



What does the gamer do?

Rebecca Davnall¹

© The Author(s) 2020

Abstract

The 'Gamer's Dilemma' is the problem of why some actions occurring in video game contexts seem to have similar, albeit attenuated, kinds of moral significance to their real-world equivalents, while others do not. In this paper, I argue that much of the confusion in the literature on this problem is not ethical but metaphysical. The Gamer's Dilemma depends on a particular theory of the virtual, which I call 'inflationary', according to which virtual worlds are a metaphysical novelty generated almost exclusively by video games. Actions performed in virtual worlds really belong to the kinds of action they appear to—'virtual murder' is a kind of murder. Inflationary theories are contrasted with 'deflationary' theories which, in effect, consider video games purely as systems for generating images, and thus the gamer as (merely) a consumer of media images. Inflationary theories struggle to explain why video games produce this unique metaphysical novelty; deflationary theories fail to do justice to the intuition that there is some significant difference between the gamer and the consumer of other media forms. In place of either, I sketch a theory of the gamer as performer, primarily by analogy with stage and cinema actors, which I suggest captures more of the moral complexity of the gamer's action.

Keywords Gamer's dilemma · Video games · Virtual ethics · Virtual metaphysics · Performance

Morgan Luck's 'Gamer's Dilemma' (2009) crystallised many of the key issues in virtual ethics.¹ Like all good philosophical problems, though, there is some confusion about precisely what the dilemma concerns. Claims about the ethics of the virtual are entangled with a variety of—sometimes incompatible—theories of action in virtual worlds.

Luck's original framing of the dilemma is as follows (please note that because one of the canonical examples used in the gamer's dilemma literature is virtual child sex abuse, this paper unavoidably includes frequent mentions of child sex abuse):

Is it immoral for a player to direct his character to murder another within a computer game? The standard response to this question is no. This is because no one is actually harmed as a result of a virtual harm ... Yet this argument can easily be adapted to demonstrate why virtual paedophilia might also be morally permissible, as no actual children are harmed in such cases. (2009, p. 31)

The dilemma for the gamer is whether to accept that there is a good argument for the permissibility of virtual child sex abuse, or to reject that argument and in so doing also render impermissible the virtual violence that typifies many video games. This framing assumes that the relationship between virtual child sex abuse and real child sex abuse is the same as between virtual murder and real murder; that the extent to which virtuality attenuates the moral significance of an act is independent of the content of the act.

Another, more technical, way of putting this is that for Luck the moral taxonomy of the virtual is isomorphic to that of the real, and virtuality merely attenuates it to some degree that is constant across the whole structure. The question of what virtual actions are permissible, then, comes down to how strong the relationship is between real moral kinds such as murder and child sex abuse and their virtual sub-classes. If the relationship is strong, then both virtual murder and virtual child sex abuse are wrong. If it is not, then both are permissible.

We need not follow Luck in framing this in terms of harm, though he is probably correct that the 'standard response' outside academia would do so. Most ethical systems agree that murder in the real world is wrong; each system can pose

¹ Luck is not the first to have discussed this problem; it appears at least as early as McCormick (2001).

✉ Rebecca Davnall
r.davnall@liverpool.ac.uk

¹ Department of Philosophy, University of Liverpool, Liverpool, UK

the question about in-game murder in their own terms. A Kantian (Tillson 2018) can argue that it transgresses a duty to respect life; a virtue theorist (McCormick 2001) can hold that it contravenes the same virtues as real-life murder.

My claim is that, against the tendency of the existing literature, Luck's challenge will not be answered by choosing an alternative to his utilitarianism. The problem is not ethics or psychology so much as the metaphysics of morals: what are the ontological conditions necessary for the application of ethical concepts to actions, and do they apply in the case of virtual worlds?

If I am right, then we need to look closely at the assumptions Luck and others make about the nature of virtual worlds and how computers sustain them. Robert Seddon (2012) persuasively argues that the word 'virtual' is obscure and obstructive; here I take up the investigation he proposes. Physically speaking, when a computer game is being played (in archetypal cases²), complex computer software responds to a user pressing buttons on an input device by generating a sequence of visual images on a screen, usually accompanied by sound and sometimes also haptic stimulation (e.g. vibrations in a gamepad or other input device). The question is then what to make of these button-presses and the resulting images.

According to one view, which I will call 'inflationary', these images constitute or represent a 'virtual world', an entity which is ontologically similar to the real world. In its strongest form, most evident in Luck's own writing, this position approaches something like modal realism (Lewis 1986), or the theory of virtuality more recently proposed by David Chalmers (2017), effectively claiming that every video game sustains a complete world, different to and ontologically separate from our own, such that things which are *true* in that world are *virtually* true in ours. So, if it is virtually true that a character in a game (that is being played in our world) murders another, then there is a virtual world in which it is true that the character has murdered someone. Assuming that moral truths are not world-relative, the action in the virtual world *just is* murder, and the player's instruction to the character to commit it implicates them in the act.

Aside from its obvious profligacy, this view runs into several problems. How do we know what objects are present, and what states of affairs obtain, in virtual worlds? Which actions in those worlds is the player involved in? And how are virtual worlds to be distinguished from a more familiar

category, fictional worlds? The gamer's dilemma would be implausible if applied to other media—we allow far more nuance with prose and cinema violence—and inflationary theorists must explain why games are different.

This last problem, in a different form, also emerges for the alternative to inflationary views, which I will call 'deflationary'. On this view, talk of virtual worlds is superfluous, and what matters, morally, is the images produced by a game, considered *as images*. For deflationary theorists, the gamer (merely) consumes representations in the same way that the reader of a book or watcher of a film does. Some deflationary theorists, e.g. Patridge (2010, 2013), might be comfortable taking the gamer's dilemma just as an instance of a more general problem in media ethics, but for others, such as Bartel (2011, 2020), games are distinctively problematic.

Several writers raise another set of moral concerns, about the development and publication of games which afford players violent options, but there isn't space to address these here. Nor will I spend any time on the classic moral-panic question of whether playing violent video games makes a person more likely to act violently in real life, except to differentiate the positions of some deflationary theorists from it. The empirical issue here is extremely murky, and the focus of the gamer's dilemma is on whether there is anything wrong with wrongdoing in virtual worlds *as such*.³

What I aim to show is that neither inflationary nor deflationary views really do justice to the ethics of the gamer's actions. The gamer's dilemma is not a completely new problem posed uniquely by virtual worlds as inflationary theories suggest. Equally, it is not reducible to conventional concerns about media consumption, as deflationary views tend to imply. A promising alternative, I argue, is to locate the gamer and the gamer's dilemma within the ethics of performance.

