



Resisting the Gamer's Dilemma

Thomas Montefiore¹ · Paul Formosa²

Accepted: 22 June 2022
© The Author(s) 2022

Abstract

Intuitively, many people seem to hold that engaging in acts of virtual murder in videogames is morally permissible, whereas engaging in acts of virtual child molestation is morally impermissible. The Gamer's Dilemma (Luck in *Ethics Inf Technol* 11:31–36, 2009) challenges these intuitions by arguing that it is unclear whether there is a morally relevant difference between these two types of virtual actions. There are two main responses in the literature to this dilemma. First, attempts to *resolve* the dilemma by defending an account of the relevant moral differences between virtual murder and virtual child molestation. Second, attempts to *dissolve* the dilemma by undermining the intuitions that ground it. In this paper, we argue that a *narrow* version of the Gamer's Dilemma seems to survive attempts to resolve or dissolve it away entirely, since neither approach seems to be able to solve the dilemma for *all* cases. We thus provide a contextually sensitive version of the dilemma that more accurately tracks onto the intuitions of gamers. However, we also argue that the intuitions that ground the narrow version of the Dilemma may not have a moral foundation, and we put forward alternative non-moral normative foundations that seem to better account for the remaining intuitive difference between the two types of virtual actions. We also respond to proposed solutions to the Gamer's Dilemma in novel ways and set out areas for future empirical work in this area.

Keywords Gamer's Dilemma · Videogame ethics · Virtual harm · Computer game ethics · Ethics and technology

Introduction

Intuitively, many people seem to hold that engaging in acts of virtual murder in videogames is morally permissible, whereas engaging in acts of virtual child molestation, sexual assault, or rape in videogames is morally impermissible. The Gamer's Dilemma, as set out by Luck (2009), challenges these intuitions by arguing that it is unclear whether there is a morally relevant difference between acts of virtual murder and acts of virtual child molestation that makes the first permissible and the second impermissible. This leaves the gamer to face the following dilemma: either it is morally impermissible to engage in virtual murder, which would make playing many popular videogames immoral, or it

is morally permissible to engage in virtual child molestation, which is an outcome many people find morally repugnant. There are two main responses in the literature to this dilemma. First, attempts to *resolve* the dilemma by defending an account of the relevant moral differences between virtual murder and virtual child molestation. This allows one to defend the intuition that virtual murder is morally permissible whereas virtual child molestation is not. Second, attempts to *dissolve* the dilemma by undermining the intuitions that ground it by denying, in some or all cases, that virtual murder is morally permissible or that virtual child molestation is morally impermissible.

We start the paper by reviewing the state of the literature, exploring existing attempts to resolve and dissolve the dilemma, and considering further novel solutions. Drawing on this analysis, we argue for some degree of both resolving and dissolving the dilemma. This leaves behind, we argue, a *narrow* version of the Gamer's Dilemma. This provides a contextually sensitive version of the dilemma that more accurately tracks onto the intuitions of gamers. However, we also argue that the intuitions that ground the narrow version of the dilemma may not have a moral foundation, and we suggest promising alternative non-moral normative

✉ Thomas Montefiore
tom.montefiore@hdr.mq.edu.au

✉ Paul Formosa
Paul.Formosa@mq.edu.au

¹ Department of Philosophy, Macquarie University, Sydney, Australia

² Department of Philosophy, Macquarie University, Sydney, Australia

foundations, such as matters of taste, suberogatory action, narrative tropes, or aesthetic conventionality, that seem to better account for the remaining intuitive difference between the two types of virtual actions. We also outline the implications of our argument for future theoretical and empirical work on the Gamer's Dilemma.

The Gamer's Dilemma: state of the literature

In the literature on videogame ethics, there are arguably two key views concerning the moral permissibility of virtual actions enacted by gamers. The first view, which we call the 'moralist view', holds that a gamer's virtual actions are subject to moral evaluation. The moralist view is expressed, for example, in one's discomfort when asked to enact virtual rape in *RapeLay* (2006), participate in a virtual mass shooting in *Call of Duty: Modern Warfare 2* (2009), or hypothetically, molest a child in a videogame. The moralist view is often defended by appealing to the supposed harm that virtual actions, which would be immoral if enacted in real life, cause the gamer and others in their community. For example, the fact (assuming for now that it is a fact) that engaging in a virtual action has the harmful effect of causing a gamer to be more aggressive, and thus a more vicious person and more of a danger to others, *makes* that virtual action morally evaluable. The second view, which we call the 'amoralist view', holds that virtual actions are harmless and are thus not subject to moral evaluation (see Patridge, 2011; Ostritsch, 2017; Young, 2017a, 2017b). This view is associated with the idea that videogames are morally insulated from the actual world in a so-called "magic circle" (see Huizinga, 1938; Juul, 2008, Saleen & Zimmerman, 2003), and therefore virtual actions lack any moral character. This view is expressed in response to seemingly morally objectionable virtual actions via the aphorism: 'It's only a game!' (see Patridge, 2011). After all, the amoralist claims, virtual actions aren't real and don't harm anyone. For example, when a gamer commits virtual murder by running over an innocent pedestrian in *GTA V* (2013), they are not doing anything morally wrong on this view since their virtual action only results in pixels on a screen being moved around.

Luck's (2009) "Gamer's Dilemma" serves as a central problem in this literature because it clearly identifies the intuitive force and the incompatibility of both the moralist and amoralist views. The Gamer's Dilemma forces the moralist to either reject the intuition that virtual murder is morally permissible or to account for how it is morally permissible given her moralist commitments, and it forces the amoralist to either reject the intuition that virtual child

molestation is morally impermissible or to account for how it is morally impermissible given her amoralist commitments.

We can state the *general form* of the Gamer's Dilemma as follows:

- G1.** Gamers hold the justified moral intuition that *in all cases* virtual murder is morally permissible.
- G2.** Gamers hold the justified moral intuition that *in all cases* virtual child molestation is morally impermissible.
- G3.** Gamers cannot defend the claim that there is a relevant moral difference between virtual murder and virtual child molestation such that the first is morally permissible and the second is morally impermissible.

To *resolve* the dilemma is to reject G3. To *dissolve* the dilemma is to reject either G1 or G2 (or both).

Attempts at resolving the Gamer's Dilemma are the most prevalent in the literature and primarily involve identifying one or several morally relevant features of virtual child molestation that renders it morally impermissible while maintaining the moral permissibility of virtual murder. Some offered resolutions ground the immorality of virtual child molestation on: the virtue ethical consideration that it harms the gamer's character (McCormack, 2001; Bartel, 2020); the basis that virtual child molestation constitutes child pornography (Bartel, 2012); or the harm virtual child molestation causes particular social groups (Levy, 2002; Patridge, 2011, 2013). We consider this literature further in "[Resolving strategies: the moral difference between virtual murder and virtual child molestation](#)" section. Dissolving the dilemma can be attempted in several ways. Without outright denying the existence of the intuitions in the dilemma (see Haynes, 2015; Prigg, 2009; and Cunningham et al., 2011 on the popularity of violent videogames), the moralist tends to reject the justifiability of the intuition that virtual murder is permissible (Tillson, 2018), whereas the amoralist tends to reject the justifiability of the intuition that virtual child molestation is impermissible (Saleen & Zimmerman, 2003). An alternative approach has been to argue that that the necessity or '*in all cases*' clauses in G1 and G2 are false. This has been done by arguing that the intuitions of gamers regarding the moral permissibility of virtual actions are, among other things, justifiably sensitive to differences in narrative, structural context, and degree of abstractness, rather than simply to the *action types* of virtual murder or virtual child molestation (see e.g., Ali, 2015; Ramirez, 2020; Öhman, 2020). We consider this literature further in "[Dissolving strategies: adding context](#)" section.

