palgrave▶**pivot**

RESOLVING THE GAMER'S DILEMMA

Examining the Moral and Psychological Differences between Virtual Murder and Virtual Paedophilia

Garry Young

Palgrave Studies in Cyberpsychology

Series Editor
Jens Binder
Nottingham Trent University
Nottingham
United Kingdom

Palgrave Studies in Cyberpsychology aims to foster and to chart the scope of research driven by a psychological understanding of the effects of the 'new technology' that is shaping our world after the digital revolution. The series takes an inclusive approach and considers all aspects of human behaviours and experiential states in relation to digital technologies, to the Internet, and to virtual environments. As such, Cyberpsychology reaches out to several neighbouring disciplines, from Human-Computer Interaction to Media and Communication Studies. A core question underpinning the series concerns the actual psychological novelty of new technology. To what extent do we need to expand conventional theories and models to account for cyberpsychological phenomena? At which points is the ubiquitous digitisation of our everyday lives shifting the focus of research questions and research needs? Where do we see implications for our psychological functioning that are likely to outlast shortlived fashions in technology use?

More information about this series at
http://www.springer.com/series/14636

Garry Young

Resolving the Gamer's Dilemma

Examining the Moral and Psychological Differences between Virtual Murder and Virtual Paedophilia

palgrave
macmillan

Garry Young
School of Social Sciences
Nottingham Trent University
Nottingham, Nottinghamshire, United Kingdom

Cover illustration: Abstract Bricks and Shadows © Stephen Bonk/Fotolia.co.uk

Printed on acid-free paper

CONTENTS

Introducing the Gamer's Dilemma

Abstract This chapter sets out the conditions that lead to the gamer dilemma. It begins with a brief discussion on video games that permit virtual murder and contrasts these with the fact that, presently, virtual paedophilia is not permitted. While this is said to accord with our moral intuition, a more detailed analysis reveals that arguments in favour of the permissibility of virtual murder appear to support the permissibility of virtual paedophilia, and vice versa in the case of impermissibility. The gamer is therefore faced with a dilemma: either he/she must permit virtual paedophilia alongside virtual murder or prohibit both. Current US and UK legislation regarding virtual child pornography is also discussed to help contextualize the dilemma further and inform discussion in the chapters to come.

Keywords Virtual murder · Virtual paedophilia · Child pornography legislation

1.1   Virtual Murder: The Current State of Play

Within single-player video games (hereafter, video games), it is permissible to engage in simulated murder. By murder, I mean the intentional and unlawful killing of an individual. Indeed, it is far from hyperbole to say that a large percentage of violent video games contain acts of simulated killing, many of which would be categorized as murder or as otherwise

unlawful if performed for real. To illustrate, Cunningham et al. (2011) report that from a total of 1117 video games sampled, 672 were identified as non-violent and 445 violent (based on the Entertainment Software Ratings Board's (ESRB's) ratings and content descriptors). Of the 445 violent titles, 113 were considered to be extremely or, as Cunningham et al. refer to them, 'intensely' violent. Moreover, Prigg (2009) reports that, on the first day of its release, the video game *Call of Duty: Modern Warfare 2* sold 4.7 million copies in the USA and UK alone, outselling the previous best video game – *Grand Theft Auto IV* – by some distance. Both the *Grand Theft Auto* and *Call of Duty* series are held to be extremely violent games. (Before proceeding, a point of clarification: reference to 'violent video games' should be understood as short-hand for video games whose content contains simulated violence.) *Call of Duty: Modern Warfare 2* became infamous for its airport massacre scene, and *Grand Theft Auto IV* permits the gamer's character to have sex with a prostitute before mugging or even killing her. The popularity of violent or even extremely violent content does not appear to be waning. As Haynes (2015) notes:

> In 2015, we saw some of the most violent video games ever released. Plus, older violent games such as *Gears of War: Ultimate Edition* and *Resident Evil: The Definitive Edition* were re-released with visual upgrades that intensify the more violent moments, including blood and gore splattering (p. 1).

When describing the *current state of play* (meaning those games currently available to age-appropriate persons in the UK and USA), enacting murder is not only permitted but a common occurrence; some might even say 'positively encouraged'. In *Manhunt 2*, for example, I (in the form of an avatar) can bludgeon to death a stranger with a kitchen utensil. *Postal 2* allows me to set someone on fire while they are alive, douse the flames by urinating on them, before beating them to death with my boot and a shovel. More recently, the video game *Hatred* has courted controversy through its seemingly relentless enactment of random murder (Campbell 2014). In contrast, the current state of play does not permit video games to contain enactments of paedophilia.[1] One quick and easy way to account for this discrepancy is to point out that virtual child pornography, which would include the virtual enactment of paedophilic acts, is illegal in many countries, including the UK and, with qualification, the USA.

Before discussing the legality of virtual paedophilia (both for the purpose of clarification and as a means of informing the moral debate

to come), one might ask with some incredulity: why would anyone want to do *that*? By 'that', I mean why would anyone want to play a game in which they can simulate paedophilic activity and therefore, to all intense and purposes, play at being a paedophile? The intuition underlying this question and the incredulity with which it might be asked seem to appeal to player motivation. Crudely put, one might suspect that there is something wrong with someone who wants to play at being a paedophile; that their motivation to enact paedophilia stems from the fact that it vicariously satisfies, and is therefore a symptom of, their desire to engage in actual paedophilia. Or perhaps, one fears the risk of enacting this activity within a game; that, somehow, repeatedly engaging in such simulations may lead one to acquire a taste for what the simulation represents (a kind of slippery-slope argument). Of course, some people may question the motivation of individuals who play a game like *Postal 2* in which one can enact all kinds of extremely violent acts. Returning to the earlier example, they may ask with equal incredulity why anyone would want to play a game in which it is possible to set someone on fire, urinate on them to douse the flames and then beat them to death. Is enacting this kind of activity likewise a symptom of some other desire: namely, to engage in actual murder? Although there will be dissenters, I suspect the majority response would be 'no'. It is, however, a question I will return to.

1.2   THE GAMER'S DILEMMA

Virtual murder *is* permitted in the UK and USA, even when enacted with the level of violence depicted in video games like *Postal 2* (as one example among many). Given this, consider the words of Morgan Luck when introducing the gamer's dilemma:

> Is it immoral for a player to direct his character to murder another within a computer game? The standard response to this question is no. This is because no one is actually harmed as a result of a virtual harm. Such an outlook seems intuitive, and it explains why millions of gamers feel it is perfectly permissible to commit acts of virtual murder. Yet this argument can be easily adapted to demonstrate why virtual paedophilia might also be morally permissible, as no actual children are harmed in such cases. This result is confronting, as most people feel that virtual paedophilia is not morally permissible. (Luck 2009, p. 31)

According to Luck, the dilemma gamers face – or indeed anyone faces who has a view on the *selective* prohibition of video game content (Young 2013b) – is that any appeal to rudimentary arguments avowing 'no harm', used to rebut criticism of our intuitions over the permissibility of virtual murder, can also be used to challenge any intuitions we may have about the impermissibility of virtual paedophilia. If the claim is that no actual harm occurs as the result of virtual murder then, likewise, why should it not be claimed that no actual harm results from virtual paedophilia? Given the permissibility of the former, why prohibit the latter? What justifies our contrary intuition, here? Where our intuitions are shown to be inconsistent or seemingly without support, at least after a cursory examination, the gamer (or any other interested party) is faced with a dilemma. If one wishes to achieve parity, either one prohibits virtual murder *and* virtual paedophilia (resulting in the unfortunate consequence of prohibiting an activity many gamers intuitively feel is acceptable and indeed enjoy enacting: namely, murder) or one permits each of these activities (thereby creating a different unpalatable consequence: allowing the enactment of paedophilia, which many would find repugnant). Of course, one could simply admit to having inconsistent and, *it would seem*, indefensible views about different virtual content; indefensible, that is, outside of an appeal to the popularity of certain intuitions.

