

Double-Standard Moralism: Why We Can Be More Permissive Within Our Imagination

Mattia Cecchinato

Abstract

Although the fictional domain exhibits a prima facie freedom from real-world moral constraints, certain fictive imaginings seem to deserve moral criticism. Capturing both intuitions, this paper argues for double-standard moralism, the view that fictive imaginings are subject to different moral standards than their real-world counterparts. I show how no account has, thus far, offered compelling reasons to warrant the moral appropriateness of this discrepancy. I maintain that the normative discontinuity between fiction and the actual world is moderate, as opposed to one that leaves fictive engagements wholly exempt from moral evaluations. I propose a way of addressing the gamer's dilemma that is compatible with a moderate kind of discontinuity. Finally, I contend that the audience is justified in adopting deviant moral attitudes in fictional situations because their consequences largely differ from those that analogous real-world events would have.

1. Introduction

Fictive imagining—the imagining of states of affairs without a commitment to truth—seems to enjoy some measure of freedom from typical moral constraints, such that moral responses towards fiction and reality often diverge. For example, while the real-world threat of a nuclear apocalypse terrifies us, fictional representations of it may amuse. Stanley Kubrick's *Dr. Strangelove* (1964) was even ranked third by the *American Film Institute* in its list of the funniest American movies. Likewise, the scale of violence tolerated in fictional situations vastly exceeds our real-life standards, such as in *Grand Theft Auto* (1997–2013), wherein players' fictional wrongdoings are quickly forgotten. But why should we apply different moral standards across fiction and reality? Are we justified in being more permissive within our imagination?

The philosophical literature has focused on whether *any* moral standards apply to objects of imagination whatsoever, rather than on whether *special* or *different* moral standards should be applied to such objects. Defenders of *response moralism* (Gaut, 1998, 2007; Hazlett, 2009; Bartel and Cremaldi, 2018) argue that the audience's fictive imaginings and their reactions to them can be morally evaluated. The question here is not whether it is immoral for the creator to produce violent fiction but whether it can be immoral for the

consumer to participate in those fictional acts of violence. While some response moralists hold certain fictive engagements to be intrinsically wrong (e.g. Smuts, 2013), others hold fictional activities to be wrong only if they prompt the real-life embrace of a blameworthy belief or attitude (e.g. Cooke, 2014).

Jonathan Gilmore's *Apt Imaginings* (2020) further enriches the facets of the debate by distinguishing between two theses: *normative continuity* and *normative discontinuity*. Normative continuity is the thesis that the norms applying to our attitudes towards fictional representations are the same as those of real life. Normative discontinuity is the thesis that such norms vary across fictive engagements and the actual world. Gilmore also suggests that advocates of response moralism seem tacitly committed to continuity:

If one holds that we can learn of a person's defective character through the pleasure she takes in imagining ... torture..., one is likely implicitly committed to the continuity thesis. One is likely committed, that is, to the proposition that a given attitude toward a fictional character is morally permissible if and only if that attitude is morally permissible when held toward an actually existing counterpart of that character. (2020, p. 105)

In contrast, this paper argues that response moralism should be combined with discontinuity. Call this view *double-standard moralism*:

The moral standards of fictional events differ greatly from isomorphic real-world events, although both domains are subject to moral evaluation. Accordingly, the agent is justified in forming different moral responses.

I shall understand a fictional event as isomorphic to a real-world event if it is in relevant respects exactly alike the real counterpart *but for* the fact that (i) one is real and the other is fictional and/or (ii) the agent justifiably believes that one is real and the other is fictional.¹

Double-standard moralism, as a version of response moralism, says that responses to fictional contents are subject to moral criteria. Yet, given its commitment to discontinuity, these criteria are relevantly different from those applying to the actual world. Consequently, on this view, moral responses to fictional contents should typically differ from moral responses to isomorphic real contents. There is also a sense in which my view goes beyond response moralism; not only do I argue that double standards mark the moral evaluations of *responses* (or *attitudes*) towards fictional situations, but I also claim that they mark the moral permissibility of fictional *actions* and the general moral status of fictional *states of affairs*. While most of my examples include attitudes, responses, and actions directed towards fictions, Section 7 weighs fictional states of affairs against their real-world counterparts to show how broad the territory of normative discontinuity is.

I find double-standard moralism appealing for two reasons. Firstly, it explains away the intuition that fictive imaginings are immune to moral criticism. For if the moral

1 This disjunction accommodates, on the one hand, events that are fictional regardless of what the agent believes and, on the other hand, cases of radical deception where the agent justifiably believes they are engaging with fiction, but they are actually deceived.

standards that govern imagination are more permissive, then we may mistakenly infer that we are not bound by any moral standard at all. The differing moral norms of imaginative engagements might be confused with the complete absence of moral norms. Secondly, double-standard moralism represents a more robust version of response moralism. This is because any view that does not acknowledge a moral gap between fiction and reality has unacceptable consequences, or so I shall argue. By applying double standards, my view takes discontinuity seriously. (That is not to say that response moralists are normally committed to a single standard view; most of them may already implicitly assume discontinuous moral standards. But little work has been done to motivate and justify this background assumption.)

To address my primary concern—why fiction differs morally from life—Section 2 begins by illustrating how fictional contexts allow for an extension of the permissible. Section 3 qualifies the extent of these double standards by introducing a distinction between ‘strong’ and ‘moderate’ normative discontinuity. Section 4 shows how the existing accounts fail to justify moderate normative discontinuity. Section 5 maintains that this discontinuity needs justification, as the alternative view (the continuity view) has unacceptable implications. Section 6 provides such justification. It argues that the outcomes of any given event at least partly determine its moral significance, and so an asymmetry of outcomes grounds discontinuity between fiction and life. Section 7 concludes by discussing whether this asymmetry must be understood in a belief-relative or fact-relative way.

2. Moral Discrepancies

Is it ever morally permissible to commit an act in a fictional situation that would be impermissible in reality? It seems to be a mark of our engagements with fiction that we can respond to its contents in ways that diverge from how we respond to isomorphic real-life counterparts. Gamers, for example, believe striving for a virtual kill record to be appropriate and such an achievement to be good, while a real-world analogue would appal them. What is vice in life often becomes virtue in violent video games. Consumers also find it permissible to support and sympathize with fictional characters such as Tony Soprano—a violent, racist, murderous mobster—when this same audience would only feel contempt for a real-life version of Tony.² Call this moral discrepancy *the extension of the permissible*: in the transition from the actual world to the fictional domain, certain morally objectionable attitudes become acceptable, even desirable.