This focus on context sensitivity is a particular form of the more general move of aiming to limit the scope of the dilemma by limiting the range of cases in which the

intuitions that ground the Gamer's Dilemma hold. We can state the *narrow form* of the Gamer's Dilemma as follows:

- N1. Gamers hold the justified moral intuition that *in some cases* virtual murder is morally permissible.
- N2. Gamers hold the justified moral intuition that *in those same cases* virtual child molestation is morally impermissible.
- N3. Gamers cannot defend the claim that there is a relevant moral difference between virtual murder and virtual child molestation *in those specific overlapping cases* such that the former are morally permissible and the latter are morally impermissible.

Clearly, the narrow form of the dilemma comes in degrees. It could hold in just one case, or it could hold in many (but not all) cases. Narrowing can be achieved through either dissolving or resolving strategies. If it is used to limit 1 or 2, it is a dissolving strategy, and if it is used to limit 3, it is a resolving strategy.

Narrowing can be pursued via a resolving strategy as follows. Assume, for example, that in *some* specific games, enacting virtual murder will not promote violence or aggression in the player *to any degree*, whereas in that *same game* enacting virtual child molestation will *strongly promote* sexual violence against children. If this could be shown, then we have grounds for resolving the dilemma for *those specific cases*, since in those specific cases we can show a relevant moral difference between the two acts (namely, that one promotes violence to some degree and the other does not). However, this only narrows rather than fully resolves the Dilemma as it still (potentially) leaves open a range of other cases where there is no such moral difference between the two acts *and* people still have the justified intuition that *in such cases* virtual murder is morally permissible and virtual child molestation is not. For example, it might be that in some other games neither virtual murder nor virtual child molestation has *any* real-world harmful impacts *whatsoever*, and yet we still hold the justified intuition that *in those specific games* the former but not the latter virtual action is morally permissible.

However, more commonly, narrowing of the dilemma is achieved through a dissolving strategy. Luck (2018), for example, pursues such a strategy in response to the dissolution offered by Ali (2015). Ali argues that the intuitions in the Gamer's Dilemma may fail in simulation games and storytelling games. In simulation games, Ali argues, the high degree of agency that players have in determining what in-game actions they perform means that both virtual murder and virtual child molestation can be morally impermissible. In contrast, in storytelling games, the comparative lack of agency that players have in determining their in-game actions means that both virtual murder and virtual child molestation can be morally permissible. If Ali is right and

his argument can be shown to successfully hold for all cases, then both virtual murder and virtual child molestation are permissible in *all* storytelling games, and both are impermissible in *all* simulation games. If it could also be shown that all games fall neatly into the categories of simulation and storytelling games (which is highly questionable—see Formosa et al., 2016), then this argument would *completely* dissolve the Dilemma as it would leave *no cases* where we can hold the justified moral intuition that virtual murder is permissible and virtual child molestation is impermissible *in the same specific context*. In response, Luck argues that while Ali's claims may hold in *some* cases, such as those with very high or very low degrees of player agency, there is still a range of cases for storytelling games where virtual murder will be justifiably intuited as permissible and virtual child molestation as impermissible, and a range of cases for simulation games where virtual murder will justifiably be intuited as permissible and virtual child molestation as impermissible. Further, Luck (2018) argues that in other game modes, such as competition games like *Counter Strike* (2000), the intuitions that ground the Gamer's Dilemma may still hold. Therefore, Luck concedes that the Gamer's Dilemma does not hold in *all* cases (contra the *general form*), but he maintains that it still holds in *some* cases. This leaves us with a *narrow form* of the Gamer's Dilemma.

Once we have successfully narrowed the Gamer's Dilemma, the important and still largely unanswered question in the literature is *how narrow* the resulting dilemma is and what (if anything) is the moral and practical significance of the narrowed dilemma. We attempt to answer both questions below. An important upshot of our argument is that it encourages us to move beyond generic claims about the status of virtual actions, such as murder, to a more nuanced and contextualised discussion of specific actions in specific games played in specific ways in specific social contexts.

Resolving strategies: the moral difference between virtual murder and virtual child molestation

There are two promising strategies for resolving the Gamer's Dilemma by identifying a morally relevant difference between virtual murder and virtual child molestation. Our aim is to develop the best versions of these to assess their effectiveness. The *instrumental* approach focuses on the different *degrees* or *types* of *harm* done to either the gamer's moral character or to other people resulting from the gamer's in-game actions. In this case, it is not the virtual acts themselves but their different real-world consequences that morally differentiates them. The *intrinsic* approach focuses on an *intrinsic* feature of the virtual act itself (which may also have instrumental downstream effects), such as it being

inherently disrespectful toward a certain group. In this case, it is something about the virtual acts themselves that morally differentiates them. While we are not convinced that either strategy *completely* resolves the Gamer's Dilemma (without further empirical evidence), we do argue below that both have (albeit differing) potential to *narrow* it.

The instrumental approach requires, one, that playing videogames has real-world consequences and, two, that there is a moral difference between the real-world consequences of enacting virtual murder and virtual child molestation. We have robust evidence for the first claim in terms of the established literature on violent videogames and the emerging literature on prosocial games. A recent meta-analytic review claims that this research “demonstrates a consistent relation between violent videogame use and increases in aggressive behavior, aggressive cognitions, and aggressive affect and decreases in prosocial behavior, empathy, and sensitivity to aggression” (APA, 2015; cf. Ferguson, 2008). This evidence shows us that enacting violence in videogames, rather than reducing aggression by helping players to ‘blow off steam’, increases aggression through players “rehearsing aggressive and violent thoughts and actions” (Anderson et al., 2010). However, while there is a clear effect on aggression from playing violent videogames, the effect size is not large, and there are many other factors that may swamp its real-world impacts (Anderson et al., 2010), leaving its moral significance unclear and its impacts in need of further study (APA, 2015). However, the effects of videogames depend on their content, and “the same basic social–cognitive processes should also yield prosocial effects when game content is primarily prosocial” (Anderson et al., 2010; cf. Staines et al., 2019). This literature shows us that playing videogames can have, albeit limited, real-world impacts, and that those impacts could be positive or negative depending on the content of the game.

However, we lack clear evidence in terms of the crucial second claim that there is an important moral difference in the real-world impacts of virtual murder and virtual child molestation. Understandably, little empirical research has been conducted on the real-world consequences of engaging in virtual child molestation, and how such engagement may differ morally from other kinds of virtually immoral activity such as virtual murder. For example, one potential candidate to justify this argument is that child pornography has been used by child sex abusers to normalise their abuse to their victims (Woo, 2004, p. 727). It is possible that a similar concern about normalising sexual abuse could be levelled against virtual child molestation. However, it is unclear whether a similar concern, in terms of normalising murder or violence, also applies to virtual murder. It could also be that both virtual murder and virtual child molestation produce harms, but the *type* and *degree* of harms they produce are significantly different. For example, it might be

that enacting virtual murder encourages low-level aggressive behaviours and cognitions, whereas enacting virtual child molestation encourages child molestation and abuse. If that were so (and this is only speculation), then the type and degree of harms produced are much greater in the second case, which would constitute a relevant moral difference. However, there is a lack of evidence to support any such claims, although an absence of evidence is not evidence of an absence. We therefore consider this instrumental approach to be a live option that must be sidelined as there is currently insufficient empirical evidence to conclusively support or reject it.