Appeal to intuition is not a sage strategy, however (something we will return to in Section 2.1); a conclusion Luck himself acknowledges. Indeed, much of Luck's original paper on the gamer's dilemma sets out to examine "whether any good arguments can be produced to reconcile the intuition that virtual murder is morally permissible, with the intuition that virtual paedophilia is not" (2009, p. 31), thereby making such seemingly inconsistent intuitions defensible through evidence and/or argument. Luck concludes that there are none.

1.2.1 A Brief Overview

Since the introduction of the gamer's dilemma, a number of ways of resolving it have been suggested, and debate continues over their respective success. In what is to follow, I will consider each of these arguments in turn and present various responses to them: mainly in relation to competing or absent empirical findings (where certain findings are required to support an argument) or through the identification of internal inconsistencies and/or conceptual incoherence within the argument itself. On completing my critical review

in which, to a greater or lesser degree, I identify problems with all previous attempts at resolving the dilemma, I present my own thoughts on how we might approach finding a solution.

Chapters 2 and 3 will be taken up with the different ways in which Luck tries to resolve the dilemma, none of which he finds wholly convincing. In his original paper, some of his suggestions are given only cursory treatment, I therefore expand on the reasoning Luck uses in each case. My aim is to provide further support for the conclusions he draws and although, in places, I disagree with the manner of his argument, I nevertheless concur with his overall dissatisfaction with the suggested means of resolving the dilemma. In Chapter 4, I consider Christopher Bartel's attempted resolution (Bartel 2012). Here, I present a systematic critique of each of the premises on which he grounds his argument. I find each problematic in its way. In my appraisal, I draw on recent (i.e. 2013) responses to Bartel's paper from Stephanie Patridge and also Morgan Luck and Nathan Ellerby. In Chapter 5, I consider Patridge's reply in more detail, and offer some critical thoughts on her position. In many respects her argument is promising; although not without its problems, as I discuss. I also consider Rami Ali's work on the gamer's dilemma. Ali (2015) offers an original approach which, again, shows promise – particularly his thoughts on different contexts – but like all previous attempts is not without its problems.

With the exception of Ali, all other attempts at resolving the dilemma have accepted Luck's claim that there is a difference in our intuitions over the permissibility of virtual murder and virtual paedophilia. If we likewise accept (for now) this claim as our starting point, then what forms the basis for this difference? Are our intuitions tapping into and therefore describing some independent moral fact – in a moral realist sense – or are they indicative only of a difference in our moral *attitude* towards these respective virtual enactments: an attitude that neither describes nor derives any moral authority from putatively independent moral truths? If moral realism is true then it appears unable to inform attempts at resolving the gamer's dilemma, as I hope to show in my critical review throughout Chapters 2–5.

In Chapter 6, I therefore adopt an anti-realist approach and, in doing so, present my own thoughts on how the gamer's dilemma could be resolved. I discuss *constructive ecumenical expressivism*: a meta-ethical approach to moral utterances which I have previously applied to virtual gaming content (Young 2014, 2015b). I argue that constructive ecumenical expressivism

provides insight into what our moral intuitions amount to and therefore why there is a difference between our moral *attitude* to virtual murder compared to virtual paedophilia. Once the nature of this moral attitude is understood (in terms of the basis for its formation), differences that exist between our attitude towards different virtual content can be articulated in morally relevant terms, whether in the context of the gamer's dilemma specifically or selective prohibition more generally. Constructive ecumenical expressivism not only proffers a means of resolving the gamer's dilemma but, importantly, is robust enough to be co-opted as a normative ethic applicable to all forms of virtual gaming content.

I would like to finish this chapter by saying something about the legality of virtual child pornography, predominantly within the UK and USA. Initially, to illustrate ways in which legislation is similar or differs between these two countries, but more importantly to make the point that the focus of this book is on the *morality* of video game content irrespective of its legal status. In other words, irrespective of the legality of virtual paedophilia, what arguments are there for or against its *moral* prohibition, and are these able to differentiate between virtual paedophilia and virtual murder in a *morally* relevant way? That said, I believe that an understanding of some of the key legal arguments for and against virtual paedophilia will prove to be of use when debating the morality of certain activities within video games.

It is also worth noting that I consider a detailed examination of the different ways theorists have attempted to resolve the gamer's dilemma to be crucial to an understanding of what Whitty et al. (2011) refer to as *symbolic taboo activities* (STAs): basically, the virtual enactment of all activities deemed to be taboo (*qua* illegal and/or immoral) in the real world, such as assault, torture, rape, murder, paedophilia (including incest), bestiality, necrophilia and so on. As alluded to above, what we will learn by considering arguments for and against the selective prohibition of virtual paedophilia will, in turn, provide a platform for further discussion on the morality of STAs more generally and, in the case of constructive ecumenical expressivism, perhaps point the way to what an agreed normative approach to policing all video game content might look like.

Before discussing any of this, however, I will present a brief exposition of the legal status of virtual paedophilia (for now, under the umbrella term 'virtual child pornography'), noting similarities and differences between the legislation of the USA and UK, respectively. I intend to

discuss the US position first and in more detail simply because (1) more academic literature is available on US legislation, particularly in relation to freedom of expression, (2) it provides a good comparison with UK and other countries' legislation and (3) current debate on the criminalization of virtual child pornography, which I wish to use to inform my discussion on the morality of virtual paedophilia, is largely based on US legislation. First, however, a point of clarification is required.

1.2.2   *The Homogeneity of Virtual Murder*

Bartel (2012) queries what he perceives to be Luck's treatment of violent video games containing enactments of murder: that they are essentially treated (by Luck) as the same; as part of one homogenous group. Bartel claims that, morally, gamer's will approach acts of killing, including murder, within games in different ways. I accept that gamers may well do this, depending on context. This context may include the reason for the killing within the narrative/gameplay (e.g. self-defence, revenge; see Ali' s work in Section 5.3; Hartmann et al. 2010), the availability of options as determined by the game mechanics (i.e. whether different outcomes are available to the player or whether actions and/or moral constraints are imposed on them *qua* their character; see Bartel 2015; Bartle 2008; Pohl 2008; Vanacker and Heider 2012; Zagal 2009), the level of violence and graphic realism (Barlett and Rodeheffer 2009; Krcmar et al. 2011; Wood et al. 2004; Zumbach et al. 2015) and so on. But the fact remains that all of these acts, in whatever context they are presented, in virtue of the gamer being able to *choose* to engage with them or not (even if 'not' ultimately means exiting the game), are permitted. It is this fact that makes all forms of *unlawful* killing part of a homogenous group: they are all enactments of something that is prohibited in the real world and yet permitted within the gameplay. In numerous other ways, they may differ, and this may impact on the player's psychological and moral appraisal of the enactment (Sicart 2009). Nevertheless, it is their permissibility *tout court* that keeps them part of the same group; and where this group contains the intentional and unlawful killing of another person, as it does here, I will refer to these acts collectively as virtual murder. Having said that, in Section 5.2, I will consider the importance to the gamer's dilemma of Patridge's distinction between run-of-the-mill virtual murder and more extreme enactments, such as those described in games like *Postal 2*.