An extension of the permissible does not seem to arise only with respect to *negative* duties (not to harm others) but also *positive* duties (to aid others). For instance, while we normally have an obligation to rescue people—at least if we can aid them without incurring excessive personal costs or the risk of such costs—in interactive fictional contexts we lack a corresponding obligation. If you repeatedly fail a mission in a video game

2 See Noël Carroll (2004) for an analysis of our deviant moral responses towards *The Sopranos* (1999–2007).

due to time lost rescuing fictional characters, one may complain that you are playing inappropriately. Fictional worlds are not as morally demanding as the real world.

However, the extended permissibility we intuitively associate with virtual murder abruptly stops short of rape fantasies or virtual paedophilia.³ There seems to be something morally objectionable about a person who routinely fantasizes about cruel paedophilic acts, even though their acts of violence involve purely imaginary victims. A *prima facie* extension of the permissible occurs, but not an unbounded one. The appropriateness of moral constraints on imaginative engagements may be further supported by ‘mixed cases’—fictive imaginings based on real-world persons and events. Most of us think, for instance, that viewers of Steven Spielberg’s *Schindler’s List* (1993) ought not to feel approval for the Nazis’ acts of cruelty. Our reactions towards historically based works of fiction tend to track our reactions towards analogous documentaries. But these moral responses might well be prompted by the non-fictional component of the representation, and so we shouldn’t rely too heavily on these examples when considering the morality of fictive imagining.

Regardless of mixed cases, the norms of imaginative engagements still seem to partially soften ordinary moral obligations by (i) permitting some kinds of wrongdoings while (ii) forbidding other kinds of wrongdoings. The combination of these two claims implies that there is an ethics of imagination, but it diverges from real-world ethics. Surprisingly, this plausible view has received no explicit defence in the literature. In what follows, I develop it.

3. Continuity and Discontinuity

The question I intend to address is whether we are justified in applying double standards across fiction and reality. The fact that we act *as if* the realm of imagination is more permissive does not answer the question. We may be wrong. For instance, Jesus in the *Sermon on the Mount* alleges that ‘everyone who looks at a woman with lustful intent has already committed adultery with her in his heart’ (Matt. 5:28), suggesting that the ethics of imaginative responses must be as austere as real-world ethics. Gilmore calls this thesis *normative continuity*, which contends that:

When we respond emotionally, or with desires or *moral evaluations* to what we only imagine to be true, we are subject to the same relevant norms as those governing the fit of our beliefs, emotions, desires, and evaluations concerning the real world. (Gilmore, 2020, pp. 9–10, my emphasis)

Most of us contravene this moral picture in our everyday life (Gilmore himself denies normative continuity), and so the implications of continuity would be wide-reaching, if true. A large part of the world’s population would be committing countless wrongdoings by treating fiction as a morally peculiar domain.

3 These conflicting intuitions give rise to the gamer’s dilemma (Luck, 2009), which is discussed in Section 6.1.

Double-standard moralism aims to vindicate the widespread intuition that the standards of imagining enjoy a special moral permissibility. As Gilmore argues, we have grounds to embrace *normative discontinuity*. But there are differences between how Gilmore conceives of discontinuity in general and how I conceive of it in relation to moral matters in particular. While Gilmore's definition of 'discontinuity' is simply characterized as the denial of the continuity view (2020, p. 7), I find it helpful to distinguish between two kinds of normative discontinuity: strong and moderate. *Strong* normative discontinuity is the thesis that the moral gap in the transition from the real to the fictional domain is such that, while there are moral norms in real life, there are none in fiction. This single standard view is the denial of the appropriateness of morally evaluating the audience's responses to objects of imagination. In contrast, *moderate* normative discontinuity is the thesis that fictional events are subject to moral norms, but they differ from those governing the actual world.⁴ Double-standard moralism is a form of response moralism committed to moderate normative discontinuity, resulting in the claim that two isomorphic events, one of which is fictional and the other is real, can require two different moral responses.

For the sake of simplicity, when referring to 'discontinuity' in this paper I mean its moderate form. I will propose some ways to morally assess fictive imaginings, but I will not fully address what makes imaginative activities susceptible to robust moral criticism as this would require another paper. For example, I will not discuss whether the ethics of imagination includes imaginings beyond our control, such as dreams or involuntary responses. I shall assume that at least voluntary responses to fictional contents can be right or wrong in order to focus on justifying a shift in moral norms across fiction and reality.

4. Existing Approaches to Discontinuity

In my own experience, discontinuity strikes most people as intuitively appealing. Yet there has been little development on its justification. A number of authors provide explanations of our moral deviances in fictional contexts, offering interesting insights into our moral psychology, but failing to justify a variation in moral norms. This section briefly examines how these accounts do not address the central normative question—namely, whether there *should* be double standards.

Shaun Nichols (2006, p. 464) calls these asymmetries in our evaluative attitudes 'discrepant affects'.⁵ Consider *Dr. Strangelove*: although at the end of the film mankind is about to be destroyed by an atomic bomb, this is an amusing fictional situation. Such a response radically clashes with how we would react if we believed this situation to be actual. Nichols explains discrepant affects by appealing to the existence of different desires about what comes about in the real world and what comes about in imaginary situations. That

4 Gilmore may well have this characterization of discontinuity in mind. But his definition of 'discontinuity' leaves open whether discontinuity theorists are committed to strong or moderate discontinuity—a crucial distinction to draw.

5 With this term, Nichols refers indiscriminately to every evaluative attitude that diverges across fiction and alike scenarios in real life, including deviant moral attitudes.

is, while we normally have a consuming desire for the preservation of humanity, in black comedy we lack the same craving for the survival of a fictional population of human beings (Nichols, 2006, pp. 469–472).

Yet, as Gilmore hints (2020, pp. 181–182), even though moral discrepancies might be explained by a misalignment of conative states, it remains an open question whether this divergence is normatively appropriate. Should we have different desires across fiction and reality? Why shouldn't we desire the salvation of humanity in fictional representation as much as in real life? Nichols' account provides no answer to these questions because it merely describes the roles played by polarized desires in ordinary evaluations without providing criteria on whether there *ought to be* such a variation in conative attitudes.⁶

Gilmore properly distinguishes *normative* discontinuity from *descriptive* discontinuity (2020, p. 5). On a descriptive level, discontinuity consists in our moral, conative, emotional, and epistemic attitudinal engagements being factually different depending on whether they involve real objects or objects of imagination. For example, the psychological mechanisms of descriptively discontinuous attitudes may vary as the source of their stimuli varies. The normative question is instead whether the appropriateness of moral, conative, emotional, and epistemic attitudes towards imaginative entities is susceptible to norms distinct from those applying to attitudes towards real entities.