I begin by extracting Luck's implicit inflationary metaphysics from his own writing about the dilemma ('The Inflationary View'). I follow this with an analysis of two strategies for establishing an inflationary theory of games (i.e. a way of differentiating them from other fictive media), one grounded in the supposed 'agency' of the player ('Agency and Identification') and the other in the 'realism' experienced by players ('Representation and Realism'). All three approaches, I argue, share the fundamental claim that virtual murder is genuine murder in virtual worlds. I then survey deflationary views ('Deflationary Views'), positing a different dilemma; deflationary views end up denying that

² One problem that a lot of engagements with the gamer's dilemma have is that they overlook the extremely wide scope of the term 'video games', especially text-based games and visual novels. There may be reasons for thinking that games of these kinds merit different ethical treatment to more 'conventional' games, but in the interests of focussing on what has already been said about the dilemma, I shall leave these aside for now.

³ As McCormick (2001) has argued, absent a psychological effect much stronger than anything yet hinted at in the empirical literature, the psychological effects of any one individual of act of virtual violence must be regarded as trivial. Any 'media effects' argument will not depend on isolated incidents but on long-standing patterns better accounted for in terms of media consumption than virtual action.

the gamer's action itself is problematic, either because the gamer is no different to consumers of other media or because all wrongdoing actually falls on the heads of game designers and publishers. Finally ('Does the Gamer do anything?'), I suggest that an appropriate way to attend to the gamer's involvement in virtual wrongdoing might be to examine the gamer as a kind of performer; this account, I suggest, connects the gamer's dilemma to a more productive lineage of ethical debate than either inflationary or deflationary views.

The inflationary view

Inflationary assumptions can be subtle. To actually engage with the gamer's dilemma in the form it is conventionally presented, one has to implicitly accept some inflationary premises, and they are hard to spot in retrospect. We have already seen one hint, where Luck writes:

A player commits an act of virtual murder in those cases where he directs his character to kill another in circumstances such that, were the game environment actual, the actions of his character would constitute actual murder. (2009, p. 31)

In a footnote, Luck claims agnosticism about the metaphysics of virtuality and virtual environments, citing Brey (1999) 'if a narrow focus is required' (2009, p. 31), but the phrasing here is not as innocent as he takes it to be. Games have 'environments', and there is some sense in which these environments could be 'actual'. I take 'actual' here as an antonym of 'virtual', without the further implication that the virtual is a subcategory of the logically or physically possible.

From this, it seems that 'killing' is a kind which has both actual and virtual instances, such that virtual killing stands to virtual environments as actual killing to actual environments. This is reinforced by Luck's affirmation that features of the virtual environment moderate the ethical status of this killing-in-the-virtual-environment⁴ in just the same way as real environmental factors moderate real killing:

Note that our focus is upon murder, rather than killing in general. The difference being that, whilst the act of killing a person may be morally permissible, murder is not... Presumably a player directing his character, an Allied soldier, to kill an Axis soldier within the context of such a [game of WWII] battle, is not committing an

act of virtual murder. This is because, were the game environment actual, we would not, by most reasonable accounts, consider the soldier to be a murderer. (2009, p. 32)

This is the clearest example of the central assumption of the inflationary view, that things that happen in games belong, virtually, to exactly the actual ethical kinds they appear to. A 'virtual murder' involves an act that genuinely belongs to the category 'murder' occurring in a virtual environment, and is of interest to us whenever it occurs, in some sense, *at the player's behest*. I leave aside the long-contested question of the relationship between battlefield killing and murder.

From here, it is easy to generate the gamer's dilemma. What determines whether a virtual act—that is, a set of instructions delivered by the player to a character through the game's interface—is wrong is simply whether the resulting action of the character belongs to a morally bad kind. 'Murder' is a morally bad kind, and so is 'child sexual abuse', so directing a game character to do either is bad, *unless* the boundary of the virtual is sufficiently impermeable that no moral weight carries over from either act. This naturally invites the question of why so many people seem to believe that virtual murder as Luck defines it is morally unproblematic, but that virtual child sexual abuse is clearly wrong.

Another way to characterise this is to say that on inflationary accounts, the appearance or presentation of the act makes no difference to how it is to be classified or assessed. One of the first deflationary responses to Luck was Christopher Bartel's argument (2011) that what differentiates virtual child sex abuse from virtual murder is that virtual child sex abuse necessarily involves real child pornography. In his original statement of this claim, Bartel does not clarify whether he thinks that virtual child sex abuse should be counted as *producing* or *consuming* child pornography. Instead, he says:

the graphic depiction of a character—who is clearly depicted as an adult—engaging in sexual acts with another character—who is clearly depicted as a child—would count as an instance of child pornography. While these may be virtual instances of paedophilia, they are still actual instances of child pornography. (2011, p. 13)

To this, Luck, writing with Nathan Ellerby, responds:

Bartel states that virtual paedophilia "necessarily involves the depiction of sexual acts involving children". But this need not be the case. Suppose a computer game were created in which acts of virtual paedophilia were not depicted on screen, but nevertheless the game suggested (in such a way as to leave little

⁴ As Seddon complained, the instability of the language of the 'virtual' is already causing problems here—per Luck, 'virtual killing' is what the *player* does, and the character just kills, but here it seems necessary to specify that the killing being done by the character is confined to the game environment.

doubt) that such an act had occurred... Such a game might count as one in which players commit the act of virtual paedophilia, despite the fact the act itself is never depicted... We take it that many gamers would want to prohibit games that allowed players to commit non-depicted acts of virtual paedophilia. (2013, p. 231)

Luck and Ellerby clearly take the same to go for virtual murder. Referring back to Luck's original paper, they say:

Luck does not limit the dilemma to cases where virtual murder is graphically depicted. Luck does state that he is focusing "on those computer games, such as Grand Theft Auto, where clear instances of virtual murder are apparent." ... Clearness in this context picks out our ability to say that the act of virtual murder has occurred; it does not pick out the depiction of such acts. (2013, p. 231)

Neither Luck and Ellerby nor Bartel should be taken as using 'depiction' to refer only to visual depiction. Prose, narration and audio might also depict with or without visual images and still count for either argument. Instead, the disagreement is over how events in game environments are to be (ethically) classified. We can illustrate it by example of a couple of hypothetical games. Game A is an adventure game which starts with a prologue in which the main character trains his martial skills under his father's guidance. The prologue ends and the game's story skips forward a few years to a scene in which the main character confronts the local king and accuses "You murdered my father!", to which the king responds "So I did."

