist may point out that this is no threat to the idea that token virtual objects are identical to token digital objects. On what metaphysical grounds, however, do we distinguish between a type and a token? Likewise, how do we identify several distinct tokens as being tokens of the same type? Last but not least, might the reidentification question not still be raised with respect to the token: what are the conditions in virtue of which a token virtual object might be properly reidentified across time?

The type-token distinction could allow strong virtual digitalism to preserve a one-to-one mapping between token virtual objects and token digital objects. Equally, the strong virtual digitalist could regard multiple bodies of data across different machines as multiple digital objects and preserve a one-to-one mapping by asserting the reality of multiple virtual objects in the distributed realization scenario. However, does strong virtual digitalism ultimately have any principled means of determining whether we have multiple digital objects (and therefore multiple virtual objects) or a

single digital object distributed across multiple machines (and therefore a single token virtual object) with respect to the distributed realization scenario? As with our cross-play scenario, it seems that we are led on the royal road to metaphysical hairsplitting relative to the distributed realization scenario.

Causal exclusion

Weak virtual digitalism is not without its own set of problems. An analogy may be identified between weak virtual digitalism and non-reductive physicalism. The latter is a view in the philosophy of mind according to which the mental supervenes on but does not reduce to the physical. However, it has been argued that non-reductive physicalism is inconsistent with the following three theses:

- T1: The physical domain is causally closed;
- T2: The mental has causal efficacy;
- T3: There is no systematic causal redundancy.

If we assume that the physical domain is causally closed (T1) and the mental has causal efficacy (T2), then it follows that either the mental is identical with something physical or there is systematic causal redundancy (\neg T3). This is known as the causal exclusion argument (Kim, 1992). By analogy, we may infer that there will be a causal exclusion problem for the weak virtual digitalist: if we assume that the physical domain is causally closed (T1) and virtual entities have causal efficacy (T2), then it follows that either the virtual is identical with something physical (e.g. concretized data structures) or there is systematic causal redundancy. Since digital entities do all the causal work, we have good reason to expect that there will be nothing left (causally speaking) for virtual entities to do (McDonnell & Wildman, 2019). This will count as bad news for weak virtual digitalists. After all and according to them, some degree of causal efficacy should typically be retained by virtual entities.

Virtual irrealism and virtual fictionalism

Virtual irrealism, on the other hand, denies that virtual entities are real. According to the virtual irrealist, virtual entities, despite appearances, do not genuinely exist and virtual events that take place in the virtual environment do not genuinely occur. Given the cross-play and multiple realization problems confronting strong virtual digitalism and the causal exclusion problem confronting weak virtual digitalism, certain challenges will have to be addressed by virtual realists who rely on at least some constitution relation between virtual and digital entities to defend the reality of virtual entities. Virtual digitalism (strong or weak) incurs ontological commitments, may result in a needlessly bloated

⁶ We may distinguish here between identity and sameness. A digital object y_1 could be an exact duplicate of another digital object y_1' , in which instance y_1 is the same as y_1' in the sense that there is a clone or a copy of y_1 . However, this does not entail that y_1 and y_1' are identical. Although their respective sequences of 0s and 1s are the same, these sequences may be stored in different physical locations and this means that they are not identical. If we destroy y_1 , y_1' need not completely disappear. I am grateful to an anonymous reviewer for having pointed out this important distinction between identity and sameness.

ontology, invites metaphysical hairsplitting, fails to square with our competing intuition about virtual entities, and may be resisted on grounds of parsimony unless they provide us with compelling reasons to posit the existence of virtual entities. These worries provide us with reasons to deny that virtual entities are real.

Some virtual irrealists may accept the existence of digital entities but deny the existence of virtual entities. Instead, they might assert that virtual worlds are fictional worlds. This species of virtual irrealism is known as virtual fictionalism. Although virtual entities do not exist, we may behave as if they do.⁷ This ‘as if’ aspect of fictionalism squares better with our competing intuitions that virtual entities can interact with other virtual entities, possess properties, and bring about a significant impact within virtual environments, yet do not exist as physical entities do in the real world. Virtual fictionalism entails that all talk of virtual entities having some form of existence within the virtual environment, though convenient or useful, is ultimately a mere *façon de parler*.