With reference to our *moral* attitudinal engagements, Gilmore seems to focus on a descriptive question about the normatively deviant attitudes we take towards fictional situations. He contends that even though the norm-discordant evaluations of fictional characters are usually kept quarantined from our own, this quarantine is regularly breached in our engagements with fiction (2020, pp. 186–194). When this psychological mechanism occurs, consumers set aside their own evaluative attitudes and adopt those of the malevolent character they are simulating. So, for Gilmore, simulating Tony Soprano's immoral attitude causes the audience to genuinely adopt Tony's point of view. Our divergent moral responses are thus explained by a breakdown in the quarantine that divides how we ordinarily evaluate states of affairs and how we simulate evaluating states of affairs from the standpoint of a fictional character.

However, from a purely descriptive explanation of how moral judgements can diverge across fiction and isomorphic real-life scenarios, a normative justification for this divergence does not follow. Even granting the correctness of Gilmore's view, one can still ask if consumers should simulate a villain's perspective to the point of temporarily losing their own sense of morality. Why should we let our quarantine be breached? Perhaps the acquisition of Tony Soprano's deviant evaluative attitudes as our own is wrong in itself, and we should refuse to watch *The Sopranos*. Gilmore describes a distinctive psychological phenomenon that transpires in the audience's attitudinal engagements with fictional works, but this descriptive discontinuity does not *per se* entail any moral difference in the norms governing our responses to fictions. So, despite Gilmore's other compelling rationales for conative, emotional, and epistemic normative discontinuities (which are not

6 Facing normative questions is, however, beyond the scope of Nichols' article, which is focused on addressing problematic cases for his account of the imagination.

discussed here), we are still in need of an appropriate justification when it comes to moral discrepancies.

Another view is that an audience's moral reactions are apt if, and only if, they meet the intended response of the fictional work. In *Imagining and Knowing*, Gregory Currie puts forwards this proposal for establishing the aptness of emotional responses to fictive contents (2020, pp. 67–71). For Currie, there is a broad sense in which the normatively correct way to employ a given artefact is by its intended function. If a consumer uses laptops as pillows, or family pictures as rugs, she will not get the most out of them, for these artefacts are not designed to fulfil such purposes. By parallel reasoning, Currie argues, when a fictional representation is crafted in order to prompt a certain emotion, eliciting this emotion is the proper way to respond.

The intended response might fix not only the aptness of *rationaly* justified emotional responses but also the aptness of *morally* justified emotional responses. On such a view, deviant moral reactions are appropriate in fictional situations because of their manner of representation. In most games, killing is acceptable as they are carefully crafted to provoke this imaginative response. We express sympathy for Tony Soprano in virtue of the design of *The Sopranos*. So, a given moral response is apt if it is a response to a fictional work constructed to elicit such a response. Since the intended response of a fictional representation often differs from the appropriate reaction to an isomorphic real-life situation, double standards follow.

Although this view captures a normative dimension, it suffers from another problem. Not all artefacts *should* be used as they are designed to be used—certainly not from a moral point of view. Consider a torture machine. It would be absurd to believe that, since this machine has been designed to torture people, it is morally appropriate to perform this task. Let's also consider a case of interactive fiction. *Rape Day* (2019) is an unreleased video game where the rape of women is incentivized to progress the plot. The representational aspects of *Rape Day* are crafted to provoke a specific response, which is not supposed to be moral disgust—the developer states that the game is targeted at those members of the population who are sociopaths. Is the intended response to *Rape Day* also the morally appropriate one? Surely not. The appropriate response seems to be strong condemnation. The mode of presentation of *Rape Day* induces support for fictional rapes, but this is not plausibly the morally appropriate attitudinal engagement.

Morally speaking, even if a fictional representation is meant to produce a particular imaginative response in the audience, its design is not sufficient to make the response right. Consumers can recognize the intended response of a fictional work and correctly judge it to be normatively defective. It seems fair to conclude that the intended response is not what fixes the aptness of our morally evaluative attitudes towards fiction.

5. Rejecting Continuity

The previous accounts might have failed to justify a variation in the moral norms governing fiction and real life simply due to the absence of such variation. Following this line of thought, the continuity view contends that the criteria for the appropriateness of our moral responses hold invariantly across real and imagined domains. To my knowledge,

no theorist fully endorses continuity, and in this section I show why this view must be rejected. But my aim isn't just to show why we should reject an uncommonly held view. It is also to establish how we can explain the sources of some *apparent* moral continuities without saying that the moral standards of fictive imagining hold continuously with reality.

Initially, let us note that the thesis that fictive imaginings can be morally evaluated does not imply that they ought to be evaluated by the *same* metric of their real counterparts. To say that imagination is included in the moral domain is one thing; to say that there is no relevant difference between the moral norms of imaginative and real-life activities is another. Interestingly, Gilmore somewhat disagrees on this point. He maintains that defenders of response moralism are *likely* to be implicitly committed to continuity, while I believe response moralism to be compatible with discontinuity.

One way to see advocates of response moralism as committed to continuity is attributable to their moral disapproval of isomorphic imagined and real-life scenarios. Consider [Berys Gaut \(1998, p. 186\)](#), who makes the case of a person who is incessantly absorbed in rape fantasies but who would never conduct them themselves. Gaut holds this person still morally blameworthy for their imaginative activities.⁷ If we find compelling the rape fantasist example, then (at least in some cases) our negative evaluative responses to imaginings mirror our negative evaluative responses to the actual world in a way that seems continuous. We may explain this close parallel by appealing to common underlying norms, and so obtain an argument for continuity:

- (P1) Certain fictive imaginings are immoral [from response moralism].
- (P2) The real counterparts of these fictive imaginings are also immoral.
- (P3) If certain fictive imaginings and their real counterparts are immoral, then these fictive imaginings are immoral in virtue of the same norms that make their real counterparts immoral.

Therefore,

- (C1) Certain fictive imaginings are immoral in virtue of the same norms that make their real counterparts immoral.

Consequently, normative discontinuity would be false because at least some fictive imaginings are subject to the same norms applying to isomorphic real-world scenarios. But this argument is not sound—(P3) is false. From the negative evaluative symmetry that a given act is objectionable in both fiction and life, one cannot conclude that it is objectionable for the *same* reasons. Virtual paedophilia, for example, may be wrong as it damages one's character or could lead one to actual acts of paedophilia. But real-world paedophilia is not wrong only for corrupting a paedophile's character or feeding future wrongdoings. These reasons do not address the crucial point: real-world paedophilia is especially morally bad because it also directly harms children. The lesson we learn from this example

7 In a personal communication, Gaut stated that he does not commit himself to normative continuity. All Gaut is committed to holding is that imaginings are subject to moral evaluation. In his example, the rape fantasist is doing something wrong, but it is not wrong for the same reasons, or to the same extent, of actually committing rape.

is that different moral reasons can make a given act immoral across fiction and the actual world. Normative continuity is merely apparent.