Game B is a cartoony side-scrolling platform game in the spirit of the *Super Mario* games. In it, the 'enemy' creatures are blobs of alien goop with eyes. When the main character jumps on them, they burst, splattering goo over the terrain, their eyes falling to the ground and going blank. However, if left on-screen long enough, the blobs eventually pull themselves back together and start chasing after the protagonist again.

Per Luck, the father in Game A is clearly murdered but it is *not* clear that the enemies in Game B are even killed—Luck's own example of a lack of clarity is the ghosts in Pac-Man, who always return after Pac-Man eats them (2009, p. 33). Bartel, on the other hand, would be more against Game B than Game A, because Game B is more 'graphic' (this is because Bartel's approach is deflationary—he focusses on the ethical significance of the image that exists in our world rather than the gamer's relationship to an act that occurs in a virtual world).

So the nature of a virtual act, for Luck, is not secured by what the images on-screen are like. Similarly, Luck denies that the nature of a virtual act is determined by whether the player recognises what is being represented. This denial

appears in his response to Rami Ali's attempt to dissolve the dilemma.

Ali acknowledges Luck's differentiation of wartime killing from murder, then goes on:

With virtual acts, however, a further complication arises. This is because virtual acts have two different contexts. There is the in-game context of the act, which is the context of the game character in its virtual world, and there is the gamer's context, which is the context of the gamer performing the in-game acts. (2015, p. 269)

Here, Ali breaks the link between the act in a virtual world and the virtual act (between what happens in the game and the player's involvement in bringing it about). According to Ali, the gamer can only be held accountable for what they take themselves to be doing, irrespective of how matters stand in the virtual world and from the perspective of its occupants. A gamer might play a particular game as an esport or in some other format where they mostly ignore the game's narrative content. To such a gamer, overcoming the game's enemies is a matter of navigating and removing obstacles rather than jumping on or shooting living creatures, and so holding them accountable for 'virtual murder' is unjust.

For Luck, though, this picture is incorrect. The nature of the player's actions is fixed by the nature of the acts occurring in the virtual world, irrespective of the player's interpretation of them. This is implicit in Luck's assertion that 'a "clear instance" [of virtual murder] is simply an instance where it is apparent that the act of virtual murder has occurred' (2013, p. 231). It is made more explicit when Luck, responding to Ali, posits a game in which 'one plays a convicted sex offender, with the ability to drag children into your car, grope them, and then throw them out', and a player of said game, Cathy, who 'is quite an innocent-minded person [to whom] it never occurs... that this is what is going on'. Cathy interprets the on-screen actions as something other than child sexual assault, but Luck insists: 'This is an instance of virtual child molestation. Why? Because the circumstances of the game... are such that, where the game world actual, the act would be actual child molestation.' (2018, p. 158).

The obvious way to interpret this is inflationary—that is, to say that Luck takes the game to be an accurate but incomplete representation of a virtual world. This representation is uniquely authoritative in the sense that even where there is sufficient ambiguity that a player could form a different interpretation of the on-screen events, a player who did so would simply be mistaken, as Cathy is, about what is actually going on.

One problem with this is that, since Luck has insisted that the presentation of the game's events does not matter,

no game can ever really exercise this authority. No game's fiction is complete enough. Imagine a game in which the player controls the actions of a nurse. The nurse enters a hospital room and gives a patient an injection, after which the patient dies. It would be possible for the player (at least, without further context), to interpret this in two different ways. The player can interpret the nurse as a medical professional with unfortunate timing, whose injection did not have the desired effect of prolonging the patient's life. Or the player can interpret the nurse as a disguised assassin, injecting the patient with poison.

Either interpretation seems equally valid. Nor can more information actually settle the question. Say the game has a later scene in which an autopsy reveals that the syringe was loaded with poison. The nurse's guilt is still not secure—the poison might have been put there by someone else. But perhaps the nurse was negligent in checking for contaminants in the syringe? Perhaps the autopsy report has been falsified, because the pathologist is the actual murderer and is covering his tracks. Perhaps the nurse did indeed intend to poison the patient, and deliberately administered the poison, but the patient died of some other cause before the poison could take effect.

All of these interpretations and more besides are compatible with the actual content conveyed by the game. There is the temptation to write at least some of them off as ridiculous, as 'unreasonable' interpretations, but to do so is to deny the ultimate authority of events in the virtual world, subordinating them to what is 'reasonable' for a human in our world to interpret from a piece of fiction.

This is a significant problem if the ethics of the virtual are to be a matter of what real ethical categories events in virtual environments belong to. It also raises the spectre of another objection to inflationary views, by attaching to games the same concerns about interpretive authority over the content of fictions that have arisen in relation to other fictive media. Inflationary views require that there is an ontological novelty to the virtual, which sets it apart from all other forms, to prevent objections to video game violence spilling over onto violent classics of literature, drama and cinema.

The gamer's dilemma, at least in Luck's original form, requires that 'virtual acts'—directing a character in a virtual world to do something—relate to the resulting actions in virtual worlds in systematic, consistent ways, such that the relative moral weights of virtual acts are proportional to the relative moral weights of their real counterparts. This, clearly, is not the case for any other fictive medium—whether reading about an act of murder or child sex abuse is bad depends on a great deal more than the wrongness of the act described. I turn now to the question of whether games can be clearly and rigidly separated from all other fictive media.

Agency and identification

The standard way of isolating video games from other media is to claim that gamers are active in a way that audiences for other media forms are not. Games, we are told, are *interactive*, where books and films are passive; gamers employ genuine agency in their play. Outside the literature on the Gamer's Dilemma, this claim has been extensively contested, but in the interests of brevity I shall not draw on these discussions in any depth.

The best illustration of the problems raised by this claim about agency comes from a recent paper by John Tillson (2018). Tillson takes an unusually austere Kantian stance on virtual violence, arguing that all game portrayals of violence show disrespect to other living beings and are therefore immoral. No such principle could reasonably be applied to any other narrative medium, and so it is especially important for Tillson to establish that games are different. He writes:

A representation is a simulation in the sense I mean, iff it is agential in the sense that a player selects actions to be undertaken by the player's character(s) (i.e. it comprises agential affordances). Put another way, simulations comprise agential affordances for at least one player (usually via an avatar). To simulate wrongdoing is to direct one's character to perpetrate some wrong (something that would be wrong if it were done outside of a simulation). (2018, p. 3)

The final sentence here is a close paraphrase of Luck's 'were the environment actual', discussed at the start of the previous section. The only difference is the omission of the explicit reference to the game's 'environment'. Perhaps wary of exactly these ontological commitments, Tillson avoids talk of environments altogether. However, his sentiment still appeals to taxonomical consistency between immoral acts in the real world and—in Tillson's terms—the simulation of immoral acts.