Not all virtual irrealists are virtual fictionalists. At least some virtual irrealists deny in equal measure that virtual worlds are real and that virtual worlds are fictional (Aarseth, 2007). Instead, virtual worlds are more analogous to dream worlds and thought experiments. Most virtual fictionalists reason about virtual worlds on the basis of claims that they take to be plausible about the world of video games (Juul, 2005; Tavinor, 2009; Meskin & Robson, 2012). Nonetheless, this point does not constitute a reason for rejecting virtual fictionalism out of hand. After all, if video game worlds are sufficiently representative of virtual worlds, then virtual fictionalists will have good reason to believe that whatever holds in virtual worlds will hold in video game worlds too.

Virtual walt-fictionalism

Virtual walt-fictionalism is a species of virtual fictionalism (McDonnell & Wildman, 2019, 2020). While it observes the central tenets of virtual fictionalism, it is also characterized (as its name is intended to suggest) by a certain reliance on Waltonian tools in the philosophy of fiction to make sense of virtual worlds. Virtual walt-fictionalism gives rise to a broadly Waltonian version of virtual irrealism.

According to the Waltonian theory of mimesis as make-believe, any work of representational art in which members of the audience present, imagine, or make-believe the world as the text represents it as being (e.g. the fictional world of Rodion Raskolnikov in Fyodor Dostoevsky’s *Crime & Punishment*) becomes a resource in the game of make-believe

(Walton, 1990; Woodward, 2014). For a work of art to count as a work of fiction, it should serve as a prop in a game of make-believe. Games of make-believe are guided and determined (at least in part) by props and principles of generation (parameters that prescribe what it is we are to make-believe). Props are items whose existence will be used to guide and determine the content of these games. Principles of generation are rules that prescribe what it is that we are to imagine in these games. In certain games (e.g. children’s games of make-believe) that are termed unauthorized games, the principles of generation may be ad hoc. In yet other games that are termed authorized games, these principles may be widely shared, stable, and publicly agreed upon (Walton, 1990; Frigg, 2010).

Principles of generation may be conditional (‘If there is a tree stump in a forest, then imagine that there is a bear where the tree stump is’), stipulative (‘Imagine that there is a bear in the forest’), or supplemental (‘If we are prescribed to imagine a bear in the forest, then imagine that it is big and hungry’) (McDonnell & Wildman, 2019). By analogy, virtual walt-fictionalists may identify digital entities (bits and bytes, screen images, sounds, haptic feedback, etc) as props. Once the relevant principles of generation are in place (e.g. ‘If you see a projected image, imagine that it is a fully 3-dimensional entity’), we should be able to engage in games of make-believe in virtual worlds. Whenever we engage with a particular virtual environment in an authorized manner, props help us to make-believe that virtual objects exist (although they do not actually exist). The moment we stop playing the game of make-believe, these virtual objects will cease to exist. Whereas virtual realists will maintain that virtual entities have causal efficacy in virtue of their reality (i.e., virtual objects genuinely exist and virtual events genuinely occur), the virtual walt-fictionalist holds that virtual entities fictionally have causal efficacy or causal powers (i.e., they are causally efficacious within the relevant games of make-believe).

Virtual walt-fictionalism has a lot working in its favour. In the first instance, it avoids ontological commitments to an overly large number of real entities (physical and virtual) and permits us to maintain an important ontological distinction between the real and the virtual (as observed in the RV continuum in Fig. 1). In the second instance, it sidesteps the cross-play and distributed realization problems confronting strong virtual digitalism (and concomitant metaphysical hairsplitting) by denying the identity relation between virtual and digital entities. In the third instance, it avoids the causal exclusion problem that plagues weak virtual digitalism by denying the dependence relation between virtual and digital entities. This denial renders it disanalogous with non-reductive physicalism. At the same time, it can explain how strong virtual digitalists mistakenly ascribe causal powers to virtual entities: virtual entities appear to have causal powers within

⁷ See Vaihinger (1923) for a precursor to modern fictionalism that foregrounds the ‘as if’ aspect of fictionalism.

the relevant games of make-believe. Last but not least, virtual walt-fictionalism recognizes the representational character of the virtual environment in a manner that virtual realists and non-fictionalist virtual irrealists do not. Virtual walt-fictionalism is a natural and powerful extension of the Waltonian approach to fiction and there are good reasons to believe that it can help us to make sense of both the fundamental nature of the metaverse and its ethical implications.

The ethics of the metaverse

The import-export and magic circle principles

If virtual walt-fictionalism is correct, then the metaverse may be regarded as a prop-driven game of make-believe. While the digital props (e.g. data structures, code, bits and bytes, images, sounds, haptic feedback, etc) are real, the virtual entities that they prompt or invite us to make-believe are not. Waltonian fictionalism already challenges the everyday usage of the term ‘fiction’ by applying it not just to novels, plays, and paintings but even to children’s games of make-believe involving tree stumps in the forest, dolls, or hobby horses. Virtual walt-fictionalism extends the application of the term to the metaverse. To disambiguate between the traditional usage of the term ‘fiction’ and the scope of reference in these instances, the term ‘walt-fiction’ may be used to pick out the relevant category of games of make-believe (Friend, 2008).