Still, if we more closely examine our moral evaluations of fictional characters, we may be struck by how many similarities such evaluations bear with those of corresponding actual people. We can react positively to Tony Soprano's wrongdoings, but we never think that his gratuitous acts of violence are morally permissible. We can have a pro-attitude towards Tony, but we never think he is a good person. As Kendall Walton says:

When it comes to moral matters . . . I judge characters by the moral standards I myself use in real life. I condemn characters who abandon their children or engage in genocide, and I don't change my mind if the author . . . considered genocide or abandoning one's children morally acceptable. (1994, p. 37)

The design features of a representation cannot make repugnant principles right, no matter how insistently a fictional character or a fictional society endorse them. They remain repugnant. Because of the *imaginative resistance* we experience when entertaining morally deviant fictional worlds, Walton concludes that he is 'sceptical about whether fictional worlds can ever differ morally from the real world' (1994, p. 37). Is there a continuity in our moral beliefs in the sense that certain types of actions, principles, or character traits are invariantly wrong, whether fictional or real? If so, then why do we believe it permissible to murder in video games, brutally and violently, without remorse?

Here we need to distinguish between internal and external normative perspectives on a given fictive imagining. We can ask: If Tony was real and I was in his world, how should I respond to Tony's actions? This immersive approach is often adopted by the consumer to fully *appreciate* works of fiction, but presumably it is not the appropriate perspective we should take to *assess* the consumer's moral reactions to them. The set of conditions is different, and, taking the external perspective, our question becomes: Since Tony is not real and I am not really in his world, how should I respond to Tony's actions? Walton likely has the internal perspective in mind: the perspective of someone who treats the fiction as their own reality. It is from this viewpoint that our moral evaluations of real people continuously match our moral evaluations of fictional characters. But the discontinuity thesis is about the responses of the audience as such—it is about how *real* people should evaluate what are taken to be merely fictional events. So, it is possible to locate the sources and grounds of our (apparently) continuous moral beliefs in a way that is compatible with normative discontinuity as long as we understand discontinuity as concerning the moral appropriateness of the consumer's reception of a work of fiction. It is permissible from our special position—external viewers of a fictional world—to support Tony, but it is not permissible for Tony to perform wrongdoings in his own world. Because, for Tony, his world is real.

We should not extend continuity to include the audience's responses. This would have unacceptable implications. Let's consider Orson Scott Card's novel *Ender's Game* (1985).⁸ Ender is a child undergoing military training, where he plays a space-battle simulator to

8 I am grateful to Berys Gaut for suggesting this example to me.

develop his strategic skills. As Ender is a talented gamer, he wins every virtual battle. But after defeating all the enemies, Ender shockingly discovers that he was not merely enjoying a video game: his battles were real. Ender was leading real armed forces against living targets. Aware of the massacres he has committed, Ender's moral attitude is totally transformed. He is overwhelmed with guilt. Now, we can raise the question: Is Ender's change in moral attitudes justified? Consider the following claims:

- (1) Ender believed his fictional killings to be morally acceptable.
- (2) The moral standards of fictional killings are equivalent to those of real-world killings in all relevant respects [from normative continuity].
- (3) However, Ender believes his real-world killings to be morally unacceptable.

If (2) is true, then Ender's change in moral attitudes is normatively inappropriate. If the moral standards of fictional killings are continuous with those of real-world killings in all relevant respects, then Ender's moral response should not be altered in any way by the revelation that his killings were real. The discovery that Ender was killing real individuals must be regarded as morally irrelevant: on any moral standard, all that matters was already present in Ender's fictional killings.

The above reading of Ender's case strikes me as implausible. Ender's moral attitudes *ought to* change when he discovers his game to be real. How could such a discovery not matter? Yet, if we assign different moral significance to Ender's real killings, then there must be normative discontinuity. Hence (2) is false since there is some relevant difference between the moral norms of fictional killing and real-world killing. But where does this difference lie?

Put simply, my view is that fictional killing is not continuous with real-world killing in all moral standards because an act's moral significance at least partly depends on the act's consequences, and fictional killing does not have the same consequences as real killing. Many people plausibly hold a view like this,⁹ but so far no one has worked out its details or elucidated its central importance to a comprehensive analysis of our deviant moral responses towards fiction. It is this asymmetry of outcomes that makes Ender's change in moral attitudes appropriate. By learning that his virtual battles were real, Ender realized that his (apparently) fictional killings had outcomes of which he was unaware, and thus different moral significance.

An asymmetry of consequences nicely explains why an extension of the permissible occurs in the fictional domain, or, to put it differently, why the ethics of imagination is more permissive than real-world ethics. It is permissible to root for Tony Soprano, while wishing failure for his real-world analogue, because a fictional character's capacity for causing harm is incomparably smaller. Similarly, it is easier to justify killing in video games as opposed to real life, for fictional murder does not involve genuine death.

⁹ For example, Susan Feagin writes: 'The freedom of imagining is freedom without responsibility ... Pleasure in what one imagines can be as fickle or base as one likes, *without consequence*' (1984, p. 50, emphasis added).

We can further specify whether discontinuity is a result of different consequences, or of different *beliefs* about the consequences of our own acts. In everyday fictive engagements, we adopt deviant evaluative attitudes based also on our best judgement of the likely effects. Children play by simulating wars and shootings for they believe that no one is actually harmed. Viewers laugh at *Dr. Strangelove* because they are confident that the nuclear threat to humanity will not affect their lives outside of the movie. The high credence in the substantially different outcomes of imaginative activities represents a concrete psychological motivation driving the audience's moral deviances. I will examine further the relevance of beliefs to discontinuity in Section 7.

6. Arguing for Discontinuity

To support the hypothesis that normative discontinuity is in large part grounded in an asymmetry of consequences, consider the following imaginary case:

*The Voodoo Video Game.*¹⁰ Through futuristic technologies, an evil scientist creates a multitude of people and a planet for them to inhabit. She then manipulates their lives by mirroring every event that occurs in a popular violent video game. The consequences of the player's fictional acts are now the same as if their acts were real: whenever someone dies in the game, someone actually dies. Since the scientist has given a world-wide presentation of her planet, gamers have now the personal moral responsibility of deciding whether, and how, to engage with this voodoo-like game.