The main questions raised by this passage, though, relate to the first two sentences. Tillson does not attribute to the player of a violent game Luck's act of 'virtual murder'; instead, the player 'simulates wrongdoing' (rather than 'committing simulated wrongdoing') via a relationship to one or more specific entities within the simulation.⁵ I take it that to say a simulation 'comprises agential affordances' is to say that the player makes choices about the content of the simulation; this is supported by Tillson's later remark that:

⁵ There is a question here of how entities within a simulation are to be individuated—what is it, for example, to count as a *character* in a simulation rather than an inanimate object, or some other pattern of software behaviours. I shall return to this in the next section.

Part of simulating an action on my understanding is an interactive interface, so that it is not merely imagined. Instead, there is some external information system that affords actions and is updated with information of one's having acted. (2018, p. 3)

Here it is clearly player actions that matter, and they matter whenever an 'external information system' changes in response to them. There is nothing in this second quote to tell us about which in-game acts the player may be held accountable for. In the previous quote, there is an ambiguity: in the first sentence, a player always acts through one or more characters, but in the second, the player's wrongdoing is only 'usually' done through an avatar. If an avatar or character is essential for wrongdoing, then Tillson has two problems: first, the many interesting cases where there seems to be no 'player character', and second, a question of what makes a certain character the player's. If a player character is not essential, however, the distinction Tillson needs between games and other media is hard to maintain.

Most discussions of computer game violence focus on a very narrow subset of games in which it is possible to commit acts of violence—the games, such as first-person shooters like *DOOM* and third-person mayhem games like *Grand Theft Auto*, which have attracted the most headlines. But there are plenty of alternatives, including one which is familiar to all my students but almost completely absent from the literature, which comes from the massively popular *The Sims* series.

In *The Sims* games, the player does not directly control any characters in the environment. Instead, at the start of a game, the player creates one or more virtual people—'sims'—and is then charged with controlling these sims' environment in order to enable the sims to prosper in school, work, love and many other aspects of life. Sims thrive when their biological needs are met, and die without. A famous⁶ pattern of play in earlier titles (prevented by developers since the series' fourth major instalment) was to wait until a sim had entered a swimming pool and then delete all the ladders allowing them to get out. Sims were unable to get out of the pool without a ladder, and would drown.

There are many other, similar examples. As a child, I would carefully choreograph crashes between my own trains in *Transport Tycoon*. Players of *Theme Park* and *Rollercoaster Tycoon* sometimes deliberately build rollercoasters that exceed safety limits, throwing passengers to their deaths. In *Black and White*, the player takes on the role of a disembodied god, and one way to secure worshippers is

to terrorise them through violence. One assumes that Tillson, at least, would want to condemn such acts, but they are excluded if a player-character or avatar is essential for wrongdoing.

The other problem for Tillson, if player-characters are essential, is that of identifying *which* characters are relevant. Tillson addresses this question almost entirely negatively:

In particular, it does not matter for the distinction whether one identifies with their avatar, or identifies with them any more than they might identify with the lead character of a film. The distinction turns entirely on whether or not one is able to direct a character to undertake actions. It does not matter whether the experience is immersive, affords one a sense of 'presence', or whether one has a sense of embodiment within the game. (2018, p. 3)

These exclusions are necessary because previous attempts to appeal to them have revealed that none of them are universal (most, at very least, exclude the 'god', 'tycoon' and 'sim' genres already discussed above). All that matters is that a player 'directs' a character to take actions—but what constitutes 'directing'? Consider again game A, outlined in the previous section. The player may not know, as they complete the objectives of the game's prologue, that doing so will result in the king murdering the protagonist's father. It seems, though, as if on Tillson's picture the player has nevertheless directed the king to do so. The player has directed the events of the game to proceed, as part of which the king has killed the father.

Let us assume the player has not played the game before; according to Tillson they have just unknowingly, unintentionally committed (virtual) murder. In no other circumstance that I know of can an act be called murder if its agent neither willed in any way nor even knew of the killing that occurred. The player can hardly even be called culpably negligent, unless we are to hold that one should look up whether there are any murders in a game before playing it, and completely abstain if there are.

Tillson's rejections of narrative ('identification') and phenomenal connections between the player and specific individuals in the game's fiction leave him without a way to deny that the player is implicated in the king's action. Even if one could be found, Game A can easily be modified to pose the same question, by making it a game in which the player 'controls' multiple characters at once (as many role-playing games allow).

Say that, in the prologue of Game A, the protagonist and his father go exploring in a nearby cave. They are joined by a mysterious stranger, and for as long as the three are adventuring together, the player has full control over the stranger's actions. At the climax of the prologue, the stranger reveals himself to be the king, betrays the protagonist and murders

⁶ I cannot offer quantitative data to back up this claim to popularity, but the pattern is well-enough recognised that I can use it as an example not only in games-specific teaching but also in general philosophy classes and have most students recognise it immediately.

his father. The game cuts to black at this point and the next scene is the youth's confrontation with the king some years later. Again, Tillson would have to say that the player has committed virtual murder even if the player never again controls the king's actions.

As with Luck, it is only the content of the game's fiction that determines the moral valence of the gamer's play. It doesn't matter how the player experiences or understands what they are doing as they push buttons on their controller; if they cause the fiction to proceed and the king murders the father, they have virtually murdered, or simulated murdering, the father. This is supported by another remark of Tillson's, which implicitly abandons reliance on any specific player-character relationship:

Ordinarily there are a range of action options which the player can select from, but even if there is a choice between selecting just one action and not playing, this could still count as simulating wrongdoing in that one initiates the character's action. (2018, p. 3)

Taking Tillson at his word here collapses the very distinction between games and other media on which his overall case rests. A film or TV show, watched at home on DVD, is an information system, external to its audience and responsive at least to their pressing the 'play' button—and thereby directing the characters to take whatever actions the narrative attributes to them. Plausibly, the same is true of a book, though in this case the information is generally not stored digitally.

Tillson explicitly denies that merely reading a book which describes, or watching a film which depicts, horrific acts—his example is *Schindler's List*—is inherently morally wrong. He says:

Rather, any disrespectfulness evident in reading such descriptions would have to consist not in the fact of reading, but in the attitude that one has to the wrong represented... While one certainly could watch *Schindler's List* in a disrespectful way, perhaps cheering through murders, or jeering at lucky escapes, no such inappropriate emotions need to be had for there to be a *pro tanto* wrong in the act of playing the hypothetical videogame. (2018, p. 10)

But as we have seen, given Tillson's own restrictions, this difference between video games and books or films cannot be clearly maintained. The matter of the player's attitude to in-game events will resurface in relation to deflationary views, but at very least there is the logical possibility of a player playing a violent video game without the presence of such problematic attitudes.