Alternatively, one may attempt to resolve the Gamer's Dilemma via an intrinsic approach. An intrinsic approach focuses on the moral features of virtual representations of child molestation *in themselves* and in what way such features are morally distinct from virtual murder in terms of permissibility. Bartel (2012), for instance, argues that because virtual child molestation involves the depiction of sexual acts involving children, and such acts, while *virtual*, are nonetheless *actual* instances of child pornography which is morally wrong, then it follows that enacting virtual child molestation is also morally wrong (Bartel, 2012, p. 13). Bartel thus argues that it is the fact that virtual child molestation is real child pornography that distinguishes it morally from virtual murder and thus resolves the dilemma by identifying a morally relevant difference. One objection to this view (see Luck & Ellerby, 2013 for others) is that, even if virtual child molestation is actual child pornography, it may lack a feature that makes actual child pornography morally wrong, namely that it involves the actual abuse of or photography of real children who are thereby harmed (see Young, 2016). However, while virtual child molestation *qua* virtual child pornography may not be morally objectionable on such grounds, it may still cause “cultural harm” through the obscenity of what virtual child pornography depicts (McGlynn & Rackley, 2009).

Further, both virtual child pornography and ‘actual’ child pornography may share the troubling and ontologically demanding feature of generating actual erotic delight in those who appropriately engage with it (that is, play the game as intended by the developer—see Ali, 2015, p. 270; Zagal, 2009). This feature of virtual child molestation *qua* virtual child pornography, that to engage with it appropriately is to find it erotic, seems to capture a morally relevant intuitive dimension of virtual child molestation that is absent in (most) virtual murder when it is appropriately engaged with. But what is morally objectionable about treating a virtual depiction of child molestation erotically? Following Patridge (2011), it can be argued that one is justified in morally objecting to imagery which advances harmful social meaning via its association with incorrigible and harmful social norms towards particular social groups, such as

those levelled at women or racial minorities. Drawing on Levy (2002), one could argue that what makes virtual child molestation, when treated as erotic, morally distinct from virtual murder in terms of permissibility is the necessary eroticisation of inequality that it promotes. Levy (2002, p. 322) argues that what makes eroticising inequality morally impermissible is that it harms efforts by women to achieve sexual equality. However, the eroticisation of inequality does not harm children, Patridge (2013) argues, because virtual and certainly actual child molestation is not associated with accepted social norms towards children. Therefore, it is unreasonable to claim that virtual child molestation harms children as a group because accepted social norms promoting the unequal sexualisation of children do not exist, unlike, for example, accepted social norms which promote the unequal sexualisation of women (Patridge, 2013, p. 29).

While harm done to women via the eroticisation of inequality may justify the moral impermissibility of virtual child molestation, it arguably fails in solving the Gamer's Dilemma because it does not seem to *causally* justify what is triggering our intuitive moral discomfort towards virtual child molestation, which surely has something to do with the treatment of *children* (not women), and thereby merely *accidentally* justifies our intuition (Patridge, 2013, p. 30; Young, 2016, p. 80). In response to Patridge, we argue that one may avoid this concern by resisting the notion that children are not subject to social norms that sexualise them, and therefore claim that children *are* directly harmed via the eroticisation of inequality. For example, consider the commonplace sexualisation of teenage celebrities (Reist, 2008), or the widespread popularity of pornography that infantilises women (Turton-Turner, 2013). When coupled with the fact that actual sexual abuse of children is depressingly common (as we explore in the next paragraph), it does not seem implausible to claim that, while not explicit, child sexualisation may not be as socially objected to as we may wish to imagine.¹ As such, virtual actions that promote social norms about the sexualisation of children may be seen as inherently disrespectful and directly harming *children* (contra Patridge), thereby making such virtual actions morally impermissible.

Building on this point, we argue that rather than the impermissibility of virtual sexual violence done to children

being the morally relevant intuition that must be vindicated to provide a resolution to the dilemma, it might be more accurate to vindicate the broader intuition that virtual sexual violence per se is morally impermissible. This is because it is unclear that virtual violence against children, be it sexual or not, is the morally relevant intuitive feature at play. Consider for example the prevalence of child murder in popular videogames such as *Bioshock* (2007) and *Crusader Kings 2* (2012), and the scarcity of games that involve adult sexual assault, besides highly controversial and widely banned cases such as *RapeLay* (2006). This suggests that rather than harm done to women being a causally disconnected justification for our intuitions regarding the moral impermissibility of virtual child molestation, it could be that the intuitive concerns with virtual child molestation extend to (at least many) cases of virtual sexual violence (against adults or children).

Here we should note the massive difference in many societies between the prevalence of sexual assault and the prevalence of murder, and the accompanying difference in social attitudes towards each type of act. For example, in Australia in 2020 there were 396 victims of homicide of whom 67% were male, whereas there were 27,505 victims of sexual assault of whom 84% were female (ABS, 2020). These numbers, however, under report the prevalence of sexual assault, with a 2016 Australian Bureau of Statistics survey suggesting that 200,000 Australian adults had experienced a sexual assault in the previous 12 months, and that almost 2 million Australians had experienced at least one sexual assault since the age of 15 (ABS, 2020). The numbers of sexual assaults against children are similar, with one estimate claiming that about 1.4 million Australian adults, mainly women, have experienced sexual abuse *before* the age of 15 (AIHW, 2020), and 46% of sexual assault victims in Australia in 2020 were under the age of 15 (ABS, 2020). Given the scale of the problem, it is not surprising to find troubling attitudes expressed towards sexual assault, such as the 42% of Australians who agree with the claim that "it was common for sexual assault accusations to be used [by women] as a way of getting back at men" (AIHW, 2020). Similarly troubling attitudes do not seem to exist toward murder (or, at least, there is no evidence of this). As such, there seems to exist a clear difference between the relative prevalence of, and social attitudes towards, sexual assault (of both adults and children) and murder, and the potential for enacting virtual sexual assault to reinforce and normalise these already prevalent negative attitudes thus seems far greater than in the case of virtual murder. However, we lack empirical evidence about the actual impact of enacting virtual actions on these troubling attitudes, and thus the force of this resolution remains unclear, although it clearly warrants further empirical investigation.

Another noteworthy strategy in the literature which involves both intrinsic and instrumental elements is to focus

¹ A more fine-grained distinction between paedophilia, hebephilia and ephebophilia would be useful here in relation to the social norms around 'child' sexualisation. Clinically, paedophilia is an attraction to prepubescent children (over 5), hebephilia is an attraction to children that are going through puberty, and ephebophilia, while not a clinical category, indicates those with an attraction to children that have gone through puberty (see Blanchard et al., 2009). The above point is suggesting that ephebophilia, rather than pedohebephilia (paedophilia and hebephilia), is socially accepted in popular media (see Reist, 2008).

on the harm caused to one's own moral character by engaging and taking delight in, as well as fantasising about, virtual child molestation (see Ostritsch, 2017; McCormack, 2001; Bartel, 2020). But to count as a resolution to the Gamer's Dilemma one must not only show that enacting virtual child molestation is morally vicious, but that it harms one's moral character *more* than virtual murder to the extent that it justifies the moral impermissibility of the former and not the latter. Once again, this approach hangs on similar empirical claims, which have not been verified, to those made by the instrumental approach considered earlier regarding the differing impacts on a gamer's character by enacting various virtual actions.

