Philosophy & Technology (2022) 35: 7
https://doi.org/10.1007/s13347-022-00502-w

RESEARCH ARTICLE

Causation in a Virtual World: a Mechanistic Approach

Billy Wheeler[1]

Received: 5 August 2019 / Accepted: 15 January 2022 / Published online: 4 February 2022

Abstract
Objects appear to causally interact with one another in virtual worlds, such as video games, virtual reality, and training simulations. Is this causation real or is it illusory? In this paper I argue that virtual causation is as real as physical causation. I achieve this in two steps: firstly, I show how virtual causation has all the important hallmarks of relations that are causal, as opposed to merely accidental, and secondly, I show how virtual causation is genuine according to one influential metaphysical theory of causation: the mechanistic approach.

Keywords Virtual worlds · Virtual realism · Causation · Mechanisms · Video games · Cellular automata

1 Introduction

There is a weapon in the online multiplayer game *League of Legends* called a "hextech gunblade." It is highly sought after because it grants the player several useful abilities. It can slow an enemy's movements down by 40% and restore health to the user by "stealing" it from their opponent. On the surface, it looks as if the hextech gunblade has causal powers. But are these powers real? Can events in virtual worlds, such as video games, stand in genuine causal relationships to other virtual events? On the one hand, it looks as if this causation is merely apparent. Video games, which might be argued, are a form of fiction and, like representations of causation in other fictional media, such as movies and books, are not genuine. They might give the *impression* or *perception* of causation—but this perception is only illusory. On the other hand, virtual events are unlike other fictional media in that they seem to support counterfactuals. If the player had not used the hextech gunblade, then their opponent's movements would not have slowed down. Like causation in the physical world, virtual causation makes a difference to what happens.

Billy Wheeler
billy.w@vinuni.edu.vn

1 Faculty of Arts and Sciences, VinUniversity, Vinhomes Ocean Park, Hanoi, Vietnam

Understanding the status of causation in virtual worlds is important for two reasons. Firstly, knowing whether these causes are real forms part of a wider debate about the metaphysics of virtual worlds and the objects, events, and activities that take place within them. In this debate, the two most prominent positions are *virtual realism*, which, simply put, says virtual objects can (in the right circumstances) be just as real as physical objects, and *virtual fictionalism*, which claims that virtual objects are almost always fictions, in the same way events and characters in a novel or movie are a fiction. Causation enters this argument in a crucial way through the arguments of David Chalmers (2017). Chalmers is a self-confessed virtual realist. His most important argument for realism is that because virtual objects appear to have causal powers, and because digital objects provide the ontological ground for these powers, so virtual objects are real. I agree with Chalmers' virtual realism, but as I will explain in more detail below, the causal argument is weak and can only be convincing if virtual objects have genuine, and not just apparent causal powers, something Chalmers himself does not explicitly argue for.

This is not the only reason why it would it be useful to know the metaphysical status of causation in virtual worlds. Philosophers have studied virtual causation in the form of computer simulations. The influential accounts of Clark Glymour and Richard Spirtes[1]were developed using computer simulations to show how causal inferences are possible in the absence of interventions and manipulations. Others have investigated simulations to better understand emergent causation and the relationship between lower- and higher-level causal claims.[2] There is a meta-philosophical question here. How do these simulations support or justify the philosophical conclusions made? One common idea is that these simulations are like other kinds of simulations in the natural sciences. On this view, the causal relations would be no more real than the galaxies in a computer simulation of the early universe. Any evidence the simulation provides comes in the strength of its analogy to the physical world. However, if causation in virtual worlds is real, and this includes computer simulations, then it would require a rethinking of causal simulations in philosophy. What we would have would not be a model but something more akin to a controlled experiment—where the genuine object of study (e.g., causation) is there but analyzed under determined and specified conditions.

In this paper I will argue that causation in virtual worlds is a genuine kind of causation and is no less causation than what occurs between physical events. Virtual causation is not a kind of "accidental" or illusory form of causation, and by studying causal relations in virtual worlds, we can learn something of value about causation in the physical world. My argument will proceed in two steps. In the first step, I shall demonstrate that virtual causal relations exhibit all the important features commonly associated with physical causal relations, such as supporting counterfactuals, being asymmetric, and chance raising. Since the features of causal relations are themselves subject to debate and controversy, this list will not be exhaustive, but it will capture

[1] Glymour et al. (1987); Glymour and Cooper (1999) and Spirtes et al. (2000).

[2] Mark Bedau (1997),), Beraldo-de-Araujo & Baravalle (2017), and Holly Andersen (2017)

what many have taken to be the key difference between causal relations and mere correlations or accidental regularities. The second step will be to show how virtual causation comes out as a genuine form of causation on one substantive theory of causation. This will be done by appealing to the mechanistic theory of causation as advocated and developed by Stuart Glennan (1996, 2010, and 2017). This theory has been very influential in recent thinking about causation and has been shown to overcome a number of problems in existing "empiricist views" as well as unify causal claims across the sciences.[3]

In the next section, I outline some key definitions and arguments in the realism-antirealism debate about virtual worlds as well as Chalmers' causal argument for virtual realism. In Section 3, I provide a close examination of an example of causation in a virtual world from a well-known video game. This will serve as an example to help investigate the similarities and differences between virtual and physical causation. In Section 4, I introduce Glennan's mechanistic theory of causation and explain why this is a suitable metaphysical theory to understand causation in virtual worlds. Here I identify the mechanisms responsible for causal relations in virtual worlds. I finish by discussing some potential objections to the approach taken concerning the issue of causal relevance, the exclusion problem, and whether it ties virtual causation too closely to the hardware and architecture of computing machines.

2 Virtual Worlds and Virtual Realism

2.1 What is a Virtual World?

Many of us have experienced a virtual world. At a simple level, a virtual world is a computer-generated space or environment that can be experienced by one or more users. This includes multiplayer online games like *League of Legends* where players can interact as well as single-player games such as *Super Mario Bros*. It also includes state-of-the-art virtual reality that uses immersive technologies like head-mounted displays and motion sensors . A little reflection, however, shows that this definition is too broad to capture what is distinct about virtual worlds. On this definition, any computer-generated space could be considered a virtual world, including desktop browsers, social media webpages, online chatrooms, and even movies (provided they are displayed with the aid of a computer). For this reason, philosophers and technologists typically use a narrower definition.

According to Michael Heim, what really separates virtual worlds from other kinds of computer-generated content are their interactivity and immersiveness:

> [A virtual world is] a scene or experience with which a participant can interact by using computer-controlled input-output devices...[It] pertains to convincing the participant that he or she is actually in another place, by substituting the normal sensory input received by the participant with information produced

[3] See Glennan and Illari (2018)

> by a computer. This is usually done through three-dimensional graphics and input-output devices that closely resemble the participants normal interface with the physical world. (1993, p. 160)

Unlike a movie or a novel, virtual worlds are *interactive*. What a person does makes a difference to the content that is produced and therefore experienced in the virtual world. This aspect is essential in video games as players need to make decisions to progress further in the game. The outcome of these decisions is fed into a computer such that, following its programs, it will make necessary changes to the virtual world experienced. Interactivity is also a core component of training simulators. These are used, for example, in the instruction of pilots and astronauts, and provided their experiences are similar enough to the physical world, can allow for the transfer of knowledge and skills gained through the simulation.

This still leaves too many computer-generated spaces defined as virtual worlds: even e-mail and word processors are interactive in this basic sense. To account for this Heim and others requires that computer-generated environments be *immersive* to truly count as virtual worlds.