The special feature of *The Voodoo Video Game* is that it presents a *symmetry* of outcomes between fictional and real-world events. Even though the interactive video game remains fictional, *real* people are causally connected to it, so that for every act one performs in the game, there is a corresponding real consequence in the scientist's planet (because of her meticulous mirroring-technology). With this in mind, how should we engage with the game? Most would agree that we should neither kill nor torture its fictional characters but despise those who continue to do so. The extension of the permissible that characterizes ordinary fictional worlds has retracted. Our moral obligations persist as demanding as ever—if by playing this game we can save the lives of some fictional characters with little effort or risk, then we ought to. Ultimately, normative continuity is obtained: fiction meeting the same moral standards as real life.¹¹

Note that all other stylistic and formal aspects of *The Voodoo Video Game* are left untouched by the scientist's project. The original design may still prompt the audience to perform violence upon fictional characters. And just as any other work of fiction, *The Voodoo Video Game* is essentially populated by fictional objects. Far-reaching *indirect* consequences are the only relevant connection between the player and real people, unlike *Ender's Game*, where real people are *directly* controlled. Because *The Voodoo Video*

10 This case's title is merely an analogy to capture the game's special feature, with no reference to religion.

11 One may deny that this is a case of continuity because the scientist can be understood as an intervening agent. In principle, however, we could amend the case by replacing the scientist with a non-agent technology.

Game is ordinary fiction in all respects except for the difference in outcomes, this special feature must explain why the moral norms of a work of fiction align with those governing reality.

This imaginary case teaches us something about normative discontinuity in general. If non-standard fictional worlds become morally continuous with the actual world when they bring about the same consequences, then it seems reasonable to infer that the gap in consequences is what motivates discontinuity in the case of standard fictional worlds. While for *The Voodoo Video Game* the *symmetry* of consequences grounds normative continuity, in ordinary fiction the *asymmetry* of consequences is responsible for normative discontinuity.

Let us summarize these considerations into an outcome-based argument for discontinuity. Consequences matter: given the asymmetrical consequences across fiction and life, what the audience fictively imagines is not as morally significant as the real counterpart. Thus, the audience ought to treat the two domains differently. To put the argument in premise-conclusion form:

(P1) The real consequences of ordinary fictional events largely differ from those that isomorphic real-world events would have.

(P2) If two isomorphic events have largely different real consequences, then they have different moral significance.

Therefore,

(C1) Ordinary fictional events have different moral significance from that of isomorphic real-world events.

(P3) If two kinds of events have different moral significance, then we are justified in having double standards.

Therefore,

(C2) We are justified in having double standards between ordinary fictional events and isomorphic real-world events.

Premise (P1) is the claim that fiction does not have the *same* consequences that analogous real-world events would have—it is not the claim that engaging with fictional representations *never* has morally significant outcomes. For example, psychologists (Rosenberg, Baughman, and Bailenson, 2013; Yoon and Vargas, 2014) provided evidence that once individuals participate in virtual contexts as a superhero, they are likely to become more altruistic and prosocial in the actual world, whereas when they play the role of a villain, the opposite can occur. Yet these outcomes largely differ from those of corresponding real-life violence: no direct harm is caused, no lives are lost, and a real-world killer has a far more tarnished character than any ordinary gamer does.

Premise (P2) is a minimal commitment to the importance of consequences, according to which a large difference between the outcomes of two isomorphic events is sufficient to make a difference in moral significance. This premise does not imply consequentialism—that is, the view that *only* the outcome matters for moral purposes. (P2) is compatible with most non-consequentialist moral theories, since it is standard practice for these views to grant that consequences are morally relevant at least to

some extent.¹² For as Rawls says: ‘All ethical doctrines worth our attention take consequences into account ... one which did not would simply be irrational, crazy’ (1971, p. 30). The first conclusion (C1) holds that fictional events carry a reduced moral significance. Premise (P3) states that when an event’s moral significance varies, our moral responses should also vary. The second conclusion (C2) contends that double moral standards between fiction and life are thus appropriate.

The aim of my outcome-based argument is to identify a distinctive asymmetry in the moral properties of fictive engagements in order to provide a reason for regarding fictional events as morally *sui generis*. It doesn’t say why entertaining certain actions in our imagination is wrong; it says why it is *more* wrong to commit analogous actions in real life. But it should be noted that the very same outcome-based asymmetry may also support an alternative view. The fact that the morally relevant consequences of fictional events largely differ from those of analogous real-world events might provide a basis for a single standard view, one on which fiction is a totally moral-free domain—*strong* normative discontinuity. Double-standard moralism, in contrast, aims to justify *divergent* moral responses to fictions (not their abandonment), and so we need to ask what it is about certain fantasies that makes them wrong even though they are fictional.

6.1 *The Gamer’s Dilemma*

Why should we believe that moral standards can apply to the fictional domain at all? We do have some prima facie evidence. Cases such as *Rape Day* show puzzling limits to the extended permissiveness we enjoy in fictional contexts. But what is worse about *Rape Day* than, say, brutally murdering civilians in *Grand Theft Auto*? This section explains why a certain type of imaginative engagement is wrong, appealing to a moral difference between virtual sexual assault and virtual homicide.

The problem is an instance of the gamer’s dilemma (Luck, 2009), which arises from the apparent incompatibility of two plausible intuitions: (i) virtual murder is morally permissible; (ii) virtual paedophilia is morally impermissible. Virtual murder seems to be harmless fun, but the same could be said of virtual paedophilia (at least in principle). Both are fictional activities with comparably bad real-world counterparts. Without a relevant moral difference, Morgan Luck concludes that we should treat them equally; either we condemn them both or allow them both.

Many have tried to solve this dilemma.¹³ Here I will sketch my own alternative proposal. I will argue that *fantasies* are less distanced from reality than other sorts of imaginings, and so fantasizing about some wrongdoing, such as rape, is morally impermissible. Fantasies, on my view, are a class of imaginings that entertain their fictional content as a real-life possibility. I think the necessary and sufficient conditions to fantasize about A are

12 One reviewer noted that there are some exceptions. For example, on some versions of Kantian ethics, consequences seem to be explanatorily inert. I thus restrict my argument to those moral theories that assign at least some non-trivial moral significance to the expected outcome of a given act.

13 See Bartel (2012, 2020), Ali (2015), and Young (2016).

(i) to actively and extensively imagine A (ii) in the context of your own life and (iii) as the object of some pro-attitude, such as desire or approval. Of course, you can fantasize about being a different sort of person or having a different life, but it is still *you* who is having a different life, not someone else. And if your identity is preserved in the fantasy, then you are still fantasizing about something within the context of your own life. This isn't normally the case when, for example, you engage with *The Sopranos*. You may have a pro-attitude towards Tony, but you extensively imagine the fictional events in the context of Tony's life, not yours.