1.3   THE LEGAL STATUS OF VIRTUAL CHILD PORNOGRAPHY IN THE USA AND UK

Child pornography typically involves the sexualized image of a child (or children) which often includes the child engaged in some form of sexual activity. Where this is the case, the image amounts to a record of an *actual* event (in effect, a sexual assault) involving at least one *actual* child. In contrast, in the case of virtual child pornography, what is accepted is that the image of the child is computer generated, meaning that its creation did not involve an actual child, nor is it intentionally meant to represent a particular child, living or dead. Consequently, objections to virtual child pornography cannot appeal to any kind of argument based on abusive production (Sandin 2004). Given my interest in the gamer's dilemma, the example of virtual child pornography I intend to focus on in this and further discussion (but not to the exclusion of other examples), is the virtual representation (*qua* computer-generated image) of a child engaged in sexual activity with an adult.

1.3.1   US Legislation

In the USA, the 1996 Child Pornography Prevention Act (CPPA) was the first attempt by the US Congress to respond to the digital era by alluding (rather than making explicit reference) to the *virtual* sexual imagery of children within its definition of child pornography. The new definition sought to criminalize not only that which depicts actual sexual activity involving a minor (in the case of the USA, someone below 18 years of age) but also that which *appears* to depict a minor engaging in sexual activity, or *conveys the impression* that a minor is involved (Bird 2011; Rogers 2009; Russell 2008). In 2002, however, a ruling by the US Supreme Court (in the case of Ashcroft v. Free Speech Coalition; based on a 6-3 decision), directly challenged the CPPA, claiming that aspects of the legislation were overbroad and therefore unconstitutional, insofar as they prevented freedom of expression (Kosse 2004; Mota 2002). Thus the US Supreme Court ruled that whilst "it remains illegal to make, show or possess sexually explicit pictures of children ... [there is] no compelling reason to prohibit the manufacture or exhibition of pictures which merely *appear* to be of children" (Levy 2002, p. 319). Moreover, with regard to images of a purely digital origin – that do not involve any actual minors and therefore do not amount to a record of an actual crime – the Supreme Court ruled that as the US child pornography laws were implemented to prevent the victimization of children, and as there

is no victim in cases of virtual child pornography, there is no compelling reason to restrict such freedom of expression (however, see Goldblatt 2012, for an attempted rebuttal of this claim). It is important to make clear, though, that the 2002 ruling did not affect the continued prohibition of 'morphed' images: namely, images of real children which have been integrated with some other image or in some way altered in order to create child pornography (Karnold 2000).

The Supreme Court did acknowledge that computer-generated images may lead to actual instances of child molestation, but they ruled that, at present, there is no evidence to suggest that a causal link between these images and actual abuse is anything other than contingent and indirect (Williams 2004). They reasoned as follows:

1. Virtual child pornography is not intrinsically related to child sexual abuse in the way actual child pornography is and so cannot be linked to any actual crime.
2. Any connection with actual child sexual abuse is indirect and contingent and so cannot be said necessarily to be connected to any future child abuse.
3. Prohibition of virtual child pornography cannot be based on the *possibility* that it will cause harm to some children.

In response to this ruling, in 2003, the US Congress introduced the PROTECT Act (which stands for *Prosecutional Remedies and Other Tools to End the Exploitation of Children Today*).[2] The PROTECT Act sought to clarify the overbroad nature of terms within the CPPA (like *appears to be* or *conveys the impression*) by seeking to prohibit virtual images that are *indistinguishable from* or *virtually indistinguishable from* actual images of children. The measure of whether a virtual image is indistinguishable from an actual image of a child (or virtually indistinguishable) is based on the extent to which an ordinary person is able to tell the difference between the two. The PROTECT Act does not therefore prohibit drawings, cartoons, sculptures and paintings of child sexual activity *per se*, given that such imagery *is* distinguishable to the average person. For the same reason, it does not criminalize (*inter alia*) plays and films such as *Romeo and Juliet* or *Titanic* or *American Beauty*, which depict adult performers appearing as minors engaged in sexual activity; thereby alleviating a previous criticism levelled at the original 1996 CPPA: that it was overly restrictive.

The PROTECT Act (section 1466A) does, however, limit the permissibility of such representations where they are considered to be obscene or 'hard core' (Bird 2011). In other words, regardless of their distinctiveness from any imagery of actual children and therefore regardless of the medium used (meaning that drawings, paintings and so on, are *included*), if a virtual image of a sexualized child or of a child involved in sexual activity is judged to be obscene, then it is deemed to be a form of child pornography subject to prosecution under the law. Indeed, as Kornegay (2006) notes: perhaps "an obscenity offence is the most appropriate way of proscribing content not produced with actual children" (p. 2167).

In the USA, obscenity is based on accepted contemporary community standards (the *Miller* test); basically, what a typical community would find obscene. What counts as obscene in the USA, then, is "not based on fact or policy, or harm done, but rather on a specific moral worldview" (Russell 2008, p. 1494).

> To be obscene, as the law defines such a status, is to belong to a legal class of things, which varies over time and space. This is because attitudes and views about what is appropriate and offensive change over time in communities. (White 2006, p. 31)

Specifically, the law criminalizes:

> . . . a visual depiction of any kind, including a "drawing, cartoon, sculpture or painting" that "depicts a minor engaging in sexually explicit conduct and is obscene" or "depicts an image that is, or appears to be, of a minor engaging in . . . sexual intercourse . . . and lacks serious literary, artistic, political, or scientific value" (18 USC §1466A) (Samenow 2012, p. 19).

Permitting a visual depiction that might otherwise be prohibited under an obscenity ruling as long as it is considered to be of serious literary, artistic, political or scientific value is likewise not without its problems. It is perhaps a matter for conjecture whether the following examples would or should fall foul of the PROTECT Act:

- In the USA in the 1990s, the work of photographers Jock Sturges (e.g., *The Last Days of Summer* and *Radiant Identities*) and David Hamilton (e.g., *The Age of Innocence*), which typically involves nude adolescent

models, were accused of violating child pornography legislation, although attempts to prosecute failed (Moehringer 1998).

- In 2007, artists Zoe Hartnell and Sysperia Poppy created artworks for their online gallery, *The King Has Fallen*, depicting erotic dolls in what has been described as a Victorian "Gothic Lolita" style. After growing controversy over the depictions, the gallery was taken off-line by the artists (Lichty 2009).

- In 2009, the Tate Modern in London was embroiled in controversy when it decided to exhibit a piece by artist Richard Prince entitled *Spiritual America* (see Adler 1996). The artwork is a photograph of a photograph of actress Brooke Shields, aged 10. She is depicted naked with oiled skin and heavy make-up, staring directly at the camera in what has been described as a provocative pose. The photograph was displayed away from the other exhibits, behind a closed door, with a warning that some may find the artwork 'challenging' (Singh 2009).[3]

1.3.2   UK Legislation

In the UK, even though it is accepted that sexual images of actual children and virtual children are not the same, the 2003 *Sexual Offences Act* and the 2009 Coroners and Justice Act in many respects treat them *as if* they are (See Ost 2010, for a detailed discussion). Under the UK law, no distinction is made regarding their *criminality*. As section 6A.1 of the Sexual Offences Act (SOA) states:

> The SOA [Sexual Offences Act] 2003 makes amendments to the Protection of Children Act 1978 and the Criminal Justice Act 1988. It is now a crime to take, make, permit to take, distribute, show, possess, possess with intent to distribute, or to advertise indecent photographs or *pseudo-photographs* of any person below the age of 18 (emphasis added).