It has been argued that the import-export principle is our default fictional interpretation principle (Gendler, 2000). According to this principle, moral norms are true in a (walt-)fiction just in case they are true in the actual world. The import-export principle provides an explanation of how we learn moral norms from (walt-)fictions: when we are invited to make-believe certain norms that we had not previously considered but that we do not believe are false, then these norms could be exported from the (walt-)fictional world and into the actual world and endorsed in accordance with the import-export principle (Rosenbaum, 2016).

The import-export principle supports an epistemic connection between the actual world and the (walt-)fictional world. The magic circle principle, on the other hand, suspends this epistemic connection. According to the magic circle principle, from moral norms being true in the actual world, it by no means follows that they shall be true in a (walt-)fiction. Games appear to take place inside a magic circle in which ordinary moral rules and norms are somehow suspended (Huizinga, 1949). We tolerate a certain degree of aggressive conduct in sports such as boxing and rugby that would be considered physical assault or battery outside the magic circle. Furthermore, the idea that video games exist

inside the magic circle has been promoted by game designers and theorists (Salen & Zimmerman, 2003).

The gamer’s dilemma

The gamer’s dilemma offers a useful test case in the ethics of the metaverse (Luck, 2009, 2022). The gamer’s dilemma is a puzzle concerning the moral permissibility of two actions p and q we might perform, typically when playing computer games. p describes the performance of the action of virtual murder, which occurs when one character murders another in a virtual environment. In *Grand Theft Auto (GTA)*, a game that simulates the antics of a car thief, players may direct their characters to run over innocent pedestrians and commit virtual murder. Many consider virtual murder to be morally permissible (Pp): it is just a game and no one is actually murdered. This just a game defence of the moral permissibility of virtual murder is also known as the magic circle defence, grounded in the magic circle principle (Nguyen, 2017). After all, from moral norms prohibiting murder being true in the actual world, it by no means follows that they shall be true in a (walt-)fiction. Conversely, q describes the performance of the action of virtual child molestation, which occurs when an adult character molests a child character in a virtual environment. In a *GTA*-style game, an adult might lure a child into his car and grope the child. Were this walt-fictional world the actual world, this action would constitute actual child molestation. Many consider virtual paedophilia to be morally impermissible ($\neg Pq$). The puzzle arises because the just a game defence, though applicable with respect to virtual murder, does not appear to apply to virtual paedophilia.

This puzzle may be presented in the following manner as a paradox:

- P1: Virtual murder is morally permissible. (Pp)
- P2: There is no relevant difference between virtual murder and virtual child molestation, in respect to being permissible. ($Pp \leftrightarrow Pq$)
- P3: Virtual paedophilia is impermissible. ($\neg Pq$) (Luck, 2022)

There are at least two possible responses to this paradox: we could resolve the paradox by showing that at least one of P1–P3 is false or we could dissolve the paradox by showing that at least one of P1–P3 is actually implausible. For instance, P2 could be denied and we might assert an important difference between virtual murder and virtual child molestation. This difference could relate to the way in which participation in these virtual acts discloses or reveals something about the human perpetrator of these virtual acts. Whereas virtual murder does not really tell us anything about the murderous inclinations of human game players,

this does not appear to be the case for virtual child molestation. Why would one resort to virtual paedophilia unless one were an actual paedophile?⁸ Relations of cognitive or intentional directedness, matter and are apt for disclosure even within the virtual environment, making the thematization of the relation of aboutness in the IAO a significant move. At the same time, all versions of virtual digitalism (including the weak version relying on the concretization relation supported by the IAO) are unable to account for why ‘According to the fiction of *GTA*, a murder was committed’ is true whereas ‘A murder was committed’ is false.