Further, any such claims here are complicated by the fact that interpreting the meaning of actions in fictional spaces, as well as the motivational basis of those who engage them, is no simple task (see Gribble, 1983, p. 16; Koster, 2005, p. 84; Nader, 2020, p. 240; Young, 2013, 2020; Woodcock, 2013, p. 9; Carroll, 1990; Feagin, 1983). Young (2020), for example, distinguishes between "idle" and "surrogate" immoral fantasies. Idle fantasies are those enjoyed independently of a desire to engage in that actual immoral activity and which therefore may not warrant moral criticism, whereas surrogate fantasies are those that are enjoyed in virtue of their satisfying a desire to actually enact that immoral activity and which therefore may warrant moral criticism (but see Bartel & Cremaldi, 2018 for a critique of the claim that idle fantasies *necessarily* lack relevant desires). Drawing on this distinction, it is intuitively plausible to suppose that, on motivational grounds, virtual child molestation seems to be morally distinct from virtual murder because the kind of player who would actively seek out such an activity, as opposed to being pushed to do so by the narrative constraints of the game, is more likely to be motivated to do so for immoral reasons, such as satisfying their vicious surrogate fantasies of actual child molestation. However, any such claims about gamer motivations must be done in the context of what Ali (2015) calls the "appropriate engagement view", that is, where the gamer engages with the game as intended by the developer. This requires not only that the gamer satisfies vicious surrogate fantasies through playing a game, but that this is the experience intended to be provoked by the developer when the game is appropriately engaged with, and clearly many games will not satisfy this condition.

Somewhat easing these interpretive concerns, but not avoiding them altogether, involves a recent resolving strategy by Luck (2022). Luck appeals to the harm done to one's own moral character, but broadens the range of cases in which the Gamer's Dilemma would apply by weakening the relevant moral property that distinguishes virtual murder from virtual child molestation based on the graveness of the latter. Luck argues that when either virtual action is treated *lightly* (and, while not necessary to Luck's argument, let's

suppose "appropriately engaged" with), virtual murder will be morally permissible and virtual child molestation morally impermissible because the latter has the additional morally relevant property of being sufficiently grave, and therefore morally 'off-limits' in being treated lightly. Luck considers treating something lightly as "acting with an insouciant, carefree or light-hearted attitude towards something" (Luck, 2022, p. 12). Luck does not commit himself to any one factor in grounding what makes a virtual action such as virtual child molestation sufficiently grave to be an 'off-limits wrongdoing', or as to what makes a virtual action like virtual murder a 'fair-game wrongdoing', but instead appeals to a range of plausible intrinsic and extrinsic factors concerning the virtual actions and their actual counterparts (Luck, 2022, p. 19). Luck's approach captures many of the morally relevant contextual factors that will be discussed in the next section. However, as we will argue in "[Is the remaining intuitive difference really a moral difference?](#)" section, it is unlikely that what provokes the intuitions that virtual child molestation is an 'off-limits' wrongdoing and virtual murder is a 'fair-game' wrongdoing are grounded on morally normative foundations rather than provoked by non-moral factors.

Drawing this section together, we can now ask how an instrumental or intrinsic resolving strategy can help to narrow the Gamer's Dilemma. Even while accepting that some proposed *complete* resolutions may lack sufficient empirical support at present, we agree that some resolving strategies discussed above may in fact justify a normative distinction between virtual murder and virtual child molestation *in certain specific instances*. For example, in the case of a virtual action that is *unambiguously* an instance of virtual child pornography in which surrogate fantasies of child molestation are intentionally encouraged by the game, it could be argued that because of its eroticising the inequality of unequal groups, be it children or women, a moral distinction in terms of permissibility is identifiable. The Gamer's Dilemma is therefore resolved *in this case*. However, while we may accept that the dilemma is resolved in *these cases*, it is not clear that *all cases* of intuitively objectionable virtual child molestation will be morally objectionable on identical grounds. If such cases can be identified, then a narrow version of the Gamer's Dilemma can be coherently maintained. But to see whether such cases exist, we first need to explore the dissolving strategies found in the literature.

Dissolving strategies: adding context

The resolving strategies discussed above presuppose that gamers have (in at least some cases) the justified intuitions that 'virtual murder is morally permissible' and 'virtual child molestation is morally impermissible'. Proposed

dissolutions to the general form of the Gamer's Dilemma challenge this presupposition either by rejecting the claim that people really have those intuitions or by rejecting the claim the intuitions they do have are morally justified.

As explored in our literature review section, Ali (2015, p. 274) attempts to dissolve the Gamer's Dilemma by arguing that, when a gamer appropriately engages with either virtual murder or virtual child molestation in a *contextually equal* setting, the intuitions that preface the Gamer's Dilemma fail. However, while Ali accepts his analysis is not exhaustive and any given game may involve several contextual dimensions, we argue (as others have, e.g., Luck, 2018) that Ali's distinction between "simulation" and "story-telling" games is too blunt. While some games are more-or-less pure "simulation" games with no narrative components and others are pure "story-telling" or visual novel games with few game mechanics, most games have elements of both, with the narrative elements giving meaning to the game mechanics and affordances (see Formosa et al., 2016). This suggests that we need to focus on, not just the general "type" of game in which virtual actions occur, but the full range of relevant narrative and game design features that impact on the experiences of gamers. An approach along these lines has recently been defended by Ramirez (2020).

Ramirez is concerned with virtual actions which generate virtually real experiences. A virtually real experience is one in which the gamer's decision to enact a virtual action has similar (even if to a different degree) physiological, psychological, and behavioural impacts as the decision to enact its real-world counterpart (Ramirez, 2020, p. 147). Ramirez argues that a virtually real experience is necessarily and sufficiently constituted through structural design elements or "perspective fidelity" and in-game features or "context-realism". Perspective fidelity denotes "the degree to which the perspective provided by a simulation represents (i.e., coheres with) the perspective of the subject using the simulation" (Ramirez, 2020, p. 148). Perspective fidelity tracks onto the immersive nature of the structural elements of the simulation's design to the extent that a simulation appears as though it is actual. Structural design elements that contribute to the degree of perspective fidelity of a simulation are both physical and simulative. Physical elements may include the mechanisms for interacting with the simulation, such as a controller or headset, as well as the hardware features of the mechanism, such as haptic feedback, the screen's refresh rate, or surround sound audio. Simulative elements may include the perspective of gamers towards their actions, such as whether they are situated behind their avatar in a third-person view, which detracts from perspective fidelity, or see through their avatar's eyes in a first-person view, which enhances perspective fidelity (Ramirez, 2020, p. 148). Other features associated with the style of the game, such as the use of diegetic sound, voiceovers, and cutscenes also

add or detract from the perspective fidelity of the simulation. Context-realism captures "the degree to which the rules of a simulated universe cohere with the rules that a subject believes the actual world is governed by" (Ramirez, 2020, p. 149). For example, an NPC generating loot when killed detracts from the context realism of the gamer's experience. Further, the time and place in which the virtual environment is set can add or detract from the context-realism of the simulation. For example, *Assassin's Creed: Origins* (2017), which is set in ancient Egypt, has less contextual realism compared to a game such as *Manhunt 2* (2007), which is set in a contemporary context. For Ramirez, it is because virtual actions, including virtual murder and virtual child molestation, can generate a virtually real experience for the gamer that they are morally significant.