The Coroners and Justice Act (65:2) also broadens the definition of 'image' to include a moving or still image *produced by any means.* Pseudo-images and images produced by any means are therefore taken to include cartoons, drawings and computer-generated images (as well as other material) which depict, or *appear to depict*, a child (someone under the age of 18) engaged in some form of sexual activity (see also Section 84(7) of the *Criminal Justice and Public Order Act* 1994 which

states that a 'pseudo-photograph' means an image, whether made by computer graphics or otherwise howsoever, which appears to be a photograph). Section 6A.3 of the SOA does acknowledge some difference between actual and pseudo-images, which should perhaps be reflected in sentencing, but at the same time recognizing the potential for equivalence under the law in more serious cases:

> 6A.3 Pseudo-photographs should generally be treated as less serious than real images. However, they can be just as serious as photographs of a real child, for example, where the imagery is particularly *grotesque* and beyond the scope of normal photography (emphasis added).

Given that no children are involved and therefore directly harmed in the production of virtual or pseudo-images, in the case of UK legislation, what is driving harsher sentencing, although not criminalization *per se*, is the degree to which the imagery is judged to be obscene (in this regard it is similar to the PROTECT Act). As Williams (2004) notes, in the case of virtual child pornography: "the criminal law is linked to the indecency of the image depicted and not to the harm suffered by the child" (p. 246).

The UK Obscene Publications Act 1959 determines something to be obscene:

> [I]f its effect or . . . the effect of any one of its items is, if taken as a whole, such as to tend to deprave and corrupt persons who are likely, having regard to all the relevant circumstances, to read, see or hear the matter contained or embodied in it. (Section 1:1)

Therefore, rather than the measure of obscenity being rooted in some form of offence principle reflecting community standards (as in the USA), classification is based on whether the material is likely to deprave or corrupt those who have access to it. In other words, what is considered obscene is couched in social pathology such that there would be a tendency towards 'moral and physical harm caused to vulnerable persons by exposure to obscene writings and images' (Hunter et al. 1993, p. 138). Potentially, this could lead to what McGlynn and Rackley (2009) refer to as *cultural harm* (see Cappuccio 2012, for more detailed discussion on this issue; see, also, Section 5.1).[4]

1.4   Some Closing Remarks

In bringing this chapter to a close, I would like to reiterate my earlier point that the focus of this book is on the morality of virtual enactments within video games, with particular emphasis on representations of murder and paedophilia, irrespective of their respective legal status. That said, there will be times when moral discussion may coincide with published legal argument and jurisprudence. Where this is the case, reference will be made to the similarity between the two accounts without seeking to use this similarity to convince the reader of the importance of the particular moral position. Without wishing to sound trite, it is my view that the merits of a moral argument should be determined by the quality of the argument itself, including, where applicable, the strength and validity of the evidence it may draw on, and not on the extent to which it aligns itself with a particular legal position. Nonetheless, it is not my intention to extol the virtues of this viewpoint while overlooking much of the good work that has been done debating the legality of virtual child pornography, and therefore ignoring the value to be had from drawing on legal argument to inform and illuminate moral debate.

In conclusion, although Section 1.3. provided only a rudimentary outline of some of the legal positions and arguments regarding the criminalization of virtual child pornography, hopefully, what has been made clear is that, hypothetically (given no commercially made games are yet available), where those players who engage in virtual paedophilia are adults, and do so willingly, much discussion has gone into whether such activity should be criminalized and, at present, as we have seen, different countries hold different views. Matters of legality aside, then, what is of interest and what will ground the discussion to come is how we might respond to the following questions:

1. Irrespective of whether it is legal to do so, is engaging in virtual paedophilia something that should be considered morally wrong?
2. In light one's response to (1), is virtual murder liable to the same moral outcome for the same moral reason(s)?

How we answer these questions will likely determine whether the gamer's dilemma can be resolved.

NOTES

1. The clinical use of the term 'paedophile' is reserved for those who have a sexual interest in prepubescent children (Berlin and Sawyer 2012). Those with a sexual interest in pubescent and prepubescent children are known as hebephiles (Neutze et al. 2011). While recognizing this difference, the term 'paedophile' will nevertheless continue to be used in a manner consistent with popular rather than clinical usage.
2. In 2008, the PROTECT Act was upheld by the Supreme Court as constitutional.
3. As an aside, although it does not concern virtual images of children but is nevertheless related to the discussion on child pornography law, Scheeres (2002) reports on Internet sites, such as *Nude Boys World* and *Sunny Lolitas*, which purportedly contain 'child erotica'. The images are typically of naked children that do not violate the US child pornography laws because they are not sexually explicit.
4. It is worth noting a few more examples of legislation from around the world as a way of illustrating further the lack of consensus over virtual child pornography: both in terms of what it entails and the age of a 'minor' within the definition. Australian legislation (for example) varies from state to state, as does the age someone is classified as a child in the context of pornographic imagery (either under 16, under 17 or under 18, depending on where you live). As for what constitutes child pornography, New South Wales, Queensland, the Northern Territory, Tasmania and Western Australia are similar to each other and well as to the US PROTECT Act, insofar as child pornography amounts to "material depicting, describing or representing a child (or, in Western Australia, a part of a child), or someone who appears to be a child, in a sexual context or engaged in a sexual act in a way likely to offend a reasonable person" (Croft and Murray 2013, p. 91). In South Australia and the Australian Capital Territory, it is not a requirement that the material be offensive; rather it must be intended to be used for sexual gratification. In Victoria, child pornography is defined as "a film, photograph, publication or computer game that describes or depicts a person who is, or appears to be, a minor engaging in sexual activity or depicted in an indecent sexual manner or context" (ibid.). The definition's direct reference to computer games gives the clearest indication of the criminalization of the sort of virtual child pornography we are discussing here (see McLelland 2005; and Simpson 2009, for further discussion). By way of a further example, as indicated in section 163.1a of the Canadian criminal code, Canadian child pornography law likewise does not differentiate between virtual and actual images of sex acts involving children (i.e. anyone who is

or is depicted as being under 18 years of age). An image is classified as child pornography whether or not it was made by electronic or mechanical means, and therefore whether it is a photograph, film, video or *some other visual representation*. (Taken from http://laws-lois.justice.gc.ca/eng/acts/C-46/section-163.1.html. Accessed 11/7/16)

Social Convention and the Likelihood of Harm: Luck's Initial Attempts at Resolving the Dilemma

Abstract This chapter begins by discussing Morgan Luck's initial attempt at resolving the gamer's dilemma through an appeal to social convention, whereby we adopt a normative ethic based on societal norms. It challenges this view, through the use of Japanese manga (cartoon) imagery depicting what appear to be highly sexualized minors, by showing that not all societies prohibit virtual paedophilia, or would consider such imagery to be child pornography. Luck's next argument is then appraised, which involves differentiating between virtual murder and virtual paedophilia based on the latter's increased likelihood of leading to harm. Research looking at the effects of violent video games, as well as the relationship between child pornography and molestation, indicates that, at present, the reason for this differentiation cannot be empirically substantiated.

Keywords Moral consensus · Japanese manga · Harm as a necessary and a sufficient condition for immorality · Violent video game research · Child pornography and molestation

In this chapter, I present Luck's first two attempts at resolving the gamer's dilemma. The first is based on an appeal to a difference in social convention, and the second to differences in the likelihood of harm occurring. Both arguments are dismissed by Luck as unconvincing. Let us examine each in more detail to understand why.