Immersiveness is a difficult concept to define and appears to involve several aspects. It has been argued that something is immersive if it creates a feeling of presence (Calleja, 2014), provides a feeling of embodiment (Schultze, 2010), is like the physical world in certain respects (Brey, 2014), is logically consistent (Ropolyi 2015), and exhibits causal relationships among objects and events (Cavazza et al. 2007; Grabarczyk and Pokropski 2016). At an intuitive level, a world is immersive if the person feels that they "exist" in that world, and, for a moment or two, they feel they have left the ordinary physical world behind. Vivid dreams are an example of an experience that can be highly immersive. Immersiveness is a qualitative feature of virtual worlds that can come in degrees. Some computer-generated environments might be more or less immersive than others. This will partly depend on the quality of the designed environment as well as the individual user's own psychological response to that environment.

In the remainder of this discussion, I will use the following as a definition of a virtual world:

> **Virtual World**: A virtual world is a computer-generated environment that is both interactive and immersive.

Because immersiveness (and to a lesser extent interactivity) comes in degrees, there is no sharp boundary between those computer-generated environments that are "worlds" and those that are "non-worlds." Fortunately, most philosophers working in this area recognize a core set of cases as examples of virtual worlds, and this can help us focus our discussion going forward. Most philosophers recognize traditional video games, virtual reality, and computer-generated training simulations as virtual worlds.

Other kinds of cases are more debatable. For example, we do not normally consider computer simulations in the natural and social sciences as virtual worlds. These are often interpreted as a type of model that represents the physical world in one or more respects. Philip Brey (2014) calls computer simulations in science “weakly interactive” because the scientist has the option to change the values of the equations solved by the program. After this stage, there is often little further input provided by the scientist. Chalmers (2017) accepts computer simulations as virtual worlds but argues we should distinguish between different kinds of virtual worlds based on how they are perceived by the user and the role they play. I will follow Brey and Chalmers in accepting computer simulations as virtual worlds. Currently most computer simulations are weakly interactive and immersive. But there is no reason why they need to be, and it is likely that as technology improves, computer simulations used in fields like sociology or psychology could become highly interactive and immersive. It would seem inconsistent, therefore, to deny that they are virtual worlds just because they are used for research purposes rather than entertainment or skills training.

2.2 Realism and Fictionalism About Virtual Worlds

The realism-antirealism debate is a perennial one and appears in discussions of a wide range of phenomena, including *truth*, *unobservable entities*, *moral properties*, *causal relations*, and *possible worlds*. The debate concerning the realism and antirealism of *virtual worlds* is not as well developed as these other topics. Chalmers can be credited with helping to shape this debate and offers the following descriptions of two opposing views:

Virtual Realism:

(1) Virtual objects really exist.
(2) Events in virtual reality really take place.
(3) Experiences in virtual reality are non-illusory.
(4) Virtual experiences are as valuable as non-virtual experiences.

Virtual Antirealism:[4]

(1) Virtual objects do not really exist
(2) Events in virtual reality do not really take place.
(3) Experiences in virtual reality are illusory.
(4) Virtual experiences are less valuable than non-virtual experiences.

(2017, p. 310).

Chalmers sees these as packages of theses that naturally go together, but each thesis should be regarded as separable from the others such that it is possible to hold just one or two from each package. The four theses in each package require some explanation and qualification.

[4] Chalmers calls this opposing position “virtual irrealism”; however, since this is often used to refer to a specific form of antirealism associated with NelsonGoodman (1978), I have preferred to use the more neutral “antirealism” label

The first thesis (1) is *metaphysical* and takes a position on the status of virtual objects. An example of a virtual object would be any kind of object that can be interacted with in a virtual world and appears to have causal powers.[5] The hextech gunblade described in the introduction is one such object, as are certain things like platforms, coins, 1-ups, and player avatars. To a simple approximation, virtual realists believe that these objects really exist (they are part of the furniture of reality), whereas antirealists reject this. The most widely adopted form of virtual antirealism is *fictionalism* whereby virtual objects are given the same ontological status as fictional entities in related media like movies, novels, and some traditional (non-computer-generated) games. On this interpretation, virtual objects are no more real according to fictionalists than Gandalf from *The Lord of the Rings* or the planet Tatooine from *Star Wars*. Computers help sustain these fictions and make it easier for us to imagine and enjoy them, but they do not make them any more real.

By contrast, realists believe that virtual objects are just as real as physical objects. How this should be spelled out in detail depends on the realist approach taken. Chalmers calls his view "virtual digitalism" because he identifies the metaphysical ground of virtual objects with digital data structures in a computer:

> What are virtual objects? In my view, they are digital objects, constituted by computational processors on a computer. To a first approximation, they can be regarded as *data structures*, which are grounded in computational processes which are themselves grounded in physical processes on one or more computers. (2017, 317. My emphasis)

That virtual objects are data structures is not meant to be an obvious fact known a priori nor is it meant to be an eliminative claim. It is, Chalmers says, akin to the identification of water with H_2O or with the claim that stars are exploding balls of gas. It is an empirical fact. That is not to say that these identifications do not invite their own deep metaphysical questions, e.g., about relations of parts to wholes and lower/higher level phenomena. However these questions should be answered in the case of physical objects, Chalmers believes the same kinds of answer can be given to virtual objects. There is no metaphysical discrimination, therefore, between virtual and ordinary physical objects. Precisely what these data structures are and how they relate to virtual phenomena, I will return to again later when I propose how this should be applied to causal relations.

I endorse Chalmers' version of realism, but other kinds have been developed. Heim, from which Chalmers borrows the term "virtual realism," takes a view that is much weaker than Chalmers. Heim positions his view somewhere between realism and antirealism. According to him, virtual objects are "non-representational phenomena," and their existence is a function of their practical applications (1998, p. 33–51). He therefore offers a pragmatic or social constructivist approach to virtual worlds. Brey (2014) offers a type of piecemeal realism whereby only some virtual

[5] According to Brey "a virtual object is a digital object that is represented graphically as an object or region in a two-or three-dimensional space and that can be interacted with or used through a computer interface" (2014, p. 44).

objects are real and others are not. This can be contrasted to the wholesale realism of Chalmers, whereby every virtual object, event, and process is real. Whilst Brey accepts the existence of virtual objects in the minimal sense in that they are produced by physical processes on a computer, a virtual object becomes real in the sense of being non-fictional provided it reproduces the same properties as the object it represents in the physical world. A virtual calculator can be real in this sense provided it calculates, and a virtual door can be real provided avatars can pass through it. On the other hand, most kittens are not real but are fictional because they were not born, do not consume food, grow, reproduce, etc.

The second thesis (2) is also metaphysical but focusses on events rather than objects. From the point of view of the user experiencing virtual worlds, distinct events appear to occur. All virtual worlds have an element of time and linear progression built into them. This is helpful if the virtual world is a video game, and the player needs to complete a stage in a certain amount of time. But the temporal dimension of virtual worlds is more fundamental than this and is linked to the fact that the computations that produce them take place in real (physical) time. According to Chalmers' brand of virtual realism if a dragon in a virtual world flies through the air, then this is an event that really takes place. If a player uses a weapon to kill their enemy, then this is an event that really takes place.

Fictionalists would say that there is no more reality to the dragon flying in a virtual world than a dragon flying in a book of *The Lord of the Rings*. Though different fictional explanations exist, one of the most popular ones takes inspiration from Kendall Walton's (1990) influential account of "make-believe."[6] In his theory, physical objects can be used as "props" to represent other physical objects. For fictionalists, computers do the same; it is just that the graphics and interactivity make those props resemble the physical objects they stand for to such a high degree that it is easy to mistake them for the real thing. As we shall see, the main reason why Chalmers rejects this argument is down to the causal nature of virtual objects and events. Events that take place in virtual worlds are unlike fictional events because they can be causally connected to other virtual events.

The first and second theses are the most important ones for understanding Chalmers' causal argument so I will have less to say about three and four, despite them being interesting in their own right. Thesis (3) is *epistemological* and concerns the veracity of our beliefs based on experiences in virtual worlds. Chalmers (2017, p. 310) believes that this thesis is often conflated with (1) and (2) which drives the intuitions behind antirealism about virtual worlds. According to him, it is perfectly possible to acquire true and false beliefs from experiences in virtual worlds and are only a source of epistemological trouble if we confuse them for the real world. This thesis can be held apart from (1) and (2) as in the case of Cogburn and Silcox (2014)

6 See, e.g., Grant Tavinor (2009), McDonnell and Wildman (2019), Martin Ricksand (2020), Eric Studt (2021)

who are fictionalists but allow the possibility of knowledge from virtual worlds.[7] In a similar way, thesis (4) can also be held separately from (1) and (2). Thesis (4) raises the question of the *axiology* or *value* of virtual worlds, and antirealists typically draw inspiration from Robert Nozick's (1981) infamous "experience machine" thought experiment that he argues shows virtual worlds are less meaningful than real-life experiences. Recently defenders have countered these arguments by showing how the interactivity and sociality of virtual worlds can make them just as meaningful.[8]

2.3 Chalmers's Causal Argument for Realism

Chalmers advocates realism about virtual worlds, and one of his most important arguments in defense of theses (1) and (2) is his causal argument:

(1) Virtual objects have certain causal powers (to affect other virtual objects, to affect users, and so on).
(2) Digital objects really have those causal powers (and nothing else does).

———

(3) Virtual objects are digital objects.

(2017, p. 318)

From this, he concludes that as digital objects are real, therefore, virtual objects must be real as well. By "digital objects," Chalmers does not mean abstract data but the representation of this data in the physical components of the computer. For fictionalists, these components would typically be regarded as real. If virtual objects are (or in some sense grounded in) digital objects, then virtual objects should be regarded as real. Notice that we do not have to accept Chalmers' particular brand of realism to use the causal argument. Even if we want to stay agnostic on the exact relationship between virtual objects and the processes on a computer, we could still use a similar argument for causation.

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

———

(3) Therefore, virtual objects exist.

In both arguments, the first premise is really carrying all the work. How convincing would this be to an antirealist about virtual worlds? As Chalmers himself seems to acknowledge, this argument appears to beg the question in favor of realism:

[7] Realists approaches to the epistemology of virtual worlds that are neutral over its metaphysics can be found in James McBain (2017) and Billy Wheeler (2020).

[8] See, e.g., Peter Ludlow (2017), Pietrucha (2017); Weijers and DiSilvestro (2017), and Alexis Elder (2017)

> Fictionalists will probably deny the first premise by saying that virtual objects do not have causal powers, or better, that they have causal powers only in the sense that Gandalf has causal powers. That is, they have causal powers within a fictional world, and any effects on the real world are brought about not by the object but by a representation of the object. Still, even the nonconclusive argument from the premise that virtual objects *seem* to have these causal powers and that digital objects really have those powers is a reasonably strong one. (2017, p. 318. My emphasis)

Going back to our initial example, if a character in *League of Legends* uses a weapon to kill an enemy, fictionalists will deny that this an example of real causation. This is no more real causation, according to them, than when Gandalf in *The Lord of the Rings* casts a spell that defends the fellowship from an attacking horde of orcs. Hence Chalmers cannot argue for realism by already assuming that causal powers or relations in virtual worlds are real. Instead, as he admits, the argument must be weakened to include a modified version of premise 1, such that objects only *appear* or to *seem* to engage in causation:

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

Chalmers thinks his argument is still "reasonably strong." However, this is an assumption he never argues for in any detail. It seems reasonable for the antirealist to reject the argument based on the modified premise (1'). There are two good reasons for this. Firstly, it should come as a surprise to nobody that objects, events, and processes that take place in virtual worlds can appear similar to objects, events, and processes in the physical world. This is because these objects have been deliberately designed to give the appearance of real-world phenomena. We know from empirical studies that the appearance of causation is an important element in making virtual worlds immersive.[9]Given that developers of virtual worlds typically desire them to be as immersive as possible, it stands to reason it would include the appearance of causal powers and relations.

Secondly, we know from the long history of analyzing causation that two events can appear causally connected even when they are not. We have developed sophisticated theories and methods to detect causation and avoid making false causal inferences. The "common cause fallacy" occurs when a person infers causation between two events or processes when in fact both are caused by a third event or process. To convince us that causal relations in virtual worlds are real, one would need to at least exclude the possibility that when events take place in virtual worlds, they are not both common causes of a single event or process in the hardware of the computer. If it was, then this would seriously undermine the argument that the appearance of causation is good evidence for genuine causal relations in virtual worlds,

[9] Cavazza et al. (2007); Hoyet et al. (2012).

The causal argument relies on a non-deductive argument to support premise (1). As Chalmers says: “If there are real objects that have all the apparent properties of virtual objects, there is not much reason to suppose that virtual objects belong to a separate layer of fictional objects” (2017, p. 318). The argument in support of premise (1) is therefore a type of analogical argument that might go something like the following:

(1) Virtual objects and events have properties P_1, P_2, $P_3 \ldots P_n$.
(2) Physical objects and events that have properties P_1, P_2, $P_3 \ldots P_n$ stand in genuine causal relationships.

———

(3) Therefore, virtual objects and events can stand in genuine causal relationships.

The strength of the argument depends on the strength of the analogy (which is never spelt out by Chalmers) as well as a consideration of any potential negative analogy between the two. I now turn to investigate what the positive and negative analogy are between virtual and physical causation to see if it can indeed support an argument of this kind.

3 Comparing Virtual and Physical Causation

3.1 Causation in Mario's World

Released for the Nintendo Entertainment System, *Super Mario Bros.* is one of the most recognizable videos games of all time. The titular character “Mario” became an icon of the 8-bit video gaming generation, and games involving Mario have continued to appear alongside the most recent and advanced consoles. In the original 1985 version of the game, you play as Mario and need to defeat an evil dinosaur-type villain who has captured a princess. Along the way, you must overcome minor villains, and each level contains various power-ups such as coins, mushrooms, and extra lives, to help you.

When a person plays this video game, there will be many examples of what appear to be or seem to be objects and events. Here are some examples of events that typically take place when a person plays the game:

- Mario collects a coin.
- Mario is killed.
- Mario jumps on an enemy.
- Mario enters a pipe.
- Mario collects a mushroom.
- Mario explodes a brick.

If we focus on causation as a relation between events, then Chalmers' suggestion that there is apparent causation in virtual worlds can clearly be demonstrated in the case of *Super Mario Bros*. Many of the events above are causally related to other events. For example, if Mario collects a mushroom, then he will grow taller. If Mario enters a pipe, then he will enter a new zone. It seems natural and right to say that these first events were the *cause* of the second. It is because Mario collected a mushroom that his size increased. It is because Mario entered a pipe that he was transported to another zone.

However, as mentioned in the previous section, just because these events *appear* causally related does not mean that they are. What we need is a more detailed examination of the properties of the relationship between these virtual events and a comparison with causation as it is ordinarily understood in the physical world. This task is made difficult by the fact that physical causation is itself an object of philosophical investigation with many competing interpretations. Despite this, there is a common list of properties that are often believed to hold for relations of physical causation, and we can use this list to investigate the strength of the analogy between virtual and physical causation. That some of these properties have been denied or questioned by philosophers need not necessarily be a problem. In fact, if virtual causation is a form of genuine causation, then we should expect the same or similar disagreements to arise there. What matters is that there is a sufficiently strong analogy between the properties to justify the claim that virtual events, like physical events, can be connected by a causal relation.

3.2 The Positive Analogy

In her work on the role of analogies in the natural science, Mary Hesse (1966) makes a distinction between *positive*, *negative*, and *neutral* analogies. The positive analogy comprises those properties that are believed to be true of both the source and target model or system. The negative analogy comprises those properties that are thought to be different. And the neutral analogy comprises the property or properties that are under dispute. An argument by analogy tends to be stronger in proportion to the size of its positive and negative analogies. The greater the positive analogy and the smaller the negative analogy, the more likely the property in the neutral analogy holds for both systems. Things are little bit more complicated than this because the properties in question need to be *structurally relevant* ones (1966, p. 92). Because we are already working with properties that are commonly believed to be relevant to causation in the physical world, we can assume for the sake of argument that these are relevant properties. In our example, the neutral analogy consists of just one property—namely, the relation of causation.

Below I discuss five properties that are commonly believed to be true for physical causation that also seem to be true for (apparent) virtual causation.

(A) ***Causal Is Counterfactual Supporting***

A widespread belief about causation is that causes make a difference to the occurrence of the effect. If two events are causally connected, then there appears to be a counterfactual dependency between them. In the classic example, if Suzy throws a stone at a bottle which causes it to break, then it seems true to say that if

Suzy had *not* thrown the stone, then the bottle would *not* have broken. Although the precise connection between counterfactuals and causation is debated, most accept that any reasonable theory of physical causation should explain how knowing that there is a causal link between events warrants or justifies belief in the associated counterfactual.

The most striking difference between causation as it is experienced in virtual worlds and how it is experienced in fictional media, like novels and movies, is that causation in virtual worlds seems to support counterfactuals. Going back to our example above, suppose that Mario (when in his diminutive size) collects a mushroom and subsequently grows taller. It seems right to say that if Mario had *not* collected a mushroom, then he would *not* have grown taller. Our reason for believing this counterfactual may be based on past experience, but at a fundamental level, we know that this counterfactual holds because of the way the game is designed. At the level of the software for the game, it has been written such that whenever Mario collects a mushroom, a subroutine is initiated that leads to Mario being displayed in his larger size. This program provides the physical base that guarantees that whenever Mario collects a mushroom, he will grow taller. Notice that no such guarantee is given in the case of movies or novels. Even though it seems reasonable to say that had Gandalf not cast the spell, the attacking orc horde would have killed Frodo, there is nothing that guarantees the truth of the counterfactual in this case.

(B) ***Causation Is Productive***

A similarly strong intuition about causation in the physical world is that causes are productive of their effects. There have been many ways to spell out this intuition in detail. Some philosophers have talked about "necessary connections,"[10] others about "processes,"[11] or "mechanisms".[12] The essential idea is that causes are connected to effects in some substantial way that goes beyond regularity and counterfactual dependence. Indeed, advocates of production often explain counterfactual dependence because there exists a substantive connection. In the example of Suzy and the bottle, it seems reasonable to say that when Suzy throws the stone and the stone hits the bottle, some process or mechanism is responsible for the bottle breaking such that the stone (or better the collision with the stone) brought the breaking about. In the Salmon-Dowe conserved quantities theory of causation, this is explained by the fact that the stone transferred energy or momentum to the bottle which resulted in its structural failure.

What about virtual causation? Is there some process or mechanism that connects one virtual event to another? Once again let us take the two events Mario collecting the mushroom and Mario getting taller. If the player directs the Mario sprite close enough towards a mushroom to collect it, then the right kind of signal is sent to the hard drive of the video game console. The structure of the

[10] David Armstrong (1983), Fred Dretske (1977), Michael Tooley (1977)

[11] Wesley Salmon (1984), Philip Dowe (2000)

[12] Glennan (1996, 2010 and 2017), Machamer et al. (2000).

computer is organized such that whenever this input is given, a new output is provided. We can say that the computer contains a mechanism, in the form of a sub-routine, such that whenever Mario collects a mushroom, he will get bigger. This relation seems productive. The collecting of the mushroom is connected in a physical way (via the mechanism) to Mario getting bigger. There is much more detail that can be given here regarding this mechanism and the roles it plays in the causal process in virtual worlds. I will return to this in Section 4.

(C) ***Causation Is Chance Raising***

Probabilistic approaches to the metaphysics and epistemology of causation rely on the idea that causes raise the probability of the effect occurring. In other words, the probability of the effect-event occurring after the cause-event is greater than the effect-event occurring alone. This basic idea has been used to develop sophisticated models of causation that have been influential in the social sciences where statistical data is used to infer causal connections. If this aspect of causation is true for virtual events, then it would follow that the probability of a virtual-effect occurring is greater following a virtual-cause than by itself. With respect to the Mario example, this would mean the following must be true for the virtual world:

$$P(M_T/M_M) > P(M_T)$$

where "P" is probability, "M_T" is the state of Mario being tall, and "M_M" Mario collecting a mushroom. To fully demonstrate that this does in fact hold would require an empirical study of the game to collect values for the various variables. However, it should be easy to see that this equation does in fact hold. Suppose that the probability of Mario being tall at any given time in the game is n. We know that $n < 1$ because Mario always begins the game in his small size. Since the probably of Mario being tall after collecting a mushroom is 1, then clearly $P(M_T/M_M)$ is going to be greater than n. Just like physical causation, virtual causes therefore raise the probability of virtual effects occurring.

(D) ***Causation Is Asymmetric and Transitive***

Causal relations are asymmetric: pressure increase in the atmosphere causes a swing on the barometer needle, but not vice versa. If virtual events are causally connected, then they too need to be asymmetric. We have already seen that the program on the computer that outputs events is critical in providing the connection between them. At an abstract level, a program is a function that takes values as input and provides values as output. The relationship between virtual events is therefore asymmetric provided the function the program represents is also asymmetric. If we look at some of the examples from *Super Mario Bros.,* we can appreciate this is highly likely to be the case. The program for the game takes "Mario collecting a mushroom" as input and outputs "Mario getting taller," but not vice versa. It takes "Mario collecting 100 coins" and gives "Mario getting an extra life," but not vice versa. Asymmetry is therefore another property that both virtual and physical causation share.

A similar claim is true for transitivity. At the end of each level, Mario must defeat the boss "King Koopa." This is achieved by Mario jumping on a switch. This causes the drawbridge beneath King Koopa to collapse leading him to fall into a pit of lava. It seems reasonable to say that the cause of King Koopa falling into the pit was Mario jumping on the switch. Yet one can see this is merely the first part in a chain of events between Mario jumping on the switch, the drawbridge retracting, and King Koopa falling into the pit. As programs can be thought of as functions, then transitivity for virtual causation holds provided those programs can be combined such that the output of one function requires the calling and implementation of another function.

(E) ***Causation Can Occur Between Absent Events***

One controversial property of causation is that of absent causation or causation by omission. Sometimes an event is said to have been caused by the absence or privation of some other events. Take the example of Suzy's plants that died after she neglected to water them. It seems natural to say that the cause of Suzy's plants dying was that she failed to water them. One interpretation of absent causation is that the absent cause is one among many other background causes or conditions that are needed for the effect to occur. If the effect here is instead interpreted as "being alive," then the background conditions are sunlight, CO_2, and water. Removing water therefore removes one of the necessary conditions for keeping the plant alive.

However one chooses to interpret absent causation, it cannot be denied that we attribute causation when events are missing. The same is true for virtual causation. Just like physical causation, virtual causation takes place in a world that is continually being produced and sustained. In the case of virtual worlds, this is carried out by the computer. Most virtual worlds are designed such that they will continue to produce output in the absence of any user action or new input (so-called "looping"). This allows for many situations where absent causation is present. In *Super Mario Bros*., Mario must complete a level before a timer runs out. The timer is set at the beginning of each level. When the timer reaches zero, Mario will lose a life. If a player does not complete the level in time, it seems natural to say that Mario lost a life *because* he did not reach the level's end in time. This an example of absent causation and provides another positive analogy between physical and virtual causation.

3.3 The Negative Analogy

So far, we have seen that relations between virtual events that appear causal have many of the same properties as causal relations in the physical world. According to Hesse's account, the strength of an argument from analogy depends also on taking into consideration properties or aspects that are different. In this subsection, I review three potential negative analogies.

(F) ***Virtual Causation Is the Result of a Computer/Physical Causation Is the Result of Nature***

Perhaps the most obvious difference between virtual and physical causation is that virtual causation is the result of a computational process, whereas physical causation is not. How relevant is this difference? Does it weaken the argument from analogy? We can return to Chalmers' remarks about skepticism in the Matrix. Chalmers argues that if it turned out we had spent our entire lives in a Matrix, we would be compelled not to say that those experiences were of unreal objects and processes, but that they were virtual or computational objects and processes. If the universe is a computer (or the output of a computer), would we be willing to say that the causation we perceive is not real? If Chalmers' reasoning is correct, we would say that it is real causation but *virtual*, rather than *physical*, causation. The fact that it is computational should not, therefore, diminish it as a genuine form of causation.

It can also be claimed that (F) is not really known to be part of the negative analogy. There are many fundamental theories of reality that posit the basis of physics as computational processes of some kind or another. Computational physics is a live option for the metaphysics of our world, and some of the most successful physicists have proposed computational physics as a solution to interpreting some of the stranger aspects of the quantum world.[13] As a result, therefore, (F) is either irrelevant to whether virtual events can be causally related or is not in fact part of the negative analogy.

(G) ***Virtual Causes Are Artefacts/Physical Causes Are Not***

A similar but more general concern is that virtual causation is the result of human activity and an artefact, whereas physical causation is not. Another way to put this is to say that mankind has no choice or power to decide which events are causally connected in nature. We cannot just decide, for example, that one event E_1 causes another event E_2 or that some object or another has specific causal powers. We must work with what nature gives us, and it takes a lot of effort to discover (not create) causal powers. Things are clearly different with virtual causation. Though virtual worlds may take inspiration from nature when deciding on the causal relations that exist, the designer is only limited by their imagination. In *Super Mario Bros.*, it was the designer Shigeru Miyamoto who decided that mushrooms should have the power to make Mario grow taller. However, Miyamoto could easily have decided that mushrooms have the opposite effect or any number of other effects in the game.

As with the previous negative analogy, the question becomes whether this is a difference that is relevant. That is, whether being an artefact or not is a crucial property of causation. To show that it is not, consider the following thought experiment. Suppose in the future, some humans develop a mutation that gives them the ability to bestow novel causal powers on objects. Suppose one of these mutants gives gold the ability to heal individuals instantly from the common cold virus. Would we not say of somebody who wore a gold ring and was healed from the cold virus that the cause of their recovery was the gold ring? I think it is natural to say that it was. Whether or not a relation is causal, therefore, does

[13] Most notable examples include Wheeler (1990), Zuse (1970), Stephen Wolfram (2002), Seth Lloyd (2006), Gell-Mann (1987), T' Hooft (2016).

not depend on whether it is manmade or natural. As a result, being manmade seems not to be a relevant property for making inferences about causation.

(H) ***Virtual Causes Depend on Physical Processes/Physical Processes Do Not Depend on Virtual Processes***

A final negative analogy arises out of the fact that virtual events, objects, properties, and causal relations depend on physical processes occurring in the hardware of the computer. In Chalmers' version of realism, the virtual objects are identified with specific data structures. But this relation of ontological dependence is only one way: facts about virtual goings on depend on physical facts, but physical facts do not seem to depend in any way on facts about virtual events. This might be put by saying that virtual events are ontologically inert—we could remove all reference to them and still be able explain every physical state and event that occurs in the universe without gaps.

It might be questioned whether this observation is in fact true. There are cases where facts about what happened inside a virtual world have had very real physical world consequences, and these consequences cannot be explained only by reference to the physical basis of the goings on inside the computer. For example, there have been reported cases of "virtual sexual harassment" where one person's avatar has made unwanted actions of a sexual nature towards another user's avatar.[14] This created real-world physical feelings of pain, embarrassment, and trauma. If asked to explain what it was that caused these physical feelings in the victim, descriptions of the physical basis of the virtual world in terms of silicon chips, switches, and electrical charges would miss something out. Hence even if it is correct that the ontological basis of virtual objects and events are physical ones on a computer, this does not imply that virtual objects and events are ontologically redundant and facts about them are essential to explaining some physical events.

To summarize the findings of this section, there is a strong positive analogy between events that are connected in virtual worlds with those that are causally connected in the physical world. This positive analogy provides good *prima facie* evidence to suggest that such connections between virtual events may indeed be a genuine form of causation. I now turn to show how, at least on one influential theory of the metaphysics of causation, virtual events can have genuine causal connections.

4 Mechanisms and Virtual Causation

4.1 Glennan's Mechanical Theory of Causation

David Hume famously argued that although we can perceive regularity, we can never perceive a tie or necessary connection between causes and their effects. Most

[14] Julie Carry Wong (2016)

subsequent thinkers about causation have tended to agree, although they make different implications about what this means for the nature of causation. Stuart Glennan has argued that Hume's claim about the failure to perceive connections between causes and effects is overstated. Whilst he agrees that no connection can be found in the causal processes at the level of fundamental physics, most cases do possess a kind of connection that can be observed. This connection is not a mysterious power or necessary connection but a physical mechanism:

> When I claim that some event causes another event, say that my turning the key causes my car to start, I do not believe this simply because I have routinely observed that turning the key is followed by the engine starting. I believe this because I believe that there is a mechanism that connects key turning to engine-starting. I believe that the key closes a switch which causes the battery to turn the starter motor and so forth. Furthermore, this is not a "secret connexion". I can look under the hood and see how the mechanism works. (1996, p. 50)

Mechanisms are behind not just the causal processes in machines but also nature. Here Glennan appeals to a more general conception of mechanism. He offers the following definition:

> **Mechanism:** A mechanism underlying a behavior is a complex system which produces that behavior by the interaction of a number of parts according to direct causal laws. (1996, p. 52)

According to Glennan, whenever there is a causal connection between two events, this is because there is a mechanism connecting them. A few remarks are in order regarding Glennan's general definition of a mechanism. A mechanism is always associated with a particular *behavior* or *regularity*. So, if the behavior is the constant flowing of blood through the body, then the right mechanism associated with this includes the parts of the cardiovascular system (heart, veins, arteries, capillaries, etc.,) and the laws governing their interaction with one another. It might be said that by appealing to laws of interaction the mechanistic approach presupposes the concept of causation and will therefore be circular. However, the causal laws that govern the behavior of the parts of the mechanism are themselves reducible to behaviors that are the result of further mechanisms embedded within (2010, 2017).

Glennan is keen to distinguish his idea of mechanism from that associated with the mechanists of the seventeenth century, such as Hobbes, Descartes, and Boyle, who assumed that the physical behavior we see in the world can be explained as the result of corpuscles colliding with one another. Glennan's mechanisms make no assumptions about what the basic kinds of objects are or the type of interactions that they have. For Glennan, behavior that is the result of electromagnetic or gravitational forces is just as "mechanical" in his sense. Because Glennan broadens what is traditionally thought of as "mechanical," it raises the prospect of explaining causation in a wider range of natural phenomena, including that which arises because of electromagnetic phenomena. This makes Glennan's mechanical account ideally suitable to explaining causation in computational processes including virtual worlds. As I shall now argue, Glennan's mechanical definition of causation allows us to show that events in virtual worlds can be causally connected in much the same way as physical events.

Before that, it is worth commenting on one important issue with the mechanistic approach to causation. Glennan believes all causal processes can be explained via mechanisms between causes and effects except for fundamental physical processes. It might be said that this shows the mechanical approach is incomplete and needs supplementing with some other notion, such as regularity or counterfactual dependency. At the most fundamental level, it seems we cannot posit mechanisms to explain the most basic forms of interaction because, by definition, there are no further parts by which the action could be explained. How should we explain these cases? Glennan's preferred choice is to accept the irreducibility of fundamental laws without being committed to precisely how they should be analyzed leaving it an open empirical question whether it is "mechanisms all the way down" (2017, 188). We might also suppose at the most fundamental level the behavior is so strange that applying the concept of causation to it would be incorrect (2017, 93). As most causal claims made involve non-fundamental, higher-level phenomena, the mechanical theory works perfectly well for these kinds of claims. As we shall now see, since virtual events are not fundamental, we can explain their causal connection by appealing to more basic mechanisms in a computer. I agree, therefore, with Glennan that we do not need to take a stance on the issue of what causation amounts to at the level of fundamental physics to utilize the benefits of the mechanical approach.

If virtual events are connected via a mechanism, then this would imply that they could be causally connected. To show this requires identifying the mechanism in question and showing that it has the right propertiesGlennan associates with causal mechanisms, namely, that it is productive of regular behavior and is composed of parts that interact in accordance with a causal law. Of course, there are other theories of causation, and so the reality of virtual causation does depend on the success or failure of the mechanistic approach. It is an interesting question whether other metaphysical theories support causal relations in virtual worlds. For example, the Salmon-Dowe conservative quantities view might also recognize virtual causation if *information* is taken as a basic quantity that is conserved during computations. On the other hand, an account like that of Armstrong that ties causal relations to natural necessity (N) would likely imply that virtual causation is not real, as only natural properties connect by N are causal. Whilst this is an important question to explore, I will not attempt to do so here. Instead, I will focus exclusively on Glennan's version of the mechanistic theory. If this can be shown to imply that causal relations in virtual worlds are genuine, then it provides at least one metaphysical underpinning for virtual realism.

4.2 Mechanisms, Programs, and Functions

Virtual events are outputs of a computer and so whatever mechanism is responsible for their connection must belong to the system hardware and its programing. To illustrate, let us return to our *Super Mario Bros*. example. It seems reasonable to believe that there is a causal connection between the following virtual events:

(E_1) Mario collecting a mushroom
(E_2) Mario getting taller

where E_1 is the cause and E_2 is the effect. The first thing we need to recognize is that the descriptions we have given for these events tell us very little about the computational processes responsible for them. These descriptions exist at the level of the user's experience that can vary between each play of the game. Behind this is the computer code, the data structure, which in turn is embodied in the electrical components of the computer. Hence, to understand the mechanism, we first need to know what computer code provides the virtual events E_1 and E_2 and how they are connected. *Super Mario Bros*., like almost every modern video game, consists of a stored program that can be read and executed on dedicated hardware. The program contains instructions that tells the hardware what visual experiences to generate (such as images and sounds) in response to further instructions and input provided by the user. The character of Mario is represented by a two-dimensional bitmap image called a "sprite." Various versions of the sprite are stored in the program and can be set as the image for Mario depending on various factors. Sprites are also generated for other objects and enemies.

At the level of the program, "Mario collects a mushroom" occurs when the value of the location of the sprite for Mario overlaps with the sprite for a mushroom. This happens when the player directs the sprite to a location where a mushroom is present. At this point, the following function or subroutine is called:

Shroom_Flower_Pup:			
d820: ad 56 07	lda	PlayerStatus	;if player status = small, branch
d823: f0 1b	beq	UpToSuper	
UpToSuper:			
d840: a9 01	lda	#$01	;set player status to super
d842: 8d 56 07	sta	PlayerStatus	
d845: a9 09	lda	#$09	;set value to be used by sub-routine tree (super)

(McFadden, 2020).

If we examine the lines of code, we can see how the program responds to the instructions and changes the player's status. For instance, when the sprite of Mario is moved to a specific position and jumps, the subroutine "Shroom_Flower_Pup" is called. This subroutine can be thought of as a mini-program stored in the computer or game cartridge. When this subroutine is invoked, the first thing it does is check the player's status (i.e., the height of Mario). The default status is set to "small" when the game is launched. After that, an "if-else" condition is used to update the player's status: if the player's status is "small," then a subroutine "UpToSuper" is called when the player collected a mushroom or other powerup; otherwise, nothing happens, and the computer will proceed with its next task. Once "UpToSuper" is invoked, the player's status will be changed to "super," which causes the height of Mario to become tall, and a new bitmap image for the sprite is sent to the visual output.

By looking at the code, we can see that there is a program such that whenever Mario collects a mushroom (and in his small size), he will be supersized and rendered tall. This is enough to guarantee that there is a lawlike, functional relationship, between the virtual events E_1 and E_2. It is not enough yet, however, to establish a causal relationship. What we have described is a program written in assembly language. This is not a physical mechanism but rather the instructions for implementing the right kind of relationship. Fortunately, we know that such a mechanism does in fact exist, and the developers of *Super Mario Bros*. found a way to create it on the hardware for this video game. It is not necessary to go into precise detail about the hardware specifics of the video game console to demonstrate that it meets Glennan's definition of a causal mechanism.

The video game console that runs this program is a variation on the standard von Neumann machine and includes the following parts: a memory, a central processing system (CPU), a control unit, an input device, and an output device. In our example, the observed behavior is the regular association between Mario collecting a mushroom and Mario getting tall. This behavior is produced by the parts of the von Neumann machine. Those parts are of course physical in nature and composed of electrical circuits, transistors, capacitors, diodes, signal generators, etc. Those components have been arranged in such a way that they produce the desired behavior, and this is achieved by manipulating electrical charges in accordance with the laws of electromagnetism. Because of this, we know that whenever those components are arranged in the right way, then they will produce the observed behavior associated with the causal connection.

4.3 What Is Virtual Causation?

We are now able to say more precisely what virtual causation is. This can be done by following closely Glennan's definition for physical causation.

> **Physical Causation**: Event A is causally connected to event B if, and only if, there is a mechanism such that, in the right circumstances, when A occurs, the mechanism will produce B in virtue of its parts and causal laws.

The proviso "in the right circumstances" is needed because a mechanism will not always produce an effect given the cause if there is some other mechanism that interferes or overrides it or if the mechanism itself responds differently to the presence of other events. This is the familiar issue of causes being context specific and is one of the reasons why causal claims can rarely be reduced to brute universal generalizations. This standard definition for physical causation can be extended to cover virtual causation when we interpret the events as virtual:

> **Virtual Causation:** Virtual event A is causally connected to virtual event B if, and only if, there is a mechanism in a computer such that, in the right circumstances, when the computer reads the code that produces virtual event A, a program is run, and virtual B is produced as output.

Going back to our example, we can see that this is what happens in events E_1 and E_2. When the machine reads the code for event E_1 (Mario collecting a mushroom), the subroutine "Shroom_Flower_Pup" is run, and, subject to the right conditions (when Mario is "small"), another piece of code is implemented that ultimately produces the virtual event E_2 (Mario getting tall) as output. In this example, two separate subroutines are run, but we can think of each of these as parts of the program that produce the required observed behavior. In this definition I have replaced Glennan's reference to parts and causal laws with "programs." However, it should be appreciated that the parts and causal laws are still there, since they provide the physical mechanisms in questions that realize or implement these programs. By referring to programs, I hope to make explicit what is unique about virtual causation as a form of causation. Just like all forms of causation, it depends on a mechanism being present, but that mechanism is computational and can be described in terms of a program.

4.4 Objections

At this point, several objections might be raised. It might be said that the proposed definition is too weak to capture only those events that are causally connected because any line of code is "connected" in some way via a program to a particular output. Take the following example. Suppose Suzy always completes levels with exactly 12 coins. Now, anyone who has played *Super Mario Bros*. knows that there is no causal relationship between collecting 12 coins and finishing a level. That Suzy always completes a level with 12 coins is a mere accidental regularity. Nonetheless, there is arguably a program that connects them. If we take as input the value for the variable "no. of coins" as 12 as well as the input provided by Suzy's movement of the Mario sprite to the end of the level, then the machine clearly contains a program that produces the "level complete" screen as output. However, we would not say that gaining 12 coins was a *cause* of Suzy's completing the level.

I do not think that this problem should be considered fatal to the proposal as it is a manifestation in the virtual world of a more general problem for mechanistic accounts of causation known as the "problem of causal relevance." Consider an analogous case. Suzy is playing football and in the last few moments of the game strikes the ball into the back of the goalkeeper's net. Her team goes into the lead and wins the match. According to the mechanistic approach, Suzy's striking of the ball is what caused it to go into the net because there is a mechanism connecting her striking of the ball with the ball going into the net. Suppose also that when Suzy strikes the ball, she marks it with the mud on her boot. This marking of the ball is also connected via a mechanism with the ball going into the back of the net. But we would not say that the marking of the ball is causally connected to the ball going into the back of the net.

There are several different ways advocates of mechanisms can respond to the problem of causal relevance. One idea is to combine mechanisms with counterfactuals and use the counterfactual dependency between causal events to sort out the relevant from the irrelevant ones. This approach has been taken by James Woodward (2002, 2011) under a manipulationist framework. Glennan himself does not think that mechanisms need supplementing with counterfactuals and instead attempts to explain away supposed examples of irrelevant causation as either falling into two groups of cases. In the first group the supposed irrelevant event was in fact causally relevant, but its causal contribution is overlooked (2017, p. 195). For example, in the case of Suzy's football match, for all we know the addition of the mud to the ball affected its trajectory through the air and therefore did form part of the "input" into the mechanism that produced the goal. In the second group of cases, Glennan complains that we make an unfair comparison between the specifics of a concrete token events with the idealized or abstract description of the effect (2017, p. 196). When we describe the effect from an abstract perspective, e.g., the ball going into the back of the net, we already presuppose that properties of the ball (such as whether or not it is marked) do not make a difference to the effect occurring.

Whether Glennan's response to the problem of irrelevant causation succeeds or whether mechanisms need supplementing with counterfactuals to pick out the genuine from the non-genuine causal events will not be answered here. Whichever solution works for causation ordinarily understood, we can apply to virtual causation as well. We have already seen that virtual causes appear to support counterfactual conditionals and so if this provides a satisfactory solution to causation in the physical world a similar approach can be taken for virtual worlds.

A second problem comes from McDonnell and Wildman (2019). They question whether virtual events and objects really do have causal powers over-and-above those of the hardware on a computer. Taking Chalmers's view that virtual objects and events are real but grounded in physical states leads to what they call an "exclusion problem". Suppose a player uses the hextech gunblade to inflict damage on an opponent. We can give different causal stories to explain what happens to the health of an opponent's avatar. One story runs entirely through the digital states on a computer, and another runs through digital states and then flips to describing virtual objects and events. Which is correct? To accept both leads to causal overdetermination. According to McDonnell and Wildman, as the virtual metaphysically depends on, or is grounded in, the physical, therefore, "the digital (non-virtual process) does all the causal work, leaving none for the virtual processes to do" (2019, 384). If pushed, I would side with Chalmers's (2019) own response to this objection. According to Chalmers, ascribing causal powers to both virtual and digital phenomena is no more different than ascribing causal powers to both biological and chemical phenomena. We accept the existence of both human organs and the molecules that they are made of, and we ascribe causal powers to both. Biologists will typically say that the causal interactions of the molecules "gives rise to" the causal powers of the organ without feeling a need to choose one over the other. It seems reasonable to suppose we can take a similar line with virtual and digital causal powers.

A final problem arises because it might be said that this definition ties virtual causation too closely to the architecture of the von Neumann machine. The definition presupposes that the virtual world is generated from a machine that reads lines of codes, stores code in a register, looks up addresses, calls subroutines, etc., which are all computational activities associated with variations on the von Neumann machine. But there are other models of computation that also appear to generate virtual worlds with causal connections. Recently, Holly Andersen (2017) has brought attention to the causal relations that manifest in cellular automata (CA). What is interesting is that Andersen argues that productive accounts of causation, such as the mechanistic approach of Glennan, are insufficient to explain all the causal relations that arise in these worlds. To help illustrate her point, consider what goes on in the most famous example of a CA—John Conway's *Game of Life* (GoL). This CA consists of a two-dimensional array of cells, and each cell can exist in one of the two possible states (usually represented as "off/on," "0/1," or "dead/alive"). The state of any cell at time *t* depends on the states of the neighboring cells as time *t-1*. The state of a cell is calculated according to the rules:

(1) If the cell state at *t-1* was 0 (dead), the cell state becomes 1 (alive) if exactly three neighbors were 1 (alive) at *t-1*.
(2) If the cell state at *t-1* was 1 (alive), the cell state is still 1 if either two or three neighbors were 1 (alive) at *t-1*.
(3) If the cell state at *t-1* was 1 (alive), the cell state becomes 0 (dead) if either fewer than two or more than three neighbors were 1 (alive) at *t-1*.

(Berto and Tangliabue, 2017)

When the GoL is run, complex stable structures emerge that seem to move across the screen and interact with one another. One such structure is a "glider." If gliders collide at an angle of 90°, they will produce a two-by-two block of cells in their place.[15] What is interesting about the GoL is that these objects and their interactions were not programed into it. They emerge from the three simple rules above. According to Andersen, this shows that productive accounts cannot fully explain causation in virtual worlds as there is no obvious mechanism or program in the three rules that connect them to the behavior of the gliders. We could say that since the gliders emerge from these rules, then there is such a program. But since *all* the behavior of the objects in GoL emerges from these three rules, this would fail to distinguish the causal from the non-causal interactions and take us back to the causal relevancy problem.

It might be argued in response to Andersen that GoL (along with all other existing examples of CA) are not true CA in the sense that they are typically designed

[15] For a detailed examination of the objects in GoL and their interactions see William Poundstone (1985)

and run on von Neumann machines. However, this response would miss the point. Although CA are for the most part hypothetical forms of computation, we know theoretically that any program that can be run on a von Neumann machine can also be run on a CA and vice versa (the so-called Church-Turing thesis). This means that not only should causal relations emerge in CA as they do in von Neumann worlds, but the causal relations should also, in principle, be one and the same. If it is a fact that there is a causal relation between Mario collecting a mushroom and getting tall in a von Neumann machine, then such a causal fact should exist in a CA world as well.

Glennan points out that his mechanisms are hierarchical such that parts of mechanisms can be yet smaller mechanisms. Eventually we know that this bottoms out for him in the fundamental entities and laws of physics. In terms of virtual worlds, mechanisms "bottom out" in the electronic components of the computer, or what other hardware is used to implement the program. But just like the physical world it stands to reason that some of these programs can be combined to produce new "emergent" virtual objects and events. This is what I believe is happening in the case of CA. Here, all emergent virtual objects arise from the basic three laws, but these objects can combine with other objects and the basic three laws to produce new objects and events. How we identify these programs will be different, just as running the same program on different hardware is different for standard von Neumann machines. If *Super Mario Bros*. was reproduced on a CA, there would still exist causal relations between Mario collecting a mushroom and him getting taller. If we investigated the underlying code, it would of course look very different, as too would the hardware implementing it. But we would still find that there is a connection between the code or codes that produce Mario collecting a mushroom and the code of codes that produce Mario getting taller, and this connection is clearly one mediated by a program—even if that program is composed of other smaller programs.

5 Conclusion

I set out to show that what appears to be causal relationships in virtual worlds are in fact genuine types of causation on a par with causal relationships we see in the physical world. I have given two arguments for this. The first is based on a similarity between the commonly held properties of the causal relation with relations in virtual worlds. This analogy is sufficiently strong to at least suggest that virtual causes are genuine. However, it is unlikely to convince virtual antirealists, those who believe that what goes in virtual worlds is mostly fictional. To fit this gap, I have argued that virtual causes come out as real on at least one influential theory of causation: Glennan's mechanistic view. I recommend, therefore, that advocates of virtual realism avail themselves of this theory to place it on a more convincing metaphysical foundation.

Author's Contributions The corresponding author is the sole contributor.

Funding The study did not receive any funding that needs acknowledgement.

Availability of Data and Material N/A.

Declarations

Ethics Approval and Consent to Participate N/A

Consent to Publication The author gives consent to publication.

Competing Interests There are no financial or non-financial conflicts of interest to declare.

References

Andersen, H. (2017). Patterns, Information and Causation. *Journal of Philosophy, 114*(11), 592–622.

Armstrong, D. M. (1983). *What Is a Law of Nature?* Cambridge University Press.

Bedau, M. (1997). Weak Emergence. In J. Tomberlin (Ed.), *Philosophical Perspectives: Mind, Causation, and World* (pp. 375–399). Blackwell.

Beraldo-de-Araujo, A., & Baravalle, L. (2017). The Ontology of Digital Physics. *Erkenntnis, 82*(6), 1211–1231.

Berto, F., & Tagliabue, J,. (2017). Cellular Automata. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab: Stanford University. https://plato.stanford.edu/archives/fall2017/entries/cellular-automata/.

Brey, P. (2014). Virtual Reality and Computer Simulation. In R. Sandler (Ed.), *Ethics and Emerging Technologies* (pp. 315–332). Palgrave Macmillan).

Calleja, G. (2014). Immersion in Virtual Worlds. In M. Grimshaw (Ed.), *The Oxford Handbook of Virtuality* (pp. 222–236). Oxford University Press.

Cavazza, M., Lugrin, J., & Buehner, M. (2007). Causal Perception in Virtual Reality and Its Implications for Presence Factors. *Presence: Teleoperators and Virtual Environments*, 16 (6), 623–42.

Chalmers, D. (2017). The Virtual and the Real. *Disputatio, 9*(46), 309–352.

Chalmers, D. (2019). The Virtual as the Digital. *Disputatio, 11*(55), 453–486.

Cogburn, M., & Silcox, M. (2014). Against Brain-in-a-Vatism: On the Value of Virtual Reality. *Philosophy & Technology, 27*(3), 561–579.

Dowe, P. (2000). *Physical Causation*. Cambridge University Press.

Dretske, F. (1977). Laws of Nature. *Philosophy of Science, 44*, 248–268.

Elder, A. (2017). "Figuring Out Who Your Real Friends Are" in *Experience Machines: The Philosophy of Virtual Worlds* ed. Mark Silcox. Lanham: Rowman & Littlefield International, 87–98.

Gell-Mann, M. (1987). "Simplicity and Complexity in the Description of Nature". Talk delivered to the The Caltech Assoication, Pasadena, October 1 1987.

Glennan, S. (1996). Mechanisms and the Nature of Causation. *Erkenntnis, 44*(1), 49–71.

Glennan, S. (2010). Mechanisms, Causes, and the Layered Model of the World. *Philosophy and Phenomenological Research, 81*, 362–381.

Glennan, S. (2017). *The New Mechanical Philosophy*. Oxford University Press.

Glennan, S., & Illari, P. (2018) “Introduction” in *The Routledge Handbook of Mechanisms and* Mechanical *Philosophy* ed. Stuart Glennan & Phyllis Illari. London: Routledge, 1-11.

Glymour, C., Scheines, R., Spirtes, P., & Kelly, K. (1987). *Discovering Causal Structure: Artificial Intelligence, Philosophy of Science, and Statistical Modelling*. Academic Press.

Glymour, C., & Cooper, G. (1999). *Computation, Causation, and Discovery*. MIT Press.

Goodman, N. (1978). *Ways of Worldmaking*. Indianapolis: Hackett.

Grabarczyk, P., & Pokropski, M. (2016). Perception of Affordances and Experience of Presence in Virtual Reality. *Avant, 7*(2), 25–44.

Heim, M. (1993). *The Metaphysics of Virtual Reality*. Oxford University Press.

Heim, M. (1998). *Virtual Realism*. Oxford University Press.

Hesse, M. (1966). *Models and Analogies in Science*. University of Notre Dame Press.

Hoyet, L., McDonnell, R., & O' Sullivan, C. (2012). Push It Real: Perceiving Causality in Virtual Interactions. *ACM Transactions on Graphics, 31*(4), 901–909.
Lloyd, S. (2006). *Programming the Universe*. Belknapp Press.
Ludlow, P. (2017) "Cypher's Choices: The Variety and Reality of Virtual Experiences" in *Experience Machines: The Philosophy of Virtual Worlds* ed. Mark Silcox. Lanham: Rowman & Littlefield International, 13–32.
Pietrucha, D. (2017). "Intuition and Imaginative Failure" in *Experience Machines: The Philosophy of Virtual Worlds* ed. Mark Silcox. Lanham: Rowman & Littlefield International, 33–42.
Machamer, P., Darden, L., & Craver, C. (2000). Thinking about Mechanisms. *Philosophy of Science, 67*(1), 1–25.
McBain, J. (2017) "Epistemic Lives and Knowing in Virtual Worlds," in *Experience Machines: The Philosophy of Virtual Worlds* ed. Mark Silcox. Lanham: Rowman & Littlefield International, 155–168.
McFadden, A. (2020). "Super Mario Bros. Disassembly". https://6502disassembly.com/nes-smb/. Accessed 30–09–2021.
McDonnell, N., & Wildman, N. (2019). Virtual Reality: Digital or Fictional? *Disputatio, 11*(15), 371–397.
Nozick, R. (1981). *Philosophical Explanations*. Belknap Press.
Poundstone, W. (1985). *The Recursive Universe: Cosmic Complexity and the Limits of Scientific Knowledge*. Oxford University Press.
Ropolyi, L. (2015). Virtuality and Reality—Toward a Representation Ontology. *Philosophies, 1*(1), 40–54.
Ricksand, M. (2020). Walton, Truth in Fiction, and Video Games: A Rejoinder to Willis. *The Journal of Aesthetics and Art Criticism., 78*(1), 101–105.
Salmon, W. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
Schultze, U. (2010). Embodiment and Presence in Virtual Worlds. *Journal of Information Technology, 25*(4), 434–449.
Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, Prediction and Search*. MIT Press.
Studt, E. (2021). Virtual Reality Documentaries and the Illusion of Presence. *Studies in Documentary Film., 15*(2), 175–185.
T'Hooft, G. (2016). *The Cellular Automation Interpretation of Quantum Mechanics*. Heidelberg: Springer.
Tavinor, G. (2009). *The Art of Video Games*. Wiley-Blackwell.
Tooley, M. (1977). The Nature of Laws. *Canadian Journal of Philosophy, 77*(4), 667–698.
Walton, K. (1990). *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press.
Weijers, D., DiSilvestro, R. (2017). "The Morality of Experience Machines for Palliative and End-of-Life Care" in *Experience Machines: The Philosophy of Virtual Worlds* ed. Mark Silcox. Lanham: Rowman & Littlefield International, 183–202.
Wheeler, J. A. (1990). Information, Physics, Quantum: The Search for Links. In W. H. Zurek (Ed.), *Complexity, Entropy, and the Physics of Information* (pp. 354–368). Addison-Wesley.
Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media
Wong, J. C. (2016). "Sexual harassment in virtual reality feels all too real", *The Guardian*. Originally published October 26 2016. Available at https://www.theguardian.com/technology/2016/oct/26/virtual-reality-sexual-harassment-online-groping-quivr. Accessed 30–09–2021.
Woodward, J. (2002). What Is a Mechanism? A Counterfactual Account. *Philosophy of Science, 69*, S366–S377.
Woodward, J. (2011). Mechanisms Revisited. *Synthese, 183*, 409–427.
Zuse, K. (1970). *"Calculating Space", MIT Technical Translation AZT-70-164-GEMIT, Massachusetts Institute of Technology (Project MAC)*. Mass.
Wheeler, B. (2020) Truth Tracking and Knowledge from Virtual Reality. *Logos & Episteme 11(3),* 369-388.
Steinhart, E. (1998). Digital metaphysics. In T. Bynum & J. Moor (Eds.), The digital phoenix: How computers are changing philosophy (pp. 117–134). New York: Blackwell.