While it is difficult to think of other reasons to engage with virtual rape beyond the purpose of fantasizing, virtual murder is different. When we kill in a video game, we are not normally fantasizing about committing murder. It is a more *impersonal* kind of imagining, which concerns a more distant event. We entertain the fictional representation of *a* life, not a possibility in *our* own life. It is in this subjective sense that fantasizing about sexual assault is often more connected to reality than other kinds of imaginings. Such fantasies do not need to capture objectively close possible worlds: it may still be wrong for rape fantasies to take place in improbable magical worlds, such as *The Lord of the Rings*. For a fantasy to be morally wrong, it is sufficient that the imaginer intentionally and extensively projects their identity into the imaginative context where the immoral act is performed, embracing their imagining with a pro-attitude. Consider two cases:

The Uncorrupted Player. A psychologist performs virtual sexual molestation in computer games to impersonally observe the typical features of this activity, without any intention to fantasize.

The Degenerate Player. A morally degenerate player who fantasizes about murdering others ... and so plays the game with the sole purpose of enacting his fantasies. (Ali, 2015, p. 269)

Not all instances of virtual murder are permissible and not all instances of virtual sexual assault are impermissible. It does not seem particularly wrong for *The Uncorrupted Player* to commit virtual sexual molestation. He is not conceiving of immoral fictional contents as real-life possibilities but maintains them wholly separated from the context of his own life. Conversely, when *The Degenerate Player* uses the gaming platform as a source of satisfaction for his murderous fantasies, something feels wrong. Whether of rape or murder, the wrongness of an imagining increases proportionally to its subjective closeness to the actual world. The gamer's dilemma is then explained away by the fact that it is quite common to entertain murder as an impersonal imagining, while it is comparatively rare to imaginatively engage with sexual violence if not to fantasize.

My proposal differs from that of Erick Ramirez (2020), who argues that video games involving virtual sexual molestation often have a high degree of perspectival fidelity that enables the subject to experience such acts as close to reality. Facts about the design of the game, I think, do not justify a difference in moral permissibility. If *The Degenerate Player* has a virtual experience of murder that is very poor in terms of realism and perspectival fidelity, but obsessively fantasizes about it, he is still doing something intuitively wrong. Vice versa, if *The Uncorrupted Player* has a realistic virtual rape experience, but also has a

cognitively enhanced brain capable of separating very precisely fantasies from impersonal imaginings, then it would be difficult to say what is wrong about his imaginative engagement. The simulation's design features often affect the subject's fictional experience, but it is ultimately the subject's internal motivational state and degree of identification that bear the crucial moral significance.¹⁴

But why is fantasizing about something evil morally impermissible? What are the wrong-making features of immoral fantasies? If we look for a moral justification for the intuition that we should permit (most instances of) virtual murder, but not (most instances of) virtual sexual assault, it is not enough to identify a psychological difference between the two. So, I will now show in addition how immoral fantasies are problematic on a wide range of views.

First, imaginatively entertaining a given wrongdoing as a fantasy may unacceptably increase the likelihood of actually committing said wrongdoing. Or so the instrumental-wrongness argument goes. After all, it is hard to see how anybody would find the contents of *Rape Day* appealing unless they were inclined to realize them in real life if they could. But perhaps it is the potential rapist who produces rape fantasies, rather than the rape fantasies that produce a potential rapist. While this may well be true, such fantasies could nonetheless reinforce pre-existing dangerous inclinations or help the would-be rapist explore hypothetical scenarios concerning how to fulfil his dispositions. This argument, however, ultimately comes down to the available empirical evidence on the harms and benefits of violent works of fiction, which is problematic and, at best, inconclusive when it comes to proving whether such fictions are overall harmful or beneficial.¹⁵

It is nonetheless worth noting that immoral fantasies may also bring about a different kind of negative consequence—not towards others but towards oneself. The rape fantasist's imaginative inclinations are likely to have a harmful impact on his own well-being and personal flourishing, so that he would be better off without them. He would be able to engage with healthier human relationships, integrate more into societal contexts, and significantly improve his emotional life. If so, immoral fantasies could be condemned as cases of *self-harm*.

Independently of any possible consequence, imaginatively entertaining something immoral in the context of one's own life might be *intrinsically* wrong. Views of this sort have been defended by a number of authors. Gaut (1998), for instance, contends that fantasizing about rape expresses an intrinsically immoral attitude towards women as a kind. But why is it that virtual homicide does not equally amount to expressing an intrinsically immoral attitude towards, say, humanity as a kind? It seems plausible to say that only immoral *fantasies*—imaginings which regard a wrongdoing as a real-life possibility for the imaginer to undertake—are authentically expressive of an immoral attitude.

My explanation of the gamer's dilemma also nicely fits with Christopher Bartel and Anna Cremaldi's view (2018). On their account, our real-life immoral desires can be

14 In this respect, my view owes a lot to Bartel (2020).

15 See Luck (2009) and Bartel (2020).

cultivated through our fantasies and, just as cultivating racist beliefs, cultivating these immoral desires is wrong even if they remain unacted upon. But again, why can sexual fantasies cultivate real-life immoral desires while more impersonal imaginings of murder cannot? This is because, in the case of fantasies, the imaginative distance between fiction and life decreases significantly. Fantasies are thus uniquely suited to reveal and exercise real-world moral sensibilities. Bartel (2020) proposes a similar way of addressing the gamer's dilemma. Contrary to virtual homicide, Bartel says, most ordinary cases of virtual rape are not redeemed by benign motives. There is a worrying real-world motivation to engage with virtual sexual violence, hence the difference in moral permissibility (2020, p. 140). I find this view compatible with my take on malevolent fantasies. Fantasizing about sexual assault as a real-life possibility is normally driven by immoral motives, and these motives are especially morally bad because they conceive of their object of imagination as a real-life possibility.

To remain compatible with as many options as possible, I do not characterize double-standard moralism as committed to any particular view about whether fantasies can be instrumentally or intrinsically morally wrong (or both). For double-standard moralism to be true, what matters is that at least a class of fictive imaginings can be morally evaluated, and I suggested a few ways to do so. But note that the gamer's dilemma is about the moral differences between two types of *fictional* acts: virtual sexual violence and virtual murder. When it comes to discontinuity—that is, the moral differences between analogous *real-world* and *fictional* acts, such as real and virtual homicide—I still believe that differing consequences constitute a major part of the explanation for the additional comparative wrongness.

7. Belief-Relative and Fact-Relative Discontinuities

I have proposed that what justifies normative discontinuity is an asymmetry of consequences, but I have not specified whether the difference-maker is a disanalogy in our *beliefs* about consequences or brute *facts* about asymmetrical consequences. In this section, I make two points: (i) the discontinuity in the appropriateness of our moral responses across fiction and reality is *belief-relative* or *fact-relative* depending on the possibility of moral luck; (ii) regardless of the agent's attitudes, the badness (or goodness) of fictional states of affairs never amounts to the badness (or goodness) of isomorphic real-world states of affairs, justifying a further fact-relative axiological discontinuity.