2.1   AN APPEAL TO SOCIAL CONVENTION

Luck (2009) first considers justifying selective prohibition through an appeal to social convention. We are told that it is socially acceptable to play video games in which one can virtually murder but not socially acceptable to engage in virtual paedophilia. (Tacit reference is perhaps being made here to a Westernized social convention; see, by way of contrast, discussion on Japanese manga and anime cartoons, particularly Hentai, in Section 2.1.1.). Given that this is the case, or at least is *said to be* the case, if we are to endorse social convention as a moral guide – as some form of moral wisdom (Kass 2002), perhaps even indicative of an independent moral truth – then the following must hold:

1. Social convention is a form of moral wisdom; that is, a measure or indicator or what is morally wrong and morally right (or, at the very least, not morally wrong), and therefore what should and should not be permitted
2. Social convention finds it acceptable to engage in virtual murder but unacceptable to engage in virtual paedophilia
3. Therefore, it is morally wrong to engage in virtual paedophilia but not morally wrong to engage in virtual murder (although this is not necessarily saying that it is morally right to do so).

When considering virtual murder, to claim that within one's society it is socially acceptable to enact murder implies that there is a high level of consensus; but in fact there have been and continue to be numerous dissenting voices making rival claims and generally disputing this view, which Luck does acknowledge. Using the UK and USA as our focus, a cursory search through newspaper headlines (for example) quickly reveals a regular questioning of the relationship between violent video game content and violent or otherwise antisocial behaviour (see e.g. Casey 2015; Orland 2013; Schreler 2015; Singer 2007). Later, in Section 2.4, we will see how this lack of consensus is reflected in the findings of empirical research on violent video games. Of course, with a relatively simple amendment, one could accept such dissent, particularly in the case of extreme violence, without altering the essential thrust of Luck's description and therefore his appeal to social convention. One could hold that within the UK and USA it is *more* acceptable to engage in virtual violence than virtual paedophilia. But even if social convention is described in these

terms, what evidence is there to support this? I am not aware of the publication of findings on the public's perception of virtual murder compared to virtual paedophilia – that the former is more acceptable than the latter – or even media debates on the matter. Instead, such a claim seems to be tapping into some kind of tacit acceptance of this difference which ultimately conforms to our (most people's) intuitions about the two forms of virtual enactment. These intuitions are then relied on, in the absence of other more empirically grounded and verifiable evidence, and used as the bedrock for the social consensus or majority view that Luck holds up for examination as a possible candidate for resolving the dilemma.

Interestingly, whether social convention is as Luck describes it is not that important: for, whatever the truth of the matter, it is a contingent truth. Luck's description either happens to be true or it happens not to be true. It is not that important because Luck does not need his description of the social convention on video game content to be true in order to make his point that, ultimately, social convention is an unsatisfactory basis for moral judgement and therefore an unsatisfactory reason to prohibit certain content. Instead, he can simply invite us to imagine that it is acceptable or merely *more* acceptable to engage in virtual murder than virtual paedophilia. In effect, he can ask, for the sake of argument, to allow that such a social convention exists. Where this is the case, he can then argue, as indeed he does, that while this fact (if we *allow* that it is a fact) might explain the relationship between our intuitions and social consensus (that, say, our intuition is in fact merely the expression of what a given society holds to be the case), is not sufficient to justify a *normative* position regarding selective prohibition. And this is the point Luck is making.

2.1.1   Manga: Reflecting Japanese Social Convention

To understand why, consider the example of the Japanese *manga* and *anime* illustrative forms which are popular across all ages in Japan (Norris 2009; Sabin 1993; Wilson 1999). Pertinent to this discussion is a classification or genre known in the West as *Hentai*, which translates as 'changed' or 'strange figure' or, in some translation, 'pervert'. Hentai imagery typically involves some form of metamorphosis or aberration indicative of a sexual perversion or abnormality (Ortega-Brena 2009). Masuchika (2015) alludes to a standard Western view of manga's sexualized imagery (although he does not use the term Hentai) when he states: "Japanese manga have an unsavoury reputation of containing seemingly pornographic, or even obscene, material.

News reporters have written about the proliferation of manga that could be classified as child pornography" (2015, p. 57).

One such news report – by Justin Norrie, reporting in *The Sydney Morning Herald* – illustrates Masuchika's point well.

> [Tokyo's] largest electronics and entertainment quarter also has an abundance of something far more unsavoury to Western sensibilities: a range of manga comics, anime films and video games that would be regarded in many countries as the biggest collection of child pornography anywhere. Lining the shelves of several stores in Akihabara and other Tokyo districts are thousands of manga stories such as *Junior Rape, Under Nine* and *Unfinished School Girl* that hint ominously at the explicit images within. The illustrations in the most hardcore titles . . . are legal under the country's child pornography laws. (Norrie 2010, p. 1)

While acknowledging their sexualized and often explicit content, Masuchika nevertheless tries (for the benefit of those less familiar with Japanese culture) to contextualize manga imagery:

> There is no doubting that there are highly sexualized manga being produced and distributed in Japan . . . This perception of manga being pornographic is partly derived by western eyes looking at the culture of Japan. The artwork found in manga is standardized . . . The large eyes often give the females an adolescent look that the Japanese call 'Kawaii' or 'cuteness' . . . There is [also] more of an acceptance of nudity in Japanese culture than in western cultures, and this is reflected in the manga. (2015, p. 57)

Similarly, Ortega-Brena (2009) notes how manga (although she tends to use the term Hentai), and the concept of animated pornography, challenges Western attitudes towards what we might crudely regard as 'cartoons'. For Ortega-Brena, this is because, in the West, animation is widely regarded as childlike, and we find it difficult to shake off this association (although see Taylor 2009, for a newspaper report on erotic comics and UK child pornography legislation). In contrast to this view, the erotic material characteristic of Hentai needs to be understood within the context of Japan's 'historical and cultural approach to both aesthetics and sexuality' (p. 18), as well as their tradition of sexual explicitness and long held belief in the naturalness of sexuality (see, also, Gwynne 2013, for a post-feminist critique of manga). Galbraith (2011) likewise discusses a genre of manga and anime (as well as games) referred to as 'Lolicon'

(derived from 'Lolita complex') which features characters, seemingly 'underage', depicted in various sexual and sometimes violent situations. Galbraith describes how critics have claimed that there is only *one way* to interpret this material – implying that is has only prurient appeal – but nevertheless defends the genre against such criticism. In a review of Lolicon culture, he suggests that consumers have a much more nuanced understanding of the imagery than critics would have us believe and that many desire *shōjo* characters (meaning 'young girl') precisely because they are unreal; or, as Honda (2010) informs us, because they are evanescent.

The difference in social consensus between Japan and the West over sexualized cartoon imagery of children, perhaps shaped by historical cultural differences, is evidenced by the unwillingness of the Japanese authorities to amend further their child pornography laws. In 1999, it became illegal to produce and distribute child pornography involving actual children, but not to own it. In 2014, the possession of child pornography became criminalized but, significantly, not virtual images of 'child' sexual activity characteristic of Hentai/Lolicon (Hellmann 2014). Interestingly, Kinsella (1998) describes how a moral panic gripped Japan in the late 1980s and into the 1990s on account of the growing popularity of the *amateur* manga movement and subsequent subculture. The majority of amateur manga artists were women in their teens and early twenties, and the movement was organized by and for people of this age group (see also Shigematsu 1999).[1] Kinsella provides an insight into their work:

> A limitless secret world of smoldering underground clubs where baby girls in bikinis weild Uzi submachine guns and Russian Eskimos D.J in Elizabethan court dress. Grey catacombs of deserted rain-swept streets where beautiful women in impeccable Nazi uniforms sport unexpected erections . . . Such is the stuff that the amateur manga is made of. (1998, p. 289)

The moral panic was caused more by the authority's perception of an alienated and disenfranchised youth than by the sexualized content of their work, which did not differ in its explicitness from professional manga. Tokyo police attempted to censor sexual images in unpublished amateur manga and prevent their wider circulation, even trying to enforce a by-law prohibiting the sale of sexually explicit published material to minors (those under the age of 18), despite the fact that a large proportion of the amateur manga was produced and sold by those under 18, often

distributed through comic markets (known as *komiketto*). The moral panic was therefore more about perceived anti-establishment or anti-traditional behaviour on the part of a large section of Japanese youth than about work featuring sexualized and putatively perverse content.