To demonstrate how perspective fidelity and context-realism matter, we introduce the hypothetical *Medieval Game*, which we use to show that *in some games*, contra to the way the Gamer's Dilemma is presented, virtual child molestation can be intuitively morally permissible and virtual murder can be intuitively impermissible in the same context (and vice versa). In the hypothetical *Medieval Game*, imagine that the gamer plays as a 20-year-old prince. The prince is to be married to a 14-year-old princess from a neighbouring principality to cement a political alliance, but to do that the prince must consummate the marriage. Imagine that the prince is in a room with the King's councillors, and he is told that he must either consummate the marriage that night with (let's assume) the princess's consent and thereby prevent war by cementing the alliance, or end the marriage that night by brutally and sadistically murdering the princess and thereby starting a bloody war. Further, imagine that the first choice is made through a simple dialogue option, and after the choice is made the game immediately cuts to a scene the next day where the political alliance has been formally secured. Nothing is seen of what occurs in between. In contrast, imagine that if the second choice is made, then the gamer faces a very graphic, brutal, and lengthy scene that requires them to continue to push multiple buttons to slowly torture their bride amidst her screams while she begs for the prince to stop.

In this game, the player is faced with choosing to engage in virtual child molestation, since sexual intercourse with the underage princess would to us (if not to the prince in the game) constitute child molestation due to the princess's age, or to engage in virtual murder of the same princess. As presented in the game, the former choice (child molestation) constitutes an instance of low perspective fidelity that does not produce a virtually real experience, whereas the latter choice (murder) constitutes an instance of high perspective fidelity that produces a virtually real experience. Here we have an identical game and narrative context, but with different degrees of perspectival fidelity and context realism. But many people would, we assume, find enacting

the first choice as presented in the game to be morally permissible, and many would find enacting the second choice in the game to be morally questionable at best and perhaps impermissible. In any case, many would find enacting the former virtual choice morally preferable to the latter. If so (and all such claims require empirical investigation), then *in some contexts* virtual child molestation might be intuitively permissible and virtual murder might be intuitively impermissible.² These constitute an opposite set of intuitions to the ones that ground the Gamer's Dilemma, which reinforces our key point that at best a narrow Gamer's Dilemma holds across some but certainly not all cases.

But by changing the specificities of the scenario and switching around the perspective fidelity, we could easily promote alternative intuitions. For example, imagine that the bride is now unwilling, and the choice is between a violent, sadistic, and brutal rape, which (as with the previous murder example) is shown in very graphic detail and requires the gamer to push multiple buttons to continue to assault the princess amidst loud screams from her to stop. In contrast, the murder is done off screen, silently and swiftly. Further, by choosing to consummate the marriage, the prince will anger the princess's family who no longer wish to cement the alliance, and this will lead to a long and bloody war, whereas the murder (let's assume) will allow both parties to walk away from the alliance without further bloodshed. In this alternative version of the game, the player is again faced with choosing to enact virtual child molestation or virtual murder. However, we imagine that many people would have the alternative intuition that in this case the virtual child molestation (with its high perspective fidelity that generates a virtually real experience) is impermissible, and the virtual murder (with its low perspective fidelity and failure to generate a virtually real experience) is morally permissible. If so, then it looks like *in some contexts* virtual child molestation might be intuitively impermissible and virtual murder might be intuitively permissible, which is in line with the grounding intuitions of the Gamer's Dilemma.

The *Medieval Game* thus demonstrates that the intuitions that ground the Gamer's Dilemma are shaped heavily by contextual considerations. However, we stop short here of accepting the normative component of Ali's or Ramirez's dissolutions whereby player agency or virtually real experiences respectively are the *sole* morally determining feature at play in the Gamer's Dilemma *in all cases*. While it is certainly plausible that the degree of player agency or whether a virtual action causes a virtually real experience

are contextually relevant features that may *sometimes* ground a moral distinction between virtual murder and virtual child molestation, it is unclear that these features will be able to do so in *all* cases. This is because there could be instances which have limited player agency or fall short of generating virtually real experiences, but which still justifiably trigger the intuition in gamers that, in such cases, the enactment of virtual child molestation is impermissible whereas the enactment of virtual murder is not. We attempt to show that such a range of cases exist in the next section.

Drawing this section together, we have argued here in support of the claim that the intuitive permissibility of virtual murder and virtual child molestation depend on a host of contextual features. Further, we suggest that considering many contextual features, rather than any one in isolation, is important since several contextual features (including game type, context realism, and perspective fidelity) are all simultaneously at play when virtual murder or virtual child molestation are enacted, albeit to varying degrees. However, we have resisted thus far accepting that focusing on context *completely* dissolves the Gamer's Dilemma, since it is possible, as we show in the next section, that there may exist a *narrow* version of the Dilemma with robust intuitive foundations (although we question below whether those foundations are genuinely moral or not).

The narrow Gamer's Dilemma: *how low will you go?*

While we have shown that the intuitive permissibility of both virtual murder and virtual child molestation can depend upon a range of contextual factors, we show in this section that, intuitively, virtual murder seems to be morally permissible *in a wider range of cases* than virtual child molestation (where the contextual features are sufficiently similar). This leaves us with what appears to be a narrow Gamer's Dilemma. To make this point, we shall consider a game inspired by Luck's (2018) hypothetical version of *Counter Strike*, which we call *How Low Will You Go?*. Whereas in the *Medieval Game* we imagined *contrasting* contextual factors, with virtual child molestation being positioned with comparatively low levels of perspective fidelity compared to virtual murder and vice versa, we now focus on varying *equivalent* degrees of contextual factors. In *How Low Will You Go?* there are two game modes, 'murder mode', which involves enacting virtual murder, and 'molestation mode', which involves enacting virtual child molestation. Both modes are always equally positioned contextually, or as closely as possible, to the other (i.e., various contextual features, including degrees of player agency, perspective fidelity, context-realism and so on, are all kept the same for both modes). As you progress through the game, say from

² Since context matters on our view, it is not surprising that we can change these intuitions by changing the context. For example, people's intuitions about this case might change if the princess's age was five years old.

level one to level ten, the contextual scales are dialled up for both modes at the same time. On Ali's in-game context scale, the game type moves from being an unobjectionable storytelling game with almost no player agency, to a somewhat objectionable storytelling game with limited agency, and eventually becoming a simulation game granting gamers almost unlimited agency. On Ramirez's scale of simulation design, the virtual actions move from being positioned in a highly pixelated two-dimensional arcade style game, with eight-bit music and low overall perspective fidelity, through a series of gradual intermediary steps, to a highly simulative game with no non-diegetic sound or game mechanics that undermine its highly detailed context-realism. This involves a gradual shift from a game that does not produce virtually real experiences for the gamer to one that, at higher levels, does so to a very high degree.

On the first few levels of *How Low Will You Go?*, both murder and molestation modes would seem to be intuitively morally permissible because all the relevant contextual scales are shifted down to a very low level. Here we have a storytelling game mode with almost no player agency and simulative design features that do not induce a virtually real experience. The story could be initially told from a third-person viewpoint and the game could make it clear that, while the actions depicted are morally objectionable, they are necessary to complete a meaningful narrative arc. Playing such a simple visual novel game would be akin to reading a book such as Nabokov's *Lolita* (for molestation mode) or Puzzo's *The Godfather* (for murder mode). If readings such books (or making both low perspective fidelity choices in the *Medieval Game* outlined above) is morally permissible, then presumably playing the first level (or first few levels) of *How Low Will You Go?* in both murder and molestation modes should be intuitively permissible too.