As virtual walt-fictionalism stipulates (as all versions of virtual fictionalism do) that we behave as if virtual entities exist and distinguishes between real and fictional contexts, it is able to account for both relations of cognitive or intentional directedness and the distinction between the truth values of ‘According to the fiction of *GTA*, a murder was committed’ (true) and ‘A murder was committed’ (false). Furthermore, our virtual walt-fictionalist analysis of the gamer’s dilemma may be supported by the incorporation of multiple and possibly competing principles of interpretation at work in the background of walt-fictions (e.g. the import-export principle, the magic circle principle). Virtual paedophilia seems morally problematic, not least because it appears to disclose or reveal the presence of real-world dispositions to commit morally repugnant acts. The just a game defence does not appear to apply in the virtual paedophilia case, because there is a worry that we might export, in accordance with the import-export principle, certain questionable moral norms from a walt-fiction authorizing virtual paedophilia into the actual world. There is a greater relative weight on the import-export principle in the virtual paedophilia case and the magic circle principle is insufficiently powerful to support the just a game defence. By contrast, virtual murder is not especially informative with respect to the real-world dispositions of human game players. This results in a greater relative weight on the magic circle principle in the virtual murder case, allowing the just a game defence to apply and resulting in a greater tolerance of virtual murder.

To recapitulate, the gamer’s dilemma identifies a conflict between the moral import of two actions in the virtual environment: the murder of a virtual human character (viz. virtual murder) and the molestation of a virtual child (viz. virtual paedophilia). Virtual murder seems morally permissible, whereas virtual paedophilia does not. This dilemma is puzzling because it seems that virtual murder (e.g. the murder of a child) would surely count as worse than virtual molestation. However, we tolerate virtual murder but not

virtual paedophilia. Our response to the gamer’s dilemma involves a denial of P2: we invoke relations of cognitive or intentional directedness, the revelatory aspect of walt-fictions (i.e. how actions in the virtual environment might reveal or disclose aspects of our real-world dispositions, tendencies, and intentions), the competing principles of interpretation at work in the background of walt-fictions (e.g. the import-export principle, the magic circle principle) and the different weights that get assigned to these principles on the basis of what gets revealed about our real-world dispositions, tendencies, and intentions.

Here is a possible counter: does this walt-fictionalist response to the gamer’s dilemma not merely recast the puzzle in a different form, as opposed to resolving it completely? Instead of a puzzle centred on norms, do we not have instead a puzzle centred on weights assigned to different principles of interpretation? How can the gamer’s dilemma be resolved by an appeal to the differential weighting of the import-export principle and the magic circle principle? What determines the relative weighting of these principles? In addition, might the differential weighting not turn out just to be a matter of social convention? If so, then why not resolve the gamer’s dilemma simply by appealing directly to matters of social convention? This counter, however, fails to recognize several key aspects of the gamer’s dilemma: the distinction between the real and the virtual, the nature of the interface between the real and the virtual, and the aboutness of actions (real or virtual). Walt-fictionalism addresses these key aspects by parsing the real-virtual distinction in terms of a real-fictional distinction, describing how certain principles of interpretation operate at the interface between the real and the fictional or virtual, and foregrounding the aboutness of actions and the revelatory or disclosive character of certain actions or choices in fictional or virtual environments.

Virtual theft

In the Supreme Court of the Netherlands case 10/00101 J, two older boys used physical violence to coerce another boy to hand over a couple of virtual entities (viz. a virtual amulet and a virtual mask) and transfer them from his account to theirs within RuneScape, a fantasy massively multiplayer online role-playing game (MMORPG). The victim was kicked and threatened with a knife until he transferred these virtual entities.

Although the two perpetrators were initially convicted in lower court rulings, legal appeals were made against these convictions. The argument in which these appeals were grounded centred on the ontological status of the virtual entities and the nature of the principles governing the use of problems in the game of make-believe that is RuneScape. In the first instance, as the virtual mask and the virtual amulet were not real, they could not be

⁸ I am grateful to an anonymous reviewer for providing clarity on this aspect of the gamer’s dilemma.

considered as goods. If they were not goods, then they could not be considered as objects of theft. Furthermore, as RuneScape is an authorized game of make-believe whose principles of generation are widely shared, stable, and publicly agreed upon, it was argued that the transfer of virtual entities from one account to another was consistent with these stable principles of generation. In 2012, the Supreme Court of the Netherlands upheld the conviction and dismissed the appeals.

The ruling of the Supreme Court generates a puzzle for virtual irrealism: how can you steal something that does not exist? This puzzle has been termed the puzzle of virtual theft. According to the Dutch Supreme Court ruling, virtual entities can be regarded as goods and can therefore be the subjects of property offences (Wolswijk, 2012). While it is obvious that digital entities (i.e., the bits and bytes encoded on silicon chips) were stolen by the two boys in the RuneScape case, the Supreme Court did not appeal to the reality of these digital entities but referred directly to the virtual entities (viz. the virtual mask and the virtual amulet). On behalf of virtual irrealism, McDonnell and Wildman (2020) argue that this is the result of an ontological mistake from the Supreme Court: digital entities are real rather than imaginary, whereas the virtual entities that digital props prompt or invite us to make-believe are imaginary rather than real. The value of a virtual and walt-fictional experience causally hinges on digital props that support the imaginative enterprise. If the younger boy is deprived of these digital props, then he is deprived of both the material value of these digital props and the experiential value of advancing in the RuneScape game of make-believe.