The inflationary attempt to claim that the gamer employs genuine agency, in contrast to the consumer of other media, fails because it is not possible to give a clear account of what

this agency consists in. Any attempt to offer an explicit and precise formal definition inevitably excludes at least some relevant examples, or becomes so general that it fails to exclude the watcher of a movie or reader of a book. Many gamers, of course, report that playing a game *feels* different to watching a film; the last attempt to establish an inflationary metaphysics which I will discuss, Philip Brey's, relies on this feeling.

Representation and realism

Luck calls Brey's account of virtuality 'suitable' (2009, p. 31) if specificity is required, probably recognising that Brey focusses on not just virtual worlds but the kind of virtual worlds sustained by 'virtual reality' technologies (that is, computers that use a headset display rather than a monitor screen). However, the details of Brey's ontological picture, though they reveal further flaws in Luck's account, stand at some distance from it.

Brey's central concern is the question of *realism*. For Brey, what makes virtual reality systems ethically distinctive is that:

There is an experienced similarity between real-world actions and actions in VR. This similarity is greatest in highly-immersive VR systems (1999, p. 8)

Because this realism is cast here explicitly as an *experiential* phenomenon, there is already a sense in which we have moved away from inflationary metaphysics; one can read Brey as more concerned with how the images produced by a game affect the player as experiential subject. However, the case Brey builds from here, across a number of articles, clearly attributes a great deal of work to details of how the game software operates rather than to player psychology, and so I think it is still appropriate to class his view as inflationary.

The similarity Brey appeals to is one part audiovisual—a VR headset allows a game's display to take up more of the player's visual field, and to blindfold the remainder, and will generally come with earphones which allow at least limited surround sound effects—and one part proprioceptive—VR systems often (and Brey anticipates this becoming ubiquitous) have control systems which track the position of player limbs directly rather than relying on buttons or joysticks. There is nothing about these technological differences which intrinsically connects with the *content* of VR experiences; in Brey's definition, a player of a VR game can 'realistically' cast spells or pilot exotic alien spacecraft.

Luck, as we have seen, insists that an event that happens off-screen may be just as morally problematic as one displayed with the highest degree of 'realism'. For Brey, this is not the case; it is only those in-game actions experienced

by the player as 'realistic' which can be morally assessed. Indeed, Brey at times goes further, insisting that realism requires a 'first-person perspective', defined thus:

A first-person perspective, finally, entails that the environment is perceived and interacted with from a single locus. A first-person perspective suggests a degree of immersion in a world, rather than the experience of the world as an object that can be (partially) controlled from the outside. (1999, p. 6)

This is Brey's attempt to define the kind of link between the player and a specific character in the game that Tillson's account stumbled over. Similar problems arise here, especially the problem of overzealous exclusions. The *Grand Theft Auto* series, probably the most commonly-discussed example of video game violence, has never used a first-person perspective, at least as the term 'first-person perspective' is conventionally used. Taking Brey's specification more generally, in terms of 'a single locus', risks trivialising it, since a player can ultimately only take in one display at a time regardless of what kind of images the game puts on it.

The examples from *The Sims* and other god games discussed in the previous section also fall outside this definition. It is possible to grow 'immersed' in playing these games, in the sense that one's attention can be captured so that one loses track of time and so on, but this is not an intrinsic power of any particular technology—a huge range of other media experiences are immersive in this sense.

As for the more general phenomenal similarity between on-screen events in a video game or VR game and 'real' experience, this certainly was not as compelling in 1999 as Brey claims. Even now, only the most expensive VR headsets can display images at a high enough resolution to trick the human eye, and doing so with any degree of fluidity requires enormous computing power. The more detail that is required in a scene, the higher this requirement becomes.

Coupled with Brey's claims about interactivity, this limitation is a problem. He defines this long-contested concept as follows:

Interactivity entails that the represented environment must allow for manipulation and navigation. Manipulation implies the modification of aspects of the environment in a fairly direct way, for example by clicking on them with a mouse or by grabbing them through the use of data gloves. Navigation implies the ability to change the location from which one perceives and manipulates features of the environment, including the ability to perceive objects from different angles. (1999, p. 6)

A few pages on, Brey offhandedly claims that "represented items in VR environments are (usually) interactive" (1999, p. 10). And 4 years later, he says:

At first glance, virtual worlds contain many of the same kinds of entities found in the real world. They may contain spaces, trees, desks, chairs, pencils, dogs, written messages, conversations, money, words, and so forth... [These] virtual objects are not just fictional objects, because they often have rich perceptual features and, more importantly, they are *interactive*: they can be manipulated, they respond to our actions, and they may stand in causal relationships to other entities. (2003, p. 276–7)

Unless we take his definition of interactivity in the thinnest possible sense, these claims cannot be reconciled with the realities of video game environments. In many video games, the vast majority of things in any given scene will not be interactive. A flowerbed or lawn is unlikely to consist of individual leaves or petals, or even individual plants—instead, it will be a flat colour pattern which looks about right until inspected closely. The drawers of a filing cabinet, or the books on a shelf, will be effectively painted on to the flat surface of a solid cuboid. Video games which do not avail themselves of these shortcuts must either look extremely sparse, render all their objects in low detail, or place impossible demands on their hardware.

To put this slightly differently, we could differentiate between the perspective of an 'immersed' player and the game's designer (or anyone else who has reason to pay attention to the structure of the game's software). To an immersed player, the books on a shelf are objects—but they are not interactive except in the minimal sense that the player can move their view along the rows of spines. To the game's designer, the books are not individual objects, but much of the immersion on which Brey places such high importance is lost. So are there books in the game or not?

If this problem seems trivial in terms of books, it is one tip of a dangerous iceberg (another is Luck's claims about 'clarity' discussed above). Its general form is the question of how we determine what kind of action has occurred. What makes it the case that a sequence of game images constitutes a virtual murder?

Murder, roughly, is the deliberate killing of a person. Where we place the 'virtual' in transferring this definition to the virtual makes quite a difference to the act we are talking about. If virtual murder is the deliberate virtual killing of a *person*, then it can only ever occur in virtual worlds where there are multiple people—with present technology, this would be limited to multiplayer games, and thus only tangentially related to the kinds of cases captured by the gamer's dilemma (see Ryland 2019 for discussion on this). To speak of 'killing' a virtual person seems an inflation, since the life that is ended is only a virtual life.

Instead, for absolute clarity, we must speak of the deliberate virtual killing of a virtual person (even the deliberation

may be virtual—most in-game killing has no alternative except not playing). This poses the question of what determines that a certain in-game entity is a (virtual) person, and a certain in-game event a (virtual) killing.