Prima facie, it makes a moral difference whether one believes that one's acts have bad outcomes. This is especially plausible in cases where we assess the appropriateness of one's moral responses. If you are in a virtual world but believe it is real, intending to kill people and doing so virtually seems very blameworthy. Conversely, in *Ender's Game*, where Ender is deceived into unintentional killings because they are depicted as fictional, we do not blame him as if he intentionally murdered these people. By appealing to such cases of radical deception, one may argue for:

Belief-relative discontinuity. The different justified beliefs of the agent about the results of their own acts make appropriate a different moral evaluation between what the agent takes to be fictional or to be real, regardless of the act's actual outcome.

I prefer to characterize belief-relative discontinuity in terms of *justified* beliefs, since in ordinary fictive engagements we have good epistemic reason to think that fictional acts do not bring about analogous real-life consequences. For when one kills a fictional character in a video game, or hopes for Tony Soprano to succeed, one justifiably believes that no person will truly suffer. By contrast, if one intentionally participates in a real-life murder, or supports Tony Soprano's real counterpart, one is typically cognisant of the harmful results.

Unjustified beliefs might also matter to the moral evaluation of fictive engagements. Consider Albert, who unreasonably believes in a conspiracy theory (with no evidence whatsoever) that claims that a video game's characters are genuine persons—sentient, rational creatures. If Albert were then to enjoy killing what he believes and hopes are real people, he would be doing something blameworthy, even though Albert's acts are certain to cause no harm on the available evidence.

To test the plausibility of justified-belief-relative discontinuity, consider the following case:

The Deceived Rapist. A player has been tricked into believing *Rape Day* to be reality and his virtual acts of non-consensual sex to be authentic. He continues to perform his fictional wrongdoings, but no person is actually harmed as a result.

From his point of view, The Deceived Rapist has every reason to believe that he is actually committing rape. The fact that he is not doing so is only a result of a deception—a fact, let's assume, completely outside of his control. Many have the intuition that if two agents form qualitatively identical intentions to maliciously harm an innocent person and have the same level of justification for believing in their success, they perform equally blameworthy acts, even if one person succeeds while the other fails by fortune. This intuition is captured by the Kantian view that:

(A) The agent's blameworthiness cannot depend on mere luck.¹⁶

If two people ought to be assessed in the same way when all things are equal except for factors beyond their control, then Ender is not morally blameworthy, and The Deceived Rapist is just as morally blameworthy as an actual rapist. Since The Deceived Rapist couldn't have known about the fictional status of his own actions, he doesn't deserve better or worse assessment than his real counterpart. Any moral difference in the assessment of The Deceived Rapist's imaginative responses would be entirely a matter of luck from his point of view. The relevant implication of (A) for discontinuity is that the moral appropriateness of our responses depends on whether these are responses to things we justifiably *believe* to be real or justifiably *believe* to be fictional, irrespective of whether they actually *are* fictional or real.

In contrast, some theorists deny (A) and accept that facts beyond the agent's control are relevant to the moral evaluation of the agent (e.g. Browne, 1992). If The Deceived Rapist can be evaluated on a different level of blameworthiness due to elements of 'moral luck', then maybe we should blame him less than a comparable person in real life. He didn't

16 This principle is famously proposed by Kant (1784/1998, 4:394).

actually harm anyone, even though this wasn't up to him. For The Deceived Rapist has as much justification as his real counterpart to believe he is doing harm, moral luck provides grounds for:

Fact-relative discontinuity. The different outcomes of the agent's acts make appropriate a different moral evaluation between isomorphic fictional and real-world acts, regardless of whether the agent justifiably believes to be doing something fictional or not.

On this view, we ought to blame more actual rapists than would-be rapists, even if the failure in harming the intended victim is outside the rapist's control.

For the purposes of this paper, it suffices to say that whether the moral discontinuity in *appropriateness* of the agent's imaginings is sensitive to their beliefs depends on the possibility of moral luck. But even if we judged the fictional wrongdoings of The Deceived Rapist to be as morally inappropriate as the wrongdoings of a real rapist (and him to be just as a horrible person), we may further ask if his fictional acts are *equally bad*. A strong case can be made for the claim that a discontinuity in badness applies independently of the rape fantasist's ill-intentioned internal states. We will still prefer the wrongdoings of The Deceived Rapist to occur in fictional situations rather than real life—we will likely consider his fictive engagements to be less bad. Because this fact-relative discontinuity applies in spite of symmetrical beliefs, I argue, even the worst imaginings meet different moral standards from those applying to real life.

Do not get me wrong. The imaginative responses of The Deceived Rapist are morally bad whatever ethical theory one embraces—whether it be virtue ethics, deontology, or consequentialism. Even without believing it to be real, utilizing *Rape Day* for fantasizing purposes indicates the virtual rapist's character flaws, reveals a misogynistic attitude that diminishes his capacity to treat others with respect, and may unacceptably raise the risk of future real-world consequences. Unsurprisingly, *Rape Day* was refused public distribution. But since I am interested in normative discontinuity, I am not asking whether enjoying *Rape Day* is bad. I am asking whether enjoying *Rape Day* is *as bad as* committing real-world non-consensual sex.

When we make axiological comparisons between isomorphic fictional and real-world states of affairs—by claiming that one is *better* or *worse* than the other—discontinuity is often implicitly assumed. For example, Neil Levy (2002, p. 321) argues that virtual paedophilic material may reduce the harm to real people and save lives by supplying an alternative outlet for satisfying paedophilic tendencies. This forward-looking argument is also proposed by Dannielle Cisneros (2002), who suggests that virtual child pornography may protect actual children from abuse, since potential paedophiles could practice their activities in victimless virtual worlds as a substitute for real violence. It is debatable whether virtual sexual violence can help release immoral desires—it might instead feed or reinforce them. But suppose tomorrow's research proved that playing *Rape Day* is a reliable preventative measure against rape. Would it not be *better* for an evil thing to occur in fiction? It is challenging to disagree. A fictive imagining, however immoral, is not as bad as performing analogous activities in real life. And if any real-world wrongdoing outweighs the isomorphic fictional wrongdoing, then there is evidence of a deep moral discontinuity that differentiates the objects of imagination.

Call this *The ‘Better Than’ Test*: if it is better for evil (or good) things to transpire in fiction rather than in life, then double standards sharply divide the moral status of these two domains. The ‘Better Than’ Test is a test to detect moral discontinuity. If the test shows a clear preference for something bad to happen in imaginative contexts, then this is evidence for discontinuity. If, instead, a work of fiction fails the test because, normatively speaking, this choice between options is indifferent, then we would obtain evidence for continuity.¹⁷ In a case of continuity, we would have no reason to prefer a given bad (or good) to happen in the fictional domain. It would be arbitrary to display any preference since the two options would be equally bad (or good) in all relevant respects. But we do display a strong axiological preference. The ‘Better Than’ Test thus provides grounds for embracing double standards between fiction and life.