We can strengthen the virtual walt-fictionalist response by invoking the import-export and magic circle principles. In the first instance, it by no means follows from a consistent application of the principles of generation in an authorized game of make-believe (viz. RuneScape) that no worries shall arise with the principles of interpretation. The import-export principle supports an epistemic connection between the actual world and the (walt-)fictional world. The actions of the two perpetrators threaten the integrity of the import-export principle, as morally deviant actions are performed in the real world in pursuit of digital props that support a more valuable virtual experience. The import-export principle does not support the use of morally questionable means in the real world in pursuit of various ends or purposes in the virtual world and will recommend certain moral restrictions on means-ends reasoning when it is applied across real and virtual worlds. In the second instance, the magic circle defence is of limited efficacy in the RuneScape case: physical violence was inflicted by the two perpetrators on the victim outside the magic circle of the walt-fiction and neither of the perpetrators showed remorse for what they did.

The epistemology of the metaverse

Scientific modelling as make-believe

On the moral epistemological front, we have already observed how we are at liberty to export moral norms from the (walt-)fictional world and into the actual world in accordance with the import-export principle. More generally, relatively recent developments in the philosophy of science have attempted to relate the Waltonian theory of mimesis as make-believe to scientific modelling. These attempts may be harnessed in support of a virtual walt-fictionalist epistemology of the metaverse.

In the literature on scientific modelling, physical systems are known as target systems and models may or may not represent these target systems via the mediation of model systems. What all walt-fictionalists (direct or indirect) have in common is their belief that the practice of scientific modelling is best understood in terms of scientists engaging in games of make-believe (e.g. in model descriptions of ideal gases, frictionless planes, the actions of perfect rational agents, etc). At minimum, scientists want to learn about the models themselves. At best, scientists want to learn about the target systems that are represented by these models (Poznic, 2016). According to indirect walt-fictionalism about scientific models, model descriptions are props that prescribe imaginings about model systems (Godfrey-Smith, 2009; Frigg, 2010, 2010a; Frigg & Nguyen, 2020). For instance, the ideal gas model prescribes imaginings about a model system (viz. the ideal gas), which is taken to represent a target system (viz. a real gas under normal conditions of pressure and temperature) (Iranzo-Ribera, 2022). According to direct walt-fictionalism about scientific models, on the other hand, the model description directly prescribes imaginings about a target system without the mediation of a model system (Toon, 2010, 2012; Levy, 2015). For instance, the model description of the ideal gas prescribes imaginings of a real gas described somewhat differently from the way it really is and there is no hypothetical model system (Iranzo-Ribera, 2022).

Our intention here is not to resolve the scientific modelling debate in favour of either direct walt-fictionalism, indirect walt-fictionalism, or a non-walt-fictionalist alternative. Rather, we shall pursue the more modest aim of suggesting how an appropriate extension of walt-fictionalism could help provide the relevant foundations for the epistemology of the metaverse. Props—whether they might be the bits and bytes of the metaverse, tree stumps in children's games of make-believe, or model descriptions in scientific models—prescribe an attitude of imaginative engagement: we mean 'It is to be imagined that p' when




Fig. 2 Screenshot of the Foldit program showing a folding puzzle in progress

we say ‘It is (walt-)fictional that p ’. Furthermore, this attitude of imaginative engagement is a propositional attitude (Currie, 1990; Salis & Frigg, 2016; McLoone, 2019; Iranzo-Ribera, 2022). In scientific models, the principles of generation that prescribe what it is we are to make-believe include the following: logical rules of inference, mathematical knowledge, and other field-specific theoretical knowledge. Each model is not complete: it only allows us to assign determinate truth values to those propositions that lie within the scope of that model. Those propositions that lie outside the scope of the model are indeterminate, since the model does not include facts that would make these propositions true or false. Last but not least, where model systems feature as intermediaries in indirect walt-fictionalism, model systems and target systems typically share a common structure. More specifically, this common structure is a causal structure that permits certain patterns of interactions. This shared causal structure is what allows claims about the target, resulting from model-based inferences, to be justified (Iranzo-Ribera, 2022).