The Japanese government's resistance to align with the US, UK and other counties legislation on virtual child pornography stems from a fear of restricting its nation's freedom of expression (Hellmann 2014). (We have, of course, seen this same fear expressed in debates on the US child pornography legislation.) There is also the possibility that Western onlookers unfamiliar with the traditions of manga (and therefore Hentai), which has its own inherent stereotyped aesthetic (Lau 2013) and stylized visual language (Cohen 2016), simply fail to appreciate the artistic value of this form of imagery. (We touched on a similar point when briefly considering examples of art in relation to the US child pornography law in Section 1.3.1)

If we take Japan as a contemporary example of a society with different social conventions regarding sexualized imagery, then it is not difficult to imagine how a different culture, whether Japan or some hypothetical society, with different traditions regarding sexuality and fiction, would not necessarily (intuitively) adhere to the view that virtual paedophilia should be morally prohibited and therefore should be distinguished from virtual murder in any morally relevant way (a similar point is made by Sandin 2004). Or perhaps, it is more accurate to say that they would not accept that what we in the West judge to be virtual paedophilia, based on legislation (see e.g. Kontominas 2008), is in fact virtual paedophilia. The issue perhaps rests on a question of interpretation (a point I shall return to in Chapter 6 when discussing constructive ecumenical expressivism).

2.2   SILENCING THE MINORITY VIEW

Leaving aside the issue of inter-societal differences in norms and conventions, the premise on which the social consensus objection is based (say within countries like the UK and USA) is the fairly safe assumption that the majority of people find the idea of virtual paedophilia, to say the least, unpleasant and, in fact, more likely morally repugnant. One could even add that virtual paedophilia does not contribute to the 'market place of ideas' (Russell 2008, p. 1494). Or, as Bird (2011) puts it, the value of materials depicting children performing sexual acts is 'slight or *de minimis*' (p. 163). For these reasons, Russell conjectures, one might hold that it is not an affront to the value we place on our democracy if such imagery is

suppressed on account of the majority finding it offensive. Moreover, given our (assumed) moral attitude, one might further conjecture, even in full recognition of the fact that virtual paedophilia is a victimless crime, that those who make up the majority find it hard to understand how such (perceived) immorality could go unpunished (Bird 2011). Therefore, given the association between activity and moral response and the suggested need for punishment, perhaps one can begin to understand why the UK government believes that the criminalization of virtual child pornography is justified given the inappropriate feelings towards children such engagement might engender (Home Office 2007). Such a move no longer takes the object of our moral concern to be children and their protection but, rather, what the majority consider normal and appropriate (Eneman et al. 2009). Yet, equally, one might argue that legislation, or, perhaps more pertinent here, a normative ethic concerning the moral prohibition of certain content (virtual or otherwise), should neither be the product of, or in any way try to appease, 'gut democracy' (Russell 2008, p. 1496), nor should it be tolerated out of habit.

Debate on virtual child pornography can make ostensive our deeply held moral convictions, "where subjective opinion and moral position tend to dominate the debate rather than rational argument" (Eneman et al. 2009, p. 3). Moreover, Eneman et al. (2009) argue that there is an inherent problem with legislation that sets out merely to articulate in legal terms the preferred moral attitude of a given society (e.g. where the criminalizing of virtual child pornography becomes nothing but an attempt to turn into law the favoured position of the moral majority). Thus they state, "If moral convergence or consensus is the main reason [for] criminalising virtual child pornography, then one needs to understand that such preferences are subject to change over time and that regulations may change as a consequence" (p. 9).

One of the issues with the social consensus argument, of course, is the necessary restriction it places on freedom of expression. Freedom of expression becomes somewhat restricted and even demonstrably false if it conforms to the view of the majority; if, that is, it only applies to those ideas favourably received because they accord with accepted opinion or are considered tolerable. Genuine freedom of expression must go some way beyond this, even if it falls short of absolute freedom. Unfortunately, the area between 'beyond' and 'absolute' is somewhat nebulous; indeed, this fact grounds much of the discussion here. This remains the case whether one's concern is with the law or establishing a normative ethic.

In light of this, and borrowing from Russell (2008), consider the following pertinent point. While one may be sympathetic to the argument that few of us would be willing to march off to war, or have our sons and daughters do the same, to preserve the right of those who so desire to engage in virtual paedophilia, perhaps this is the problem. If we are unwilling to defend undesirable viewpoints then tyranny wins by default: for the same argument, based on social consensus, that is being used against virtual paedophilia has been and is still being used in some parts of the world against homosexuality, or interracial couples, and will likely be used at some future time against some other minority group.

As noted, an appeal to social convention should not be thought of as requiring unanimous agreement, either in favour of permissibility or against it. Instead, for the enactment of virtual murder within video games to be accepted or virtual paedophilia to be rejected, based on social convention, what is required is simply sufficient agreement (which could include tacit acceptance) by the members of that society. Of course, where a normative ethic seeks to establish a universal principle of right and wrong, this requires (majority) consensus among all societies, otherwise it cannot be applied (as the Japanese example helps illustrate). But even if a universal consensus were achieved, it is still vulnerable to a charge that the enforcement of a normative ethic stem simply from an appeal to the moral view of the majority, which may change over time, thus making our putative normative ethic somewhat fluid (which, of course, is Luck's point and reason for rejecting this approach). Being fluid is not necessarily a bad thing, it has to be said; although this is something I will leave for now and return to in Chapter 6 when again discussing constructive ecumenical expressivism. All I will say in anticipation of the discussion to come is that it is not the possibility of fluidity that should be of concern to us in our quest for a normative ethic but, rather, the reason(s) for this fluidity. Where the reasoning is sound then moral change should follow.

Within a given society, and as a way of acknowledging any dissent towards the permissibility of virtual murder or prohibition of virtual paedophilia, where harm is raised as an issue, the conventional position need not champion the total harmlessness of enacting virtual murder or indeed the absolute harmfulness of virtual paedophilia. One could simply concede that the former is *less* harmful than the latter; although one might reasonably ask how great the discrepancy between the two levels of harm has to be to justify a difference in moral judgement. Having rejected an appeal to social convention, for the reasons discussed, a consideration of

the role played by the *likelihood* of harm brings us to Luck's next attempt at resolving the gamer's dilemma and of achieving selective prohibition: something he calls *significant likelihoods*. A position he again rejects.

2.3   AN APPEAL TO SIGNIFICANT LIKELIHOODS

[T]he claim that it [pornography] is harmful is usually a refuge for those who, really, think that it is objectionable in itself, but believe that couching their objections in terms of harm is more likely to command agreement. (Graham 1999, p. 105)

[I]f somebody creates or possesses indecent pseudo-photographs of children, he is a potential child abuser and will abuse children in the future. (Edwards and Waelde 1997, p. 228; cited in Karnold 2000, p. 61; purportedly quoting a belief held at the time by UK Police)

An appeal to 'significant likelihoods' is a way of allowing that both forms of enactment (virtual paedophilia and virtual murder) could result in some harm – insofar as neither is nor, importantly, *has to be* harm free – while maintaining that, of the two, only virtual paedophilia is *significantly* likely to result in harm. This possible response is grounded on a consequentialist approach to morality, as should be evident from premise (i). In keeping with consequentialism, broadly construed, Luck presents the following argument:

   (i) Any act which is significantly likely to result in harm is immoral;
  (ii) Virtual paedophilia is significantly likely to result in harm;
 (iii) Therefore, virtual paedophilia is immoral.