The kind of axiological discontinuity identified by The ‘Better Than’ Test dispenses with any appeal to the agent’s attitudes; it is, I believe, fact-relative. Even in cases such as The Deceived Rapist, the test remains successful in spite of symmetrical beliefs and intentions. This residual moral discrepancy requires further justification independent of the rape fantasist’s internal states, and this is provided by asymmetrical outcomes. Immoral fictive engagements may well lead to undesirable states of affairs, but these consequences largely differ from those that corresponding real-life activities would have.

8. Conclusion

Normative continuity and discontinuity can be mistakenly seen as ‘all or nothing’: if there is continuity, the moral norms of the fictional are the same as the real; if there is discontinuity, there are no moral norms for the fictional. This paper defended the middle-ground view—*double-standard moralism*. On this view, the imaginative domain is subject to *sui generis* moral standards, ones that allow us to enjoy an extension of the permissible that the actual world cannot, and should not, accommodate. Double-standard moralism is a more robust version of response moralism since any account that is not committed to normative discontinuity has unacceptable implications, as shown by my previous example, *Ender’s Game*.

How, then, can double standards be justified? I claimed that a moral response to a fictional event can be apt, while said response would be unacceptable for the isomorphic real-world event because the consequences of fictional events largely differ, and so does their moral significance. This asymmetry can be understood in a belief-relative or fact-relative way. While our beliefs might be relevant to assessing the moral appropriateness of our fictive imaginings, a further fact-relative discontinuity characterizes our axiological judgements. On both fronts, double standards between fiction and life are morally justified.¹⁸

17 The Voodoo Video Game fails The ‘Better Than’ Test, thus indicating a case of continuity.

18 For helpful comments, I am especially grateful to Berys Gaut, Theron Pummer, Roger Crisp, and Louise Mackie, who endured many drafts of this paper. I am also grateful to Christopher Bartel, David Brink, Tom Kaspers, Julia Driver, Matthew Moore, Joel Joseph, Melania Pellitteri, Luca Stroppa, Nathan Bray, Jonas Hertel, Conor Mockler, Tomi Francis, James Glover, and two anonymous reviewers. I am finally thankful for the support of an Open Philanthropy Scholarship.

Mattia Cecchinato
 University of Oxford, UK
mattia.cecchinato@philosophy.ox.ac.uk

References

- Ali, R. (2015). 'A new solution to the gamer's dilemma'. *Ethics and Information Technology*, 17, pp. 267–274.
- Bartel, C. (2012). 'Resolving the gamer's dilemma'. *Ethics and Information Technology*, 14, pp. 11–16.
- Bartel, C. (2020). *Video games, violence, and the ethics of fantasy: killing time*. London: Bloomsbury Academic.
- Bartel, C. and Cremaldi, A. (2018). 'It's just a story': pornography, desire, and the ethics of fictive imagining'. *British Journal of Aesthetics*, 58, pp. 37–50.
- Browne, B. (1992). 'A solution to the problem of moral luck'. *The Philosophical Quarterly*, 42, pp. 343–356.
- Card, O. S. (1985). *Ender's game*. New York: Tor Books.
- Carroll, N. (2004). 'Sympathy for the devil', in Greene, R. and Vernezze, P. (eds.), *The Sopranos and philosophy: I kill therefore I am*. Chicago: Open Court, pp. 121–136.
- Cisneros, D. (2002). 'Virtual child pornography on the internet: a virtual victim?'. *Duke Law & Technology Review*, 1, pp. 1–8.
- Cooke, B. (2014). 'Ethics and fictive imagining'. *The Journal of Aesthetics and Art Criticism*, 72, pp. 317–327.
- Currie, G. (2020). *Imagining and knowing: the shape of fiction*. Oxford: OUP.
- Desk Plant. (2019). *Rape day*, [Game] PC. Unreleased.
- Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb*. (1964). [Film] Directed By Stanley Kubrick. England: Hawk Films.
- Feagin, S.L. (1984). 'Some pleasures of imagination'. *The Journal of Aesthetics and Art Criticism*, 43, pp. 41–55.
- Gaut, B. (1998). 'The ethical criticism of art', in Levinson J. (ed.), *Aesthetics and Ethics: Essays at the Intersection*. Cambridge: Cambridge University Press, pp. 182–203.
- Gaut, B. (2007). *Art, emotion and ethics*. Oxford: OUP.
- Gilmore, J. (2020). *Apt imaginings: feelings for fictions and other creatures of the mind*. Oxford: OUP.
- Hazlett, A. (2009). 'How to defend response moralism'. *British Journal of Aesthetics*, 49, pp. 241–255.
- Kant, I. (1784/1998). *Groundwork of the metaphysics of morals*. M. Gregor (ed., tr.). Cambridge: Cambridge University Press.
- Levy, N. (2002). 'Virtual child pornography: the eroticization of inequality'. *Ethics and Information Technology*, 4, pp. 319–323.
- 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, pp. 31–36.
- Nichols, S. (2006). 'Just the imagination: why imagining doesn't behave like believing'. *Mind & Language*, 21, pp. 459–474.
- Ramirez, E. J. (2020). 'How to (dis)solve the gamer's dilemma'. *Ethical Theory and Moral Practice*, 23, pp. 141–161.
- Rawls, J. (1971). *A theory of justice*. Harvard, MA: Harvard University Press.

- Rockstar North. (1997–2013). *Grand Theft Auto*, [Game] PlayStation. New York: Rockstar Games.
- Rosenberg, R. S., Baughman, S. L. and Bailenson, J. N. (2013). 'Virtual superheroes: using superpowers in virtual reality to encourage prosocial behavior'. *PLoS One*, 8, pp. 1–9. doi:[10.1371/journal.pone.0055003](https://doi.org/10.1371/journal.pone.0055003).g002
- Schindler's List*. (1993). [Film] Directed by Steven Spielberg. Los Angeles: Universal Pictures.
- Smuts, A. (2013). 'The ethics of singing along: the case of 'mind of a lunatic''. *Journal of Aesthetics and Art Criticism*, 71, pp. 121–129.
- The Sopranos*. (1999–2007). [TV Series] Created by David Chase. New York: HBO.
- Walton, K. L. (1994). 'Morals in fiction and fictional morality'. *Proceedings of the Aristotelian Society*, 68, pp. 27–50.
- Yoon, G. and Vargas, P. T. (2014). 'Know thy avatar: the unintended effect of virtual-self representation on behavior'. *Psychological Science*, 25, pp. 1043–1045.
- Young, G. (2016). *Resolving the gamer's dilemma: examining the moral and psychological differences between virtual murder and virtual paedophilia*. London: Palgrave Macmillan.