If the victim of the act is a character in the game's narrative, and the process within the game's code that is activated is titled 'death', then it is almost impossible to refer to the event as anything other than a killing. But many in-game acts are not like this; think of the goop enemies in 'Game B', described above in reference to Bartel. Even if the game's code refers to what happens when they burst as 'death', or a state variable associated with an individual enemy changes from 'alive' to 'dead', these creatures' eventual, and routine, reanimation suggests that 'death' is the wrong word.

Another challenge to the personhood of these enemies is their lack of individuality. Even in games where one can 'kill' in-game entities that look like people, such as the *Grand Theft Auto* games, in terms of the game's code the figures are only very basic automata. Video games are largely smoke and mirrors; the pedestrians in a *Grand Theft Auto* game do not have jobs, families, homes or communities, they simply walk from point A to point B and then disappear. They might leap out of the way of an oncoming car, but if hit they do not incur medical bills; if killed they leave no bereaved loved ones.

Nor are these entities ontologically persistent. When a player enters an area, the software will populate it with walking figures. If the player leaves the area, those figures will be removed from the game's state (is this virtual killing? If not, why not?), and if the player returns later, new figures unrelated to the first set will manifest.⁷ Indeed, if a player saves their game, stops playing and turns their hardware off, the game may not save the states of these space-filling figures, instead simply generating a new set when the player next loads.

Even where a game's narrative is clear that a murder has occurred, the fact that most modern games allow the saving of progress casts doubt on the moral status of the act. Imagine a game with an optional side-mission to kill a minor villain; if I save before undertaking this mission, complete it, then reload and finish the game without replaying the mission, have I killed the villain? Killed and then resurrected him? Expunged my killing of him from the record?

What all of this goes to show is that much hangs on Luck and Ellerby's (2013) 'clarity'. Video game events are hard to pin down; establishing for certain that any one virtual

event belongs to a specific moral category is always going to require an interpretive leap. An inflationary metaphysics would resolve this by providing a canonical answer: how things stand in the virtual world picked out by the game. But we could never be sure of this answer, because of how limited the information provided by even a highly-detailed game is about its world.

Further, to make an inflationary metaphysics plausible without extending the gamer's dilemma to other media, some account would be needed of why video games need to be explained in terms of virtual worlds where films and novels do not. Brey's attempt to do this in terms of a formally-constructed 'realism' fails because the results are much too narrow to capture all relevant cases. If the way that the player experiences a game is the source of moral significance for acts the player commits or directs inside the game, then this can only be accounted for in deflationary terms. We must turn from claims about 'virtual worlds' and 'virtual acts' to questions of images, and their effects and meanings, in the real world.

Deflationary views

Deflationary views are at once more heterogenous in their moral commitments than inflationary and more unified in their metaphysics. It will not be necessary to examine these diverse positions in as much depth as the inflationary writers, precisely because their metaphysics are much less a part of the picture. In general, deflationary views take video game software as producing images within our world, not connecting us to independent virtual worlds. The gamer's action, on these views, is simply to consume these images.

This is evident straight away in Bartel's (2011) response to Luck, quoted above. For Bartel, the problem with a game which features child sex abuse, which murder games avoid, is that by definition it produces pornographic images of children (this claim has been extensively critiqued, especially by Stephanie Patridge 2013). Such images are obviously morally problematic regardless of the circumstances in which they are produced.

Bartel's view expands very naturally to other media. A game that features child sex abuse is of a kind with film footage of a child being sexually abused—it generates a kind of image that one should not generate. Of course, lots of other systems can generate such images (even a kaleidoscope could, potentially, generate such an image by chance) but distinctions can perhaps be made in terms of consistency, reliability, or even intent of the designers or consumers; more on this last point later.

Stephanie Patridge takes a very similar stance when she argues that:

⁷ To be clear, this is the case with the pedestrians in *Grand Theft Auto* games, and with many kinds of bystanders in many similar games. It is not the case in all games; in some games, usually ones with less densely-populated environments, such as the *Legend of Zelda* series, non-player characters in the game's narrative may have more persistence.

Some otherwise imaginative representations, including those that we find in videogames, have what I call incorrigible social meanings. These representations operate to limit the range of reasonable interpretations of fictionalized representations. (2010, p. 307)

Here videogames are almost incidental to the point; what matters is representations and their semantic connections to the real world. When we see a human figure on a screen appearing to be, for example, stabbed, what makes this image meaningful is its resemblance to real-world wrongs. Per Patridge, some such connections are stronger than others, because they perpetuate contemporary injustices or draw on long histories of oppression.

For Garry Young and Monica Whitty, the question posed by the gamer's dilemma is

whether taboos are context-dependent, so as to match the contingent nature of the interactions available within a given space—one possible outcome being to render certain virtual spaces taboo-free zones. Or whether our attitudes toward taboos necessarily transcend worlds, because, along with our wider sense of morality, they guide those actions, attitudes, and values that, themselves, transcend worlds. (2011, p. 803)

They draw the opposite conclusion to Patridge, arguing that while the affective dimension of a taboo, the disgust or anger felt at its violation, might transfer over to an in-game act, the actual moral weight does not, because the connection between the player and agents within the game is not intrinsically strong enough to make the agents' actions the player's. Young and Whitty's only remarks about the gamer's actions relate to the psychological attitude from which the gamer engages with the images (for example, a gamer who played a violent game specifically to use it as pornography might rightly be treated with suspicion, but this is framed as an atypical case).

In an extensive subsequent body of work (2013, 2014, 2016), Young has gone on to argue that the gamer's dilemma is not an ethical but a meta-ethical problem. That is, the problem is not about whether players engage in particular virtual acts, but whether players' engagement counts as 'promotion' of certain real-life courses of behaviour. Promoting or advocating an immoral act is wrong, but there is widespread disagreement about what features a representation must have in order to count as promotional and this, according to Young, is where the gamer's dilemma truly lies.

An advantage of Young's position is that it sketches a way for deflationary theorists to avoid the old empirical question of whether playing violent video games causes people to behave violently.⁸ Because they deny that virtual acts really

are, or are morally assessable as, the real-world acts they resemble, deflationary theorists have struggled to explain what is wrong with the images they condemn. This is clearest in virtue-theoretic responses to the dilemma (most obviously McCormick (2001) as well as perhaps Thomas Nys (2010) and Bartel (2011)) which suggest that the problem is the harm players inflict on themselves by playing games which cultivate vices, or unvirtuous behaviour.