In contrast:

  (iv) Virtual murder is not significantly likely to result in harm;
   (v) Therefore, virtual murder is not immoral.

Luck actually makes a more specific connection: namely, that virtual paedophilia is significant more likely to result in actual paedophilia, while virtual murder is not significantly more likely to result in actual murder. It is worth pausing at this point to consider the phrase used by Luck: namely 'to result in'. This is ambiguous insofar as it suggests a causal connection without explicitly stating it (i.e. it does not state explicitly that virtual

paedophilia is significantly likely to *cause* harm, or that virtual murder is not significantly likely to *cause* harm). The phrase 'to result in' could arguably be describing a correlation whereby virtual paedophilia is correlated with harm, such that those who regularly engage in virtual paedophilia are *also* more likely be the cause of another's harm (that of a child in this case), to the extent that the more virtual paedophilia is enacted the higher the likelihood of harm occurring. In such a situation, while it is clear that the person who engaged in virtual paedophilia caused the child to suffer harm, it is not clear that this action was itself caused by the individual engaging in virtual paedophilia. Throughout the chapter, I shall continue to use the phrase 'to result in', while fully acknowledging its ambiguity. It is also worth noting that I am going to discuss the appeal to significant likelihoods approach with reference to harm more generally; something which should not undermine Luck's original argument or detract from the critical points I wish to discuss.

2.3.1   Deductive Fallacy, and Necessary and Sufficient Conditions

Returning to the deductive reasoning presented above. If one accepts the truth of (i) and (ii), then (iii) necessarily follows, but (v) does not necessarily follow from (i) and (iv). This is because premise (i) is describing a *sufficient* condition for a claim to immorality, not a *necessary* one. Luck appears to be aware of this but does not press the point. Formally, 'significant likelihood' as a sufficient condition (SL_s) would take the following form:

(SL_s) An act is immoral *if* it is significantly likely to result in harm.

   If virtual paedophilia is significantly likely to result in harm then, in accordance with premise (i) (as a sufficient condition), virtual paedophilia is immoral. Importantly, though, even if virtual murder is not significantly likely to result in harm, the conclusion presented in (v) does not *necessarily* follow from this (to say that it does it to commit a deductive fallacy). This is because virtual murder could satisfy some other (yet to be identified) sufficient condition for a claim to immorality. Where 'significant likelihood' (of harm) is sufficient for a claim to immorality, we cannot deduce from this and the (alleged) truth that 'virtual murder is not significantly likely to result in harm but virtual paedophilia is' that virtual paedophilia is immoral *and* virtual murder is not. We can only deduce the former moral position, not the latter. To reiterate, this is because (a) 'significant likelihood' is not presented as a necessary condition for a claim to immorality

and (b) given (a), there may be some other sufficient condition which virtual murder satisfies, thereby making it immoral. Suppose, instead, then, 'significant likelihood' is presented as a necessary condition (SL_n) for a claim to immorality.

(SL_n) An act is immoral *only if* it is significantly likely to result in harm.

In this case, given the truth of (ii) and (iv), the respective conclusions (iii) and (v) would necessarily follow, and with this would come a means of differentiating between virtual murder and virtual paedophilia.

Of course, one may object to the use of 'significant likelihood' as a necessary condition. It does not seem unreasonable to hold that an act can still be immoral even if no harm occurred as a consequence (see Section 5.1). To illustrate, consider Gutierrez and Ginner-Sorolla (2007) fictitious example of two consenting adults who are brother and sister (and knowingly so). One day, they decide to engage in a one-off act of lovemaking (a similar example is discussed by Haidt and Hersh 2001). Contraception is used, and neither party regret what they did. Afterwards, each is able to engage in meaningful relationships with other people, and their brother–sister bond is not damaged. What occurred also remains forever private. If one accepts the truth of the description, no harm resulted. Nevertheless, Gutierrez and Ginner-Sorolla found that, when presented with the fiction, participants typically referred to the event as 'disgusting' and considered what had happened to be immoral. It is also interesting to note that many of the participants found it difficult to believe that no harm had in fact occurred. Perhaps this supports the view that even if one is prepared to accept that, in this case – this somewhat contrived case – no one was harmed, typically, one would still be of the opinion that a similar type of event would more often than not result in harm.

Putting their incredulity aside, what the Gutierrez and Ginner-Sorolla fiction illustrates is that even when participants accepted, albeit reluctantly, that no harm had occurred, many still considered the act to be immoral. Such a finding is compatible with the view that 'significant likelihood' (of harm) is sufficient for a claim to immorality while also allowing that it is not necessary. The next example (adapted from Jesse Prinz) reinforces the point. Why is it wrong to sexually assault a young child who will never remember the event? One could contrive to make the assault such that no physical damage was inflicted on the child. In such a scenario, the lack of memory of the event, along with the lack of awareness of what was actually happening at the time, as well as the lack of physical damage, are all meant to convey (convince, even) that no harm occurred. In answering the

question – "Is it morally wrong to do this?" – Prinz (2007) responds thus: to say that it is just wrong is not to fail to have a reason, or to be unable to articulate it. "[W]e are not obviating reason", he declares; rather, "we are explicitly giving one" (p. 31). Saying it is wrong *is* the reason; it is *just wrong*.[2]

As an aside, it is possible that one might wish to argue that the perpetrator is harming themselves through the act of sexually assaulting the young child (perhaps by engaging in an act that is incongruent with one's psychological well-being or ability to flourish, or some such thing); therefore, even if one accepts that the child is not harmed, the event itself is not harm free. Even if this were true, it might perhaps strike the reader as perverse to state that *the reason* why one should not sexually assault a young child, even if no harm befalls them, is because one would be harming oneself (we will return to a variation on this point in Section 4.4 when discussing the argument that virtual paedophilia harms women). Prinz's point seems a reasonable one, then: it is wrong to do this irrespective of any harm.

Like Gutierrez and Ginner-Sorolla's participants when presented with the fictional example of incest, one might have difficulty accepting that the child was not harmed *in some way*. This perhaps reflects the view that one can be harmed by an event even if one is not aware of being harmed (Nagel 1979). As a philosophical position, such a view is contentious, however, although it is not my intention to enter into a discussion on that matter here.

2.4 Is Virtual Paedophilia Significantly More Likely to Result in Harm than Virtual Murder?

Much of the debate around the issue of harm centres on what is meant by harm, of course. But even if we leave open or unresolved what is meant by harm, one could still challenge the utility of 'significant likelihood' in differentiating between virtual murder and virtual paedophilia by denying either premise (ii) or premise (iv). In doing so, one is essentially denying that virtual paedophilia is significantly *more* likely to result in harm compared to virtual murder. So, however one wishes to define harm, one could simply deny that more of it is likely to occur as a result of virtual paedophilia. There are two ways this could be done. One could accept premise (ii) (that virtual paedophilia is significantly likely to result in harm) and

deny premise (iv) (that virtual murder is not significantly likely to result in harm), or one could deny (ii) and accept (iv) This latter option is the approach taken by Luck as one means of rejecting the 'significant likelihood' argument.