Gradually, as you progress through the game from level to level, various contextual factors are slowly dialled up, with the game shifting to a first-person view, and the game in both modes becoming closer and closer to constituting a highly immersive virtually real experience with enormous scope for player agency with minimal storytelling. Gamers would likely be intuitively comfortable enacting virtual murder as the dials on each of these contextual scales are slowly increased for each level. But at some point, say by level 7 or 8 (if we imagine a 10-point scale), when enacting virtual murder involves engaging in a highly realistic murder and torture simulator which, to the player, almost feels like *really* murdering and torturing someone, gamers will likely find continuing to play the game to be intuitively morally objectionable. Whatever the level at which players find murder mode morally impermissible, it seems that players would likely find virtual child molestation mode to be morally impermissible *at an earlier level* than virtual murder. Since players will find molestation mode morally

objectionable *earlier* than they will find murder mode morally objectionable, it follows that players would not find the virtual actions to become morally impermissible *at the same contextual level*.

Thus, in *How Low Will You Go?* at very high levels both modes are intuitively impermissible, at very low levels both modes are intuitively permissible, and at some intermediary levels murder mode is permissible whereas molestation mode is impermissible. It is not necessary to specify exactly where these intuitive shifts occur. All that needs to be intuitively plausible is that there are at least some intermediary levels of the game (say levels 4 to 6), where the contextual features are sufficiently similar, but it is intuitively permissible to play murder mode and intuitively impermissible to play molestation mode. It is *only* at these intermediary levels that the narrow Gamer's Dilemma seems to hold as we have an accurate articulation of the intuitions of gamers that apply to a narrow range of overlapping specific cases that are not, as Ali's and Ramirez's dissolutions require, the result of mistakenly underappreciating the context in which the relevant virtual actions are situated.

At this point, we can seek to resolve this narrow Gamer's Dilemma by rejecting **N3** through showing that gamers can defend the claim that *in those specific cases* there is a relevant moral difference that justifies both **N1** and **N2**. We will address and reject the plausibility of this approach in the remainder of this section. Alternatively, we can seek to resist the narrow Dilemma by accepting that the intuitions of **N1** and **N2** are indeed held by gamers, but that the *content* of those intuitions are not *justified on moral grounds*. We address and endorse this second approach in the next section.

By specifying a clear set of cases to ground the narrow form of the Gamer's Dilemma, we can significantly advance our ability to assess proposed resolutions to the narrow Dilemma as we are forced to be explicit about exactly what sort of contextual features we are thinking about when assessing proposed resolutions. While we will not rehash all the proposed resolutions explored above, many of which we found to be promising but lacking sufficient empirical support at present, we can make our point by reconsidering one, and arguably the strongest, candidate resolution. To this end, we return briefly to the candidate resolution that we developed above that builds off the work of Levy (2002), Bartel (2012) and Patridge (2011, 2013) that involves the eroticisation of sexual assault of socially unequal groups.

What is essential to this proposed resolution is that because of the incorrigible social norms that exist in society that promote the sexual inequality of women and, arguably, children, it is wrong to "appropriately engage" in a virtual action that depicts child molestation when that act is intended to be erotic for the gamer. The erotic nature of virtual child molestation for appropriately engaged gamers

is clearly realised, for example, when ‘virtually-generated’ depictions of child molestation are treated as pornographic in much the same way as ‘actually-produced’ depictions of child molestation is by those who engage as intended with it. In both cases, they share the feature of being actual instances of child pornography (see Bartel, 2012; Rea, 2011, p. 134). This is because virtual child molestation *qua* child pornography lacks, building off Patridge (2011), the interpretive flexibility to be regarded as *not* eroticising the inequality of unequal groups when appropriately engaged with.

When this resolution is applied to our narrow formulation of the Dilemma, it is apparent that virtual child molestation *qua* child pornography is only likely to be realised at higher levels of the game (where the narrow Gamer’s Dilemma does not hold), insofar as it is only at these higher levels that the game will feature a degree of context-realism that is (arguably) *justifiably* intuited as *morally* objectionable virtual child molestation that is akin to immoral visual child exploitation imagery. But at these comparatively high levels of the game, virtual murder would also likely be intuited, and perhaps justifiably, to be morally objectionable, insofar as it is akin to snuff imagery or an unambiguous endorsement of actual murder, which is why the dilemma does not hold at this level.³

However, if we instead focus on lower intermediary levels of the game where the narrow Gamer’s Dilemma is argued above to hold (i.e., where murder but not molestation mode is intuited as permissible), any erotic features of virtual child molestation at this level will primarily be produced by narrative elements, or highly abstract rather than realistic visual imagery. Importantly, many gamers may still find it intuitively wrong to play a game with explicit erotic narrative elements involving children, even if the game at this level lacks the visual and narrative components needed to constitute actual child pornography.⁴ But a resolution requires that this specific intuition is *morally justified* on the grounds that the relevant virtual actions in molestation mode at this lower intermediary level lack the interpretive flexibility to be understood as not immorally eroticising the inequality of unequal groups. However, this need not be the case. Consider again the moral distinction between idle and surrogate fantasies outlined above, and the interpretive flexibility available to the former in avoiding the same kind of moral criticism as the latter. Drawing on this distinction, for weaker instances of virtual child molestation, such as video game narratives akin to Nabokov’s *Lolita* with no realistic pornographic visual imagery, it seems wrong that a proper

engagement with such a video game narrative *necessarily* involves eroticising inequality or fuelling immoral surrogate fantasies when engaged with appropriately, even if it still contains erotic elements that many gamers would find intuitively impermissible to enact in the game. As such, appealing to the immorality of eroticising inequality or fuelling obscene surrogate fantasies fails to resolve the narrow Gamer’s Dilemma for lower intermediary levels of *How Low Will You Go?*. This proposed resolution is thus unable to adequately track onto the *specific* intuitions that ground the narrow form of the Gamer’s Dilemma as outlined here, even if it applies in other cases.

Is the remaining intuitive difference really a moral difference?

We now need to consider whether the remaining intuitive difference between virtual murder and virtual child molestation in the intermediary levels of *How Low Will You Go?*, and other similar cases, is really a *moral* intuition at all and, if it is, whether it is a *justified* one (i.e., of the sort that can withstand reflective equilibrium). While gamers may indeed intuit a difference between the permissibility of virtual child molestation and virtual murder in a narrow range of cases, as is captured by Luck’s distinction between ‘off-limits’ and ‘fair-game’ wrongdoings (Luck, 2022, p. 13), it could be some other *non-moral* feature that is underwriting that intuitive difference for them, such as a conventional or taste-based aesthetic norm or an implicit psychological attitude, which they are confusing with a moral basis (see Young, 2017a, 2017b, 2019; Kjeldgaard-Christiansen, 2020). Such factors, while normatively forceful, are not necessarily morally grounded (see, e.g., Young, 2017b, p. 233 on suberogatory, i.e., blameworthy but not morally impermissible, actions in videogames).