Indeed, Young's own argument against the application of real morality to in-game acts is that games do not have any *intrinsic* power to change players' behaviour. With Whitty, he writes:

It is our contention, then, that virtual taboo violation (in the context of single-player games) is not in and of itself psychologically unhealthy. (2011, p. 813)

The advantage of the 'promotion' frame of Young's later work is that promotion of an immoral act can be a bad thing in itself even if ineffective. A hate preacher standing on a street corner and calling for violence against a minority group is doing something wrong even if no-one walks past him or hears him. Patridge's 'incorrigible social meanings' do similar work with even less of a consequentialist inflection.

Two other deflationary theorists merit brief mentions here. Geert Gooskens explicitly labels "acting-in-a-virtual-environment as a form of image-consciousness" (2010, p. 66), providing a treatment of virtual violence grounded in the phenomenological tradition, drawing on Husserl's account of the 'as-if' relation. Gooskens uses this to dispute the application of conventional normative ethics to virtual acts, concluding that such acts are never morally prohibited but may still be uncomfortable "when they are performed by people who have lost track of the distinction between their actual I and their (immoral) image-world-I" (69).

Finally, Rami Ali (2015), as quoted previously, draws attention to the player's own understanding of their in-game acts, to dispute Luck's implicit claim that the game software itself is the unique authority on what takes place in its fiction. To assess players' actions, we must surely at least note what they take themselves to be doing. For Ali, we can only fault the player if they understand their own in-game actions in the way that the designers intended them to, *and* the designers intended players to consider a particular in-game action immoral.

Footnote 8 (continued)

standard point of reference on this issue for participants in the gamer's dilemma debate has been Craig Anderson and team's 2010 meta-analysis and Christopher Ferguson and John Kilburn's response (also 2010).

⁸ I suspect this question, at least formulated as I have here, is too imprecise to answer, but it is entirely tangential to my argument. The

These views are all deflationary; they agree that the problem with a video game that features child sex abuse is to do with the fact that such a game produces images, in the real world, that appear to represent child sex abuse. There is a subtle bifurcation among them, though, about what exactly the problem is, and it leads to a different dilemma.

Some writers, most clearly Gooskens, but also Bartel in more recent (2020) work, take the moral issue to arise when the player engages with the game in some inappropriate psychological state. This might be, as Gooskens has it, the blurring of the lines between player motivation and character motivation, the player taking on the straightforwardly immoral motivations of the character as their own. Or, following Bartel, it might be the player deciding to use *as pornography* images which were not intended as such.

On these accounts, the player is primarily a media consumer, approaching their consumption in the wrong way. The same problem, then, can also go for other media. A consumer who reads a story or watches a documentary about child abuse as a surrogate for watching child pornography commits the same kind of wrong.

On the other side, where Patridge's 'incorrigible social meanings' and Young's 'promotion' sit, the problem is fundamentally with the content of the images themselves, independently of anything the player does. Sebastian Ostritsch (2017) takes a similar position in protesting games whose narratives or scenarios 'endorse' morally problematic worldviews. Ostritsch goes one step further, though, in arguing that the player is not morally implicated in such endorsement unless they agree with it.

These two variations of the deflationary view thus fall away, in different directions, from the question of the *gamer's* action. On the former variation, there is nothing specific to games about the wrong; all media are subject to this problem. On the latter, if anything is specific to games then it is only in the role of the designers, publishers and distributors of games; perhaps the player is complicit in whatever way the purchaser of morally compromised goods is complicit, but this is generally held to be a relatively trivial wrong except in extreme cases.

What makes the gamer's dilemma compelling as a problem is that we sense some difference between what the gamer does and what other figures—the reader, the movie-watcher, the listener to music or radio drama—do on their sofas. On these deflationary views, this difference disappears; even if we follow Patridge and Young and direct our moral scrutiny to the designers and publishers of games, it's not clear that there is anything in what they do to single out games from other media.

What I will outline in the final section below is that either or both of these forms of the deflationary view can be sharpened in focus by treating the gamer as a kind of *performer*. A performer is more active in a performance than media

consumers are generally regarded as being, and we might well ask Bartel- and Gooskens-style questions about the relish a performer might take in playing an immoral role or character. Similarly, unlike a 'mere' consumer, a performer plays an important role in the production of media and may be critiqued as such along Young's or Patridge's lines, according to how the symbols they perform relate to existing cultural tensions and harms.

Does the gamer do anything?

While much-discussed outside the literature on the gamer's dilemma (see e.g. Kania 2018, whose arguments I do not have space to engage with here), the idea that the gamer might be a kind of performer has received little interest within it. The only author to address it that I am aware of is Sebastian Ostritsch, who ultimately argues:

"The moral fault of the actor, I believe, lies in this willing exposure of others to an immoral endorsement. It is at this point that the analogy with playing games breaks down. Since gaming is not inherently public... we cannot judge gamers who play an immoral game in the same way we judge actors who are part of an immoral movie." (2017, p. 126)

This fits Ostritsch's argument that a media act is immoral only if it promotes an immoral viewpoint. But I think it is too hasty. Imagine an actor who, in his spare time, rehearses scenes from D.W. Griffith's infamous racist propaganda film *Birth of a Nation*. We would, I think, be alarmed to learn of such a person; this behaviour at very least is indicative of a lack of respect for the injustices contained in and supported by the material. We need not infer any specific *form* of racist animus on the part of the actor; to spend time engaged in this activity, even if one does so without any intention to promote racist views or political programmes, is racist because the performance that results, regardless of audience, is racist (i.e. something like Patridge's incorrigible social meanings is at work in it).

To make this picture a little sharper, imagine that instead of rehearsing *Birth of a Nation* by choice, the actor is down on his luck and, in desperation, takes a role in a revival production backed by a rich white supremacist. Some scenes are filmed, but the production falls apart and no film is ever released. We are, I think, likely to feel pity for the actor; his circumstances forced him to compromise his character for money. Such pity assumes that his participation, *even though it came to nought*, is morally problematic.

If the actor's circumstances were different—if ambition and the desire to ingratiate himself to a powerful backer motivated him—then we might be more willing to condemn him, and even scorn his failure. We would be troubled if,

despite *asserting* a moral objection to the film, he went on to take professional pride or even personal relish in the role rather than doing the bare minimum to keep the job.

These are, I think, fairly close to the moral contours of the gamer's consumption of a game, and the issues of player psychology raised variously by Gooskens, Bartel and Ali discussed in the previous section. A gamer is seldom if ever forced by circumstance to play a violent game; we assume then that they are more like the actor who practices alone for his own motives. If we see the gamer take pride in violent play, or display enjoyment, our discomfort with them increases.