HoloFoldit

Foldit is a multiplayer online game that engages non-scientists to help derive solutions to the protein folding problem (Fig. 2). The protein folding problem is a problem about how we infer or predict the 3D structure of a protein molecule on

the basis of information about the amino acid chain comprising that molecule (Cooper et al., 2010).

Smart (2022) HoloFoldit is a hypothetical system in which the use of a conventional desktop computer and physical display device in the real-world Foldit system is replaced by the use of a HoloLens-type mixed reality device. The Microsoft HoloLens is a head-mounted mixed reality device that enables users to interact with 3D virtual objects (or holograms) that are rendered in the real-world environment of the users. These holograms can be viewed from multiple angles and manipulated using a combination of voice commands and bodily movements.

The hypothetical HoloFoldit case offers a perfect example of how metaverse-based technologies could be employed to facilitate learning, discovery, and knowledge acquisition. The target problem is the protein folding problem (viz. the question of how a protein’s amino acid sequence dictates its 3D structure) and the model is a crowdsourcing-based model that challenges players to discover low-energy protein models by exploring the protein conformational space. The shared causal structure between the target system and the model system allows for model-based inferences to be made about how the protein folding problem might be solved.

Epistemic inequalities

We have thus far characterized games of make-believe in the metaverse in terms of digital props, the virtual entities

that we are prompted or invited to make-believe, principles of generation, and principles of interpretation. At the same time, a more critical and systematic study of these metaverse-based games of make-believe is likely to reveal an underlying power dynamic within which the metaverse is situated. This power dynamic is marked by certain epistemic inequalities that will probably get exacerbated as the implementation of the metaverse becomes more widespread (Turner, 2022). The first epistemic inequality concerns digital distraction: superimposed virtual objects constitute perceptual affordances that actively ‘call out’ for the attention of users. This may exacerbate the problem of digital distraction. The second epistemic inequality concerns digital deception: technologies can generate false or misleading appearances of reality (e.g., fake news articles and deepfake videos) and the metaverse will add virtual entities that are (possibly) phenomenologically indistinguishable from real-world physical entities to our gallery of entities. The third epistemic inequality concerns digital divergence: digital consumers are increasingly living in different informational universes or filter bubbles. This digital divergence could be accelerated with the metaverse. Nonetheless, we should not be in a rush to throw the baby out with the bathwater: the HoloFoldit example illustrates how metaverse technologies, when appropriately implemented, could promote learning, discovery, model-based inferences, and knowledge acquisition. We will simply need to be more aware of the risk of exacerbating epistemic inequalities through an uncritical, naïve, and indiscriminate approach to the metaverse.

Conclusion

To conclude, we have distinguished between virtual realism and virtual irrealism, two key metaphysical positions that could be staked out relative to the metaverse. We have identified two versions of virtual realism: strong virtual digitalism and weak virtual digitalism. We have argued that the cross-play and multiple realization problems confront strong virtual digitalism, while the causal exclusion problem confronts weak virtual digitalism. We have identified a version of virtual irrealism known as virtual walt-fictionalism and characterized this position in terms of digital props, the virtual entities that we are prompted or invited to make-believe, principles of generation, and principles of interpretation. We have demonstrated that there are good *prima facie* reasons to support virtual walt-fictionalism. We have argued that virtual walt-fictionalism, through a reliance on the import-export principle and the magic circle principle as principles of interpretation, can help us to make sense of the ethics of the metaverse and solve various puzzles (e.g. the gamer’s dilemma, the puzzle of virtual theft). In addition, we have demonstrated that virtual walt-fictionalism can shed light on how we can make model-based

inferences and derive conclusions about the physical or target systems.

Is there any reason why the claims made about the ethics and the epistemology of the metaverse could not be accommodated under some version of virtual realism, even if we accept the metaphysical objections to virtual realism? For instance, how might strong virtual digitalism be unable to draw on the import-export principle or the magic circle principle? Are there reasons why virtual realism cannot deal with the ethical and epistemological issues in the same way that a virtual walt-fictionalist might? The difference between virtual walt-fictionalism and virtual realism is a metaphysical one: the former denies the reality of virtual entities whereas the latter asserts their reality. This metaphysical difference makes virtual walt-fictionalism more sensitive to the distinction between the real and the fictional, the real/fictional interface, and the principles of interpretation operating at this interface in a way that virtual realism cannot be (at the pain of inconsistency). As a result of both its denial of the reality of virtual entities and its assertion that we behave as if they do exist, virtual walt-fictionalism enjoys certain advantages (*viz.* sophistication, versatility, ontological parsimony, explanatory scope, etc) over its virtual digitalist rivals (strong or weak). Virtual walt-fictionalism, grounded in its central trope of the metaverse as a prop-driven game of make-believe, has the potential to deliver, in one fell swoop, the metaphysics, axiology (ethics), and epistemology of the metaverse.