Specifically he asks, what reason might we have to accept that virtual paedophilia is significantly likely to result in harm (specifically, to result in actual paedophilia)? *A priori*, what would be required to support this view and therefore premise (ii)? To support premise (ii), one would need to show that:

A. Engaging with virtual child pornography/enacting virtual paedophilia can cause one to form beliefs and pro-attitudes (Davidson 1980) about that which the enactment represents (minimum causal connection) and that these newly acquired beliefs/pro-attitudes contribute to a change in behaviour (a more causally forcefully requirement).
B. Engaging with virtual child pornography/enacting virtual paedophilia can cause a change in behaviour even when no change in beliefs/pro-attitudes occurs (perhaps as a result of feeling compelled to act, such that I behave in a morally proscribed way even though I believe this to be morally wrong).

Should compelling evidence for either A or B be forthcoming then one would have support for at least the first part of the 'significant likelihood' argument: that virtual paedophilia is significantly likely to result in harm (in the form of actual paedophilia). In addition, compelling evidence (beyond our mere intuition) would need to be presented showing virtual murder is not significantly likely to result in harm (in actual murder).

Of the two forms of fiction under scrutiny, over the years, numerous studies have been carried out investigating the relationship between violent fiction (e.g. murder) and actual aggressive behaviour: first in the form of film and television violence (Alia-Klein et al. 2014; Anderson and Bushman 2002; Bushman and Huesmann 2006; Huesmann et al. 2003) and later video game content (see below). A brief presentation of some of the research on violent video games should help us determine the extent to which the data are congruent with our intuitions about violent content (that it is not significantly likely to result in harm) and also dispel those dissenting voices which claim that violent content is significantly likely to result in harm of some description.

2.4.1   Evidence (Virtual Violence)

It would be erroneous to ignore findings supporting the view that playing violent video games is associated with increased violent or otherwise antisocial behaviour (Greitemeyer and Mügge 2014). Anderson et al. (2010), for example, claimed to have found that exposure to video games with violent content is a causal risk factor for increased aggressive behaviour, cognition and affect, and decreases empathy and prosocial behaviour. Ferguson (2007a, b), however, in his meta-analytic review on video game violence, warns us to treat many of the findings supporting a connection to antisocial behaviour with caution, arguing that the measures of aggression used in most studies lack validity and that often the effect sizes are close to zero. He also suggests that there is a bias in the academic literature in favour of those papers which report statistically significant differences between groups. Because of this, and based on conflicting evidence found in the literature, any attempt to posit a direct causal link between video game content and violent (real-world) behaviour should be regarded as overly simplistic, largely uncorroborated and ultimately contentious. Indeed, for dissenting voices and further critical discussion on Anderson et al.'s conclusion, see Bushman et al. (2010); Ferguson and Kilburn (2010) and Huesmann (2010) (see also Bensley and Van Eenwyk 2001 and Ferguson 2011). Moreover, Markey et al. (2015) following a meta-analytic review of data, including Federal Bureau of Investigation crime statistics and video game sales, report:

> Contrary to the claims that violent video games are linked to aggressive assaults and homicides, no evidence was found to suggest that this medium was a major (or minor) contributing cause of violence in the United States. Annual trends in video game sales for the past 33 years were unrelated to violent crime both concurrently and up to 4 years later. Unexpectedly, monthly sales of video games were related to concurrent decreases in aggravated assaults and were unrelated to homicides. (pp. 14–15)

Cunningham et al. (2016) likewise report that there is no evidence supporting an association between violent video games and increased crime in the USA; and in fact conjecture over evidence supporting a slight decrease. Furthermore, in light of the controversy (in the West) over Japanese manga (Section 2.1.1), a particularly pertinent study is carried out by Diamond and Uchiyama (1999). Using official Japanese sex crime statistics, they found that

the increased availability of pornography (including manga) in Japan since the 1990s was correlated with a decrease in sex crimes.

Putting all of this together, then, in terms of the findings of research currently undertaken, there is no consensus on what the effects of playing violent video games are (see Ferguson 2013; but also Bushman et al. 2015; Bushman and Huesmann 2014 and Krahé 2014, for a rebuttal of Ferguson's claims, and therefore as a means of reinforcing the argument for a lack of consensus). Therefore, *a posteriori*, there is no compelling reason (at least where compelling reason requires a consensus in the empirical findings) to challenge the view that enacting virtual murder is not significantly likely to result in harm generally, let alone actual murder. Subsequently, there is no compelling reason to deny premise (iv).

Let us therefore turn our attention to premise (ii). There is currently a paucity of research on the relationship between virtual paedophilia and actual paedophilia. Consequently, evidence-based argument examining the merits of premise (ii) requires that we engage in a degree of extrapolation. To illustrate, in 2008, Bryant and Linz set out to test an assumption made by the US government in defence of the 1996 Child Pornography Protection Act: "that virtual child pornography stimulates and whets adults' appetites for sex with children and that such content can result in the sexual abuse or exploitation of minors becoming acceptable to and even preferred by the viewer" (Bryant and Linz 2008, p. 35). After exposing adults to 'barely legal' pornography, Bryant and Linz concluded that although those who viewed the material were more likely to cognitively associate sexual activity to non-sexual images of minors (based on response latency), there was no evidence that exposure caused participants to be more accepting of child pornography or paedophilia. (Barely legal pornography uses models who are over 18 years of age, but who are depicted as being under or just over the legal age of consent.) Imagery of this nature is not virtual in the sense referred to within the gamer's dilemma, but it is suggestive of the absence of a connection needed to support premise (ii). I do, however, recognize that virtual paedophilia would typically afford an interactive element that is absent in most, if not all, barely legal pornography.

In the absence of research directly testing the relationship between virtual and actual paedophilia, perhaps research looking at the relationship between those who view actual child pornography and engage in hands-on molestation of minors would be informative; with, of course, the caveat that virtual child pornography is not actual child pornography insofar as the depiction is

not of an actual child (although, as we have seen, in some countries, under the law, it is treated as such; see also Section 4.3 for discussion on the lack of ontological equivalence between virtual and actual child pornography). Therefore, one must be cautious when drawing any conclusions given the amount of extrapolation occurring (see, e.g. Jewkes and Wykes 2005, critical discussion on the construction of the 'cyber-paed').

2.4.2   Evidence (Child Pornography)

Most contemporary child pornography offenders commit their offenses over the Internet (Ray et al. 2010). Among these offenders is a distinct category of individuals who restrict their offending behaviour to the voyeuristic pursuit of child abuse images. *Sexual voyeurism* refers to a group of perpetrators who have no documented history of having ever attempted to approach a child sexually (including no evidence of ever wanting to), but who nevertheless manifest a pattern of viewing child pornography, sometimes compulsively (Berlin and Sawyer 2012; Lodato 1998). This category of offender may include those who are impulsive and curious. They may not therefore have a *sexual* interest in children *per se* or their sexual interest may not be exclusively in children. But the category also includes those who use child pornography to fuel existing or developing sexual interests in children (Beech et al. 2008; Seto et al. 2010).

When the computers of *sexual voyeurism* offenders are confiscated, neither evidence is found of sexually inappropriate 'chats' with children; nor, following their arrest, are there reports of children coming forward accusing them of any form of inappropriate contact (Berlin and Sawyer 2012). In fact, Berlin and Sawyer go on to note that many in this category have children of their own who likewise do not report inappropriate sexual contact from this parent. Given this, Berlin and Sawyer conclude:

> [S]ome individuals appear to be experiencing compulsive urges to voyeuristically view such images [of child pornography], devoid of any motivation to actually approach a child sexually. In other