Consider, for example, the way that virtual murder (i.e., wrongful intentional killing) is positioned as an action type within the broader, equivalently sanctioned, class of actions that depict killing in videogames generally (Young et al., 2011; Davnall, 2020). This distinction draws on the different kinds of killing made possible in videogames, from justified killing in warfare games, such as in *Call of Duty* (2003), to unjustified intentional killing, or murder, in games like *Manhunt 2* (2007) (Luck, 2009, p. 32). We suggest that the capacity for videogame context to structure the gamer’s relationship to killing ‘bad guys’ or ‘enemies’ as morally permissible conventionally justifies the permissibility of virtual killing per se and renders instances of virtual murder interpretively flexible (c.f. Patridge, 2011). As a result, the conventional permissibility of justified killing in games may bleed into, and cause, gamers to construe many, if not most, instances of unjustified killing as being (at least

³ While we take this claim to be intuitive, our proposed narrow dilemma requires empirical support.

⁴ This point rests on a subtle distinction between ‘erotic entertainment’ and ‘pornography’; for discussion (see Levinson, 2005).

potentially) conventionally permissible. After all, is there really such a big intuitive difference, for gamers, between running over and killing a 'bad guy' NPC (non-player character), or an NPC that might be a bad guy, or an unknown but innocent pedestrian NPC in *GTA V* (2013)? Further, NPCs will typically fight back (or sometimes attack first) in many games, bringing in issues of self-defence that can also help to muddy the intuitive waters when it comes to murder in games.

Virtual child molestation, on the other hand, does not have a range of permissible, equivalently sanctioned action types in the class of sexual assault actions to draw on to justify, or at least make interpretively flexible, its intuitive conventional permissibility as an action (Young et al., 2011; Bartel, 2020; Patridge, 2011; Young, 2013, 2016, 2020). While this is surely, in part, a result of the *necessarily* unjustified nature of child molestation as an action type, it is also a result of killing, as a class of action, being deeply tied to permissible conventions of taste in videogames within the popular gaming sub-genre of 'shooters'. This conventional permissibility has developed from fixed and multi-directional shooter games like *Space Invaders* (1978) and *Asteroids* (1979), to first-person and third person shooter games such as *Doom* (1993) and *Max Payne* (2001) and beyond. Therefore, we contend that without killing being a conventionally permissible dimension of videogame actions, virtual unjustified killing, or virtual murder, would not be intuited as morally permissible *as readily* as it is by gamers. The intuitions that ground the narrow Gamer's Dilemma must therefore be situated in the context of the development of conventional norms of appropriate videogame play that have developed over time around killing (or murdering) in games, and the very different history of conventional norms that have emerged around the enactment of sexual violence and child molestation in videogames.

This story seems to offer an alternative explanation as to the normative force behind the intuitions grounding the remaining narrow Gamer's Dilemma. It suggests that the robustness of the intuitive difference may indeed have a normative basis, but that basis may lie in aesthetic and conventional norms. Further, there may be other plausible alternative normative foundations we could also appeal to, to explain this intuitive difference. For example, it might be that though it is not immoral to play a game with low context-realism narrative elements of child molestation (i.e., intermediary levels of *How Low Will You Go?*), it is still in poor taste, or not something that one is comfortable doing, or something that conflicts with the non-moral ideals of the sort of person one wants to be. After all, there are many things we may be uncomfortable or unwilling to do, such as eat insects, go sky diving, craft grammatically poor sentences, or (more relevantly here) play narratives involving child molestation, without thinking that those things are

immoral. This allows us to say that there is indeed a relevant normative difference between virtual murder and virtual child molestation in certain contexts such that the former but not the latter is appropriate behaviour, but without holding that the relevant normative difference here is a moral one.

These alternative accounts are not only clearly plausible (although in need of further elaboration and defence than we can provide here), but they also help to explain why it has proven so difficult to find a *moral* resolution to the narrow version of the Gamer's Dilemma, namely because the intuitive difference we seem to see between virtual child molestation and virtual murder in such cases may have a *non-moral* normative foundation. Further, many of the proposed resolutions explored above were shown to lack concrete empirical support, and it seems unlikely that our moral intuitions are attempting to track such fine-grained empirical facts or could plausibly be expected to do so reliably such that we should justifiably continue to rely on them. Of course, while these alternative explanations do not by themselves prove that the remaining intuitive difference in the narrow form of the Dilemma identified here cannot *possibly* be morally grounded (see Kjeldgaard-Christiansen, 2020, p. 100), they do have sufficient argumentative force to suggest that future philosophical and empirical work should pursue these alternative explanations rather than continue to seek resolutions to the husk of what remains of the narrow Gamer's Dilemma.

To summarise the novel contributions of this paper, we have argued for an approach to limit the gamer's dilemma which substantially differs from the literature in three important respects. First, other limiting approaches in the literature each focus on a particular contextual dimension—such as virtually real experiences (Ramirez, 2020) and game type/player agency (Ali, 2015; Bourne, 2019) (see also Öhman, (2020) on degrees of abstraction, and Luck (2022) on other potentially morally relevant contextual factors)—but we argue that when *all morally relevant* contextual dimensions are kept constant, there will *still* be an intuited difference in permissibility between action types, as shown by the intermediary levels of the *How Low Will You Go?* example. Second, we argue that these approaches to limiting the Dilemma, which all tend to focus on the contextual extremes to dissolve the Dilemma, fail because it is not in the contextual extremes that the intuitive moral distinction of the Dilemma is triggered, but rather in cases which fall short of what these writers think is ethically pertinent (such as Ramirez's 'virtually real experience') as again shown by the intermediary levels of our *How Low Will You Go?* example. Finally, our own suggested contextually relevant dimensions, which are non-moral, are a substantial departure from the literature as we demarcate our moral intuitions from their justificatory grounds. We do so by arguing that while our *intuitions* about the Gamer's Dilemma may be concerned

with *moral permissibility*, it is unclear that there is any moral basis to justify these intuitions, and we instead point to several non-moral foundations that may explain the intuitive differences between murder and molestation in video games. No other approach in the literature explicitly argues for this view (while Kjeldgaard-Christiansen (2020) alludes to the idea of our moral intuitions having non-moral foundations, they ultimately reject this view), especially one that is also in-keeping with our precise suggestion of how the Dilemma ought to be limited.

Conclusion

The practical significance of the Gamer's Dilemma comes from it forcing us either to accept that many of the most popular videogames, which involve murder, are immoral to play, or that child molestation and sexual assault simulators should be morally accepted. In this paper, we have argued that the *general form* of the Gamer's Dilemma can be dissolved, since sometimes virtual murder is intuitively impermissible and sometimes virtual child molestation is intuitively permissible. This claim depends on the relevance of a range of specific features of the gaming context, such as the type of game, and its perspective fidelity and context-realism. However, we also argue that a *narrow form* of the Gamer's Dilemma seems to survive because, focusing on the intermediary levels of our *How Low Will You Go?* example, there is a *narrow range of cases* where, holding the perspective fidelity and other contextual factors constant, gamers intuit virtual murder to be permissible and virtual child molestation to be impermissible, and various proposed resolutions do not seem to be able to identify a relevant moral difference *in those specific cases*. Further, the inability to locate such moral resolutions leads us to suggest that while this intuitive difference may have a robust normative foundation, that normative foundation seems more likely to be grounded in matters of taste, suberogatory action, narrative tropes, or aesthetic conventionality than in morality. Our argument thus helps us both to sharpen our intuitions through a focus on specific contextually identical cases of virtual murder and molestation, and encourages the exploration of attempts to provide non-moral foundations, such as narrative tropes, to explain any remaining intuitive difference we find between such cases. We also suggest several avenues of future empirical work that would help to further the debate.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