Acknowledgements The author wishes to acknowledge his wife (faithful Penelope) and son (brave Telemachus), who have been a perennial source of joy and inspiration. In addition, he wishes to acknowledge his students and faculty colleagues for their continued goodwill, support and encouragement.

Author contributions Not applicable.

Funding This work was supported by an intramural Nanyang Technological University (CoHASS-NBS) grant.

Data availability Not applicable.

Declarations

Competing interest The author has no relevant financial or non-financial interests to disclose.

Consent to participate Not applicable.

Consent for publication Not applicable.

Ethical approval This research does not involve human or animal subjects and no ethical approval is required.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes

were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aarseth, E. (2007). Doors & perception: Fiction vs. simulation in games. *Intermedialités*, 9, 35–44.
- 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>
- Brampton, C. K. (1964). Nominalism & the law of parsimony. *Modern Schoolman*, 41(3), 273–281. <https://doi.org/10.5840/schoolman196441356>
- Ceusters, W., & Smith, B. (2015). Aboutness: Towards foundations for the information artifact ontology. In Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO) (pp. 1–5). CEUR.
- Chalmers, D. J. (2019). The virtual as the digital. *Disputatio*, 11(55), 453–486.
- Cooper, S., Khatib, F., Treuille, A., Barbero, J., Lee, J., Beenen, M., Leaver-Fay, A., Baker, D., Popović, Z., et al. (2010). Predicting protein structures with a multiplayer online game. *Nature*, 466(7307), 756–760.
- Currie, G. (1990). *The nature of fiction*. Cambridge University Press.
- Deutsch, H., & Garbacz, P. (2022). Relative identity. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Winter 2022). Metaphysics Research Lab, Stanford University.
- Enderton, H. B. (1972). *A mathematical introduction to logic*. Academic Press.
- Friend, S. (2008). Imagining fact & fiction. In K. Stock & K. Thomsen-Jones (Eds.), *New waves in aesthetics* (pp. 150–169). Palgrave-Macmillan.
- Frigg, R. (2010a). Fiction & scientific representation.
- Frigg, R. (2010). Models & fiction. *Synthese*, 172(2), 251–268. <https://doi.org/10.1007/s11229-009-9505-0>
- Frigg, R. (2020). *Modelling nature: An opinionated introduction to scientific representation*. Springer.
- Gendler, T. S. (2000). The puzzle of imaginative resistance. *Journal of Philosophy*, 97(2), 55–81.
- Godfrey-Smith, P. (2009). Models & fictions in science. *Philosophical Studies*, 143(1), 101–116. <https://doi.org/10.1007/s11098-008-9313-2>
- Hacker, P. M. S. (1982). Events & objects in space & time. *Mind*, 91(361), 1–19. <https://doi.org/10.1093/mind/xci.361.1>
- Hacker, P. M. S. (1982). Events, ontology & grammar. *Philosophy*, 57(222), 477–486. <https://doi.org/10.1017/s0031819100056102>
- Huizinga, J. (1949). *Homo ludens: A study of the play element in human culture*. Routledge & Kegan Paul.
- Iranzo-Ribera, N. (2022). Scientific counterfactuals as make-believe. *Synthese*, 200(6), 89. <https://doi.org/10.1007/s11229-022-03949-8>
- Juul, J. (2005). *Half-real: Video games between real rules & fictional worlds*. MIT Press.
- Kim, J. (1992). The nonreductivist's troubles with mental causation. In J. Heil & A. R. Mele (Eds.), *Mental causation*. Oxford University Press.
- Levy, A. (2015). Modeling without models. *Philosophical Studies*, 172(3), 781–798. <https://doi.org/10.1007/s11098-014-0333-9>
- Lewis, D. K. (1986). *On the plurality of worlds*. Wiley.
- Luck, M. (2009). The gamer's dilemma: An analysis of the arguments for the moral distinction between virtual murder & virtual Paedophilia. *Ethics & Information Technology*, 11(1), 31–36. <https://doi.org/10.1007/s10676-008-9168-4>