And there is a metaphysical difference (at least of degree, if not kind) between the actor and the watcher of a movie. An actor is ontologically bound up with even a rehearsal of a scene to a degree that a media consumer is not. Yes, a media consumer is a necessary condition for the scene they consume to be consumed—someone must press 'play', or buy a ticket, or open a book—but that is all; media can be consumed 'in the background' or 'as background noise'. An actor's participation in a scene is a constant stream of small actions, each of which is a constituent of the scene. Decisions the actor makes—how to stand, to speak, to gesture—are essential to the work, not in the sense that what the actor chooses is transformatively significant to the narrative, but in the sense that *some* such decision must be reached for the performance to occur.

This, I think, accounts for the strong and common intuition that the gamer is a kind of agent, and allows us to get the scope of their agency right. The gamer does not any more choose to kill zombies, or aliens, or enemy soldiers than the reader of a book about the same events—generally speaking, violent encounters in games tend to be structured to require these killings just as much as reading a book requires turning a page. But the gamer has more say in *how* the killing happens, in how the act is represented (as a simple example, one gamer may prefer to use a shotgun weapon for the sense of immediate power it provides; another may use a chainsaw instead to generate gratuitous gore⁹). The gamer typically makes many of these little decisions around the restrictions of the 'script' provided by the game.

A performer is at once a limited agent and the first and most immediate consumer of their own performance, as any musician who practices regularly will attest. To perform is simultaneously to be aware of one's performance; in performing one considers potential audience perspectives through one's own. There is not space to develop this point here, but the broader understanding of 'performance' in the philosophy of gender following Judith Butler (1990) would enrich this argument still further.

Whether the performance of the gamer is different in degree or kind from the engagement of a media consumer is another question that will have to wait for further work. It is worth noting in this context Kendall Walton's (1990) argument that the reader of a book is far from passive, but in fact conducts what might be called an internal performance of the text. At very least, performers make a useful point of comparison for the gamer and bear more investigation as such.

If we take the gamer to be (like) a kind of performer, then the gamer's dilemma falls under the broader question of what sorts of performance are morally permissible to engage in. This is a complex question with a long history, including debates about freedom of expression, obscenity, propaganda, the promotion of taboo or immoral views and political programmes and many more issues besides. The commercial structures that exist around any morally problematic performance (who profits from the representation of whom) are as important as the performer's own psychology, intent, and understanding of their performance, as are the perceptions of diverse groups of outside observers.

What makes the difference, for example, between an actor taking a role in a remake of *Birth of a Nation*, and an actor taking the role of Macbeth in Shakespeare's play? Macbeth certainly commits atrocities, but we would regard as prudish in the extreme someone who called for the banning of his play. I do not propose to answer this question here, merely to point out that any answer would be far more complex than those proposed thus far to the gamer's dilemma—and that we should expect any actual satisfactory answer to the gamer's dilemma to be similarly complex.

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

- Ali, R. (2015). A new solution to the gamer's dilemma. *Ethics and Information Technology*, 17, 267–274.
- Anderson, C. A., Shibuya, A., Ihori, N., Swing, E. L., Bushman, B. J., Sakamoto, A., et al. (2010). Violent video game effects on aggression, empathy, and prosocial behavior in eastern and western countries: A meta-analytic review. *Psychological Bulletin*, 136(2), 151–173.
- Bartel, C. (2011). Resolving the gamer's dilemma. *Ethics and Information Technology*, 14, 11–16.

⁹ I am indebted to Nele van der Mosselaer for this example.

- Bartel, C. (2020). *Video games, violence and the ethics of fantasy: Killing time*. London: Bloomsbury.
- Brey, P. (1999). The ethics of representation and action in virtual reality. *Ethics and Information Technology, 1*, 5–14.
- Brey, P. (2003). The social ontology of virtual environments. *American Journal of Economics and Sociology, 62*(1), 269–282.
- Butler, J. (1990). *Gender trouble: Feminism and the subversion of identity*. London: Routledge.
- Chalmers, D. (2017). The virtual and the real. *Disputatio, 9*(46), 309–352.
- Ferguson, C. J., & Kilburn, J. (2010). Much ado about nothing: The misestimation and overinterpretation of violent video game effects in eastern and western nations: Comment on Anderson et al. (2010). *Psychological Bulletin, 136*(2), 174–178.
- Gooskens, G. (2010). The ethical status of virtual actions. *Ethical Perspectives, 17*(1), 59–78.
- Kania, A. (2018). Why gamers are not performers. *The Journal of Aesthetics and Art Criticism, 76*(2), 187–199.
- Lewis, D. (1986). *On the plurality of worlds*. Hoboken: Wiley-Blackwell.
- Luck, M. (2009). The gamer's dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology, 11*, 31–36.
- Luck, M. (2018). Has Ali dissolved the gamer's dilemma? *Ethics and Information Technology, 20*, 157–162.
- Luck, M., & Ellerby, N. (2013). Has Bartel resolved the gamer's dilemma? *Ethics and Information Technology, 15*, 229–233.
- McCormick, M. (2001). Is it wrong to play violent video games? *Ethics and Information Technology, 3*, 277–287.
- Nys, T. (2010). Virtual ethics. *Ethical Perspectives, 17*(1), 79–93.
- Ostritsch, S. (2017). The amoralist challenge to gaming and the gamer's moral obligation. *Ethics and Information Technology, 19*, 117–128.
- Patridge, S. (2010). The incorrigible social meaning of video game imagery. *Ethics and Information Technology, 13*, 303–312.
- Patridge, S. (2013). Pornography, ethics, and video games. *Ethics and Information Technology, 15*, 25–34.
- Ryland, H. (2019). Getting away with murder: Why virtual murder in MMORPGs can be wrong on Kantian grounds. *Ethics and Information Technology, 21*, 105–115.
- Seddon, R. F. J. (2012). Getting 'virtual' wrongs right. *Ethics and Information Technology, 15*, 1–11.
- Tillson, J. (2018). Is it distinctively wrong to simulate doing wrong? *Ethics and Information Technology, 20*, 205–217.
- Walton, K. (1990). *Mimesis as make-believe: on the foundations of the representational arts*. Cambridge: Harvard University Press.
- Young, G. (2013). Enacting taboos as a means to an end; but what end? On the morality of motivations for child murder and paedophilia within gamespace. *Ethics and Information Technology, 15*, 13–23.
- Young, G. (2014). A meta-ethical approach to single-player gamespace: Introducing constructive ecumenical expressivism as a means of explaining why moral consensus is not forthcoming. *Ethics and Information Technology, 16*, 91–102.
- Young, G. (2016). Violent video games and morality: A meta-ethical approach. *Ethics and Information Technology, 17*, 311–321.
- Young, G., & Whitty, M. (2011). Should gamespace be a taboo-free zone? Moral and psychological implications for single-player video games. *Theory & Psychology, 21*(6), 802–820.

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