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(4), 267–274.
- Anderson, C., et al. (2010). Violent video game effects on aggression, empathy, and prosocial behavior in eastern and western countries. *Psychological Bulletin*, 136(2), 151–173.
- American Psychological Association. (2015). APA review confirms link between playing violent video games and aggression. Retrieved from <https://www.apa.org/news/>.
- Australian Bureau of Statistics (ABS). (2020). Recorded Crime—Victims, 2020. Retrieved June 24, 2021, from <https://www.abs.gov.au/statistics/people/crime-and-justice/recorded-crime-victims/latest-release>.
- Australian Institute of Health and Welfare (AIHW). (2020). “Sexual Assault in Australia”. Retrieved June 24, 2021, from <https://www.aihw.gov.au/getmedia/0375553f-0395-46cc-9574-d54c74fa601a/aihw-fdv-5.pdf>.
- Bartel, C. (2012). Resolving the Gamer's Dilemma. *Ethics and Information Technology*, 14(1), 11–16.
- Bartel, C. (2020). *Video games, violence, and the ethics of fantasy: Killing time*. New York: Bloomsbury Publishing.
- Bartel, C., & Cremaldi, A. (2018). ‘It's just a story’: Pornography, desire, and the ethics of fictive imagining. *British Journal of Aesthetics*, 58(1), 37–50.
- Blanchard, R., Lykins, A. D., Wherrett, D., et al. (2009). Pedophilia, Hebephilia, and the DSM-V. *Archives of Sexual Behaviour*, 38, 335–350.
- Bourne, E. (2019). Players, characters, and the gamer's dilemma. *Journal of Aesthetics and Art Criticism*, 77(2), 133–143.
- Carroll, N. (1990). *The philosophy of horror*. Routledge.
- Cunningham, A. S., Engelstätter, B., & Ward, M. R. (2011). Understanding the effects of violent video games on violent crime. 11-042. *ZEW Discussion Papers*. Leibniz Centre for European Economic Research.
- Davnull, R. (2020). What does the gamer do? *Ethics and Information Technology*, 23(3), 225–237.
- Feagin, S. L. (1983). The pleasures of tragedy. *American Philosophical Quarterly*, 20(1), 95–104.
- Ferguson, C. J., Rueda, S. M., Cruz, A. M., Ferguson, D. E., Fritz, S., & Smith, S. M. (2008). Violent video games and aggression. *Criminal Justice and Behavior*, 35(3), 311–332.
- Formosa, P., Ryan, M., & Staines, D. (2016). Papers, please and the systemic approach to engaging ethical expertise in videogames. *Ethics and Information Technology*, 18(3), 211–225.
- Gribble, J. (1983). *Literary education*. Cambridge University Press.
- Haynes, J. (2015). 10 Most Violent Video Games of 2015 (and What to Play Instead). *Common Sense Media*.
- Huizinga, J. (1938). *Homo Ludens*. Angelico Press.

- Juul, J. (2008). The magic circle and the puzzle piece. *Conference Proceedings of the Philosophy of Computer Games*.
- Koster, R. (2005). *A theory of fun for game design*. Paraglyph Press.
- Kjeldgaard-Christiansen, J. (2020). Splintering the Gamer's Dilemma. *Ethics and Information Technology*, 22(1), 93–102.
- Levy, N. (2002). Virtual child pornography. *Ethics and Information Technology*, 4(4), 319–323.
- Levinson, J. (2005). Erotic art and pornographic pictures. *Philosophy and Literature*, 29(1), 228–240.
- Luck, M. (2009). The Gamer's Dilemma. *Ethics and Information Technology*, 11, 31–36.
- Luck, M. (2018). Has Ali dissolved the Gamer's Dilemma? *Ethics and Information Technology*, 20(3), 157–162.
- Luck, M. (2022). The grave resolution to the Gamer's Dilemma: An argument for a moral distinction between virtual murder and virtual child molestation. *Philosophia*: 1–22.
- 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(4), 277–287.
- McGlynn, C., & Rackley, E. (2009). Criminalising extreme pornography. *Criminal Law Review*, 4, 245–260.
- Nader, K. (2020). Virtual competitions and the Gamer's Dilemma. *Ethics and Information Technology*, 22(3), 239–245.
- Öhman, C. (2020). Introducing the Pervert's Dilemma. *Ethics and Information Technology*, 22(2), 133–140.
- Ostritsch, S. (2017). The amoralist challenge to gaming and the Gamer's moral obligation. *Ethics and Information Technology*, 19(2), 117–128.
- Patridge, S. (2011). The incorrigible social meaning of video game imagery. *Ethics and Information Technology*, 13(4), 303–312.
- Patridge, S. (2013). Pornography, ethics, and video games. *Ethics and Information Technology*, 15(1), 25–34.
- Prigg, M. (2009). Violent video game breaks records with 4.7m sales in a day. *London Evening Standard*. Retrieved 24 February 2016 from <http://www.standard.co.uk/news/violent-video-game-breaks-records-with-47m-sales-in-a-day-6705082.html>.
- Ramirez, E. J. (2020). How to (Dis)Solve the Gamer's Dilemma. *Ethical Theory and Moral Practice*, 1, 141.
- Rea, M.C. (2001). What Is Pornography? *Noûs*, 35(1): 118–45. <https://doi.org/10.1111/0029-4624.00290>
- Reist, M. T. (2008). The pornification of girlhood. *Quadrant Magazine*, 52(7/8), 10–16.
- Salen, K., & Zimmerman, E. (2003). *Rules of play*. MIT Press.
- Staines, D., Formosa, P., & Ryan, M. (2019). Morality play. *Games and Culture*, 14(4), 410–429.
- Turton-Turner, P. (2013). The porn illusion. *The International Journal of the Image*, 3(3), 17–26.
- Tillson, J. (2018). Is it distinctively wrong to simulate doing wrong? *Ethics and Information Technology*, 20(3), 205–217.
- Woo, J. (2004). The concept of “harm” in computer-generated images of child pornography. *The John Marshall Journal of Information Technology & Privacy Law*, 22(4), 717.
- Woodcock, S. (2013). Horror films and the argument from reactive attitudes. *Ethical Theory and Moral Practice*, 16(2), 309–324.
- Young, G. (2013). Enacting taboos as a means to an end; but what end? *Ethics and Information Technology*, 15(1), 13–23.
- Young, G. (2016). *Resolving the Gamer's Dilemma*. Springer.
- Young, G. (2017a). Objections to Ostritsch's argument in “The Amoralist challenge to gaming and the Gamer's moral obligation.” *Ethics and Information Technology*, 19(3), 209–219.
- Young, G. (2017b). Integrating poor taste into the ongoing debate on the morality of violent video games. *The Computer Games Journal*, 6(4), 227–237.
- Young, G. (2019). An expressivist account of the difference between poor taste and immorality. *Ethical Theory and Moral Practice*, 22(2), 465–482.
- Young, G. (2020). *Fictional immorality and immoral fiction*. Lexington Books.
- Young, G., Whitty, M., & Goodings, L. (2011). What i won't do in pixels. *Computers in Human Behavior*, 27, 268–275.
- Zagal, J. P. (2009). Ethically notable videogames. In *Proceedings of DiGRA*. DePaul University.

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