- Luck, M. (2022). The grave resolution to the gamer's dilemma: An argument for a moral distinction between virtual murder & virtual child molestation. *Philosophia*, 50(3), 1287–1308. <https://doi.org/10.1007/s11406-021-00455-y>
- Ludlow, P. (2019). The social furniture of virtual worlds. *Disputatio*, 11(55), 345–369. <https://doi.org/10.2478/disp-2019-0009>
- Maurer, A. (1978). Method in Ockham's nominalism. *The Monist*, 61(3), 426–443. <https://doi.org/10.5840/monist197861334>
- McDonnell, N., & Wildman, N. (2019). Virtual reality: Digital or fictional? *Disputatio*, 11(55), 371–397. <https://doi.org/10.2478/disp-2019-0004>
- McDonnell, N., & Wildman, N. (2020). The puzzle of virtual theft. *Analysis*, 80(3), 493–499. <https://doi.org/10.1093/analysis/anaa005>
- McLoone, B. (2019). Thumper the infinitesimal rabbit: A fictionalist perspective on some “unimaginable” model systems in biology. *Philosophy of Science*, 86(4), 662–671. <https://doi.org/10.1086/704976>
- Meskin, A., & Robson, J. (2012). Fiction & fictional worlds in videogames. In J. R. Sageng, T. M. Larsen, & H. Fossheim (Eds.), *The philosophy of computer games* (pp. 201–18). Springer.
- Milgram, P., Colquhoun, H., et al. (1999). A taxonomy of real & virtual world display integration. *Mixed Reality*, 1, 1–26.
- Milgram, P., & Kishino, F. (1994). A taxonomy of mixed reality visual displays. *IEICE Transaction on Information and Systems*, 77(12), 1321–1329.
- Milgram, P., Takemura, H., Utsumi, A., & Kishino, F. (1995). Augmented reality: A class of displays on the reality-virtuality continuum. *Telemanipulator & Telepresence Technologies*, 2351, 282–292.
- Nguyen, C. T. (2017). Philosophy of games. *Philosophy Compass*, 12(8), e12426.
- Poznic, M. (2016). Make-believe & model-based representation in science: The epistemology of Frigg's & Toon's fictionalist views of modeling. *Teorema*, 35(3), 201–218.
- Rosenbaum, J. W. (2016). *Poetic license: Learning morality from fiction in light of imaginative resistance* (pp. 165–183). Revista Internacional de Filosofía.
- Salen, K., & Zimmerman, E. (2003). *Rules of play: Game design fundamentals*. MIT Press.
- Salis, F., & Frigg, R. (2016). Capturing the scientific imagination. In P. Godfrey-Smith & A. Levy (Eds.), *The scientific imagination*. Oxford University Press.
- Schaffer, J. (2009). On what grounds what. In D. Manley, D. J. Chalmers, & R. Wasserman (Eds.), *Metametaphysics: New essays on the foundations of ontology* (pp. 347–383). Oxford University Press.
- Schechtman, M. (1996). *The constitution of selves*. Cornell University Press.
- Smart, P. (2022). Minds in the metaverse: Extended cognition meets mixed reality. *Philosophy & Technology*, 35(4), 1–29. <https://doi.org/10.1007/s13347-022-00580-w>
- Tavinor, G. (2009). *The art of videogames*. Wiley.
- Toon, A. (2010). The ontology of theoretical modelling: Models as make-believe. *Synthese*, 172(2), 301–315. <https://doi.org/10.1007/s11229-009-9508-x>
- Toon, A. (2012). *Models as make-believe: Imagination, fiction and scientific representation*. Palgrave-Macmillan.
- Turner, C. (2022). Augmented reality, augmented epistemology, & the realworld web. *Philosophy & Technology*, 35(1), 1–28. <https://doi.org/10.1007/s13347-022-00496-5>

- Vaihinger, H. (1923). The philosophy of 'as if' (C. K. Ogden, Trans.). Kegan Paul.
- Walton, K. L. (1990). *Mimesis as make-believe: On the foundations of the representational arts*. Harvard University Press.
- Wilson, R. A., & Craver, C. F. (2007). Realization: Metaphysical & scientific perspectives. In *Philosophy of psychology & cognitive science* (pp. 81–104). Elsevier.
- Wolswijk, H. (2012). Theft: Taking a virtual object in RuneScape: Judgment of 31 January 2012, case no. 10/00101 j. *The Journal of Criminal Law*, 76(6), 459–462.
- Woodward, R. (2014). Walton on fictionality. *Philosophy Compass*, 9(12), 825–836. <https://doi.org/10.1111/phc3.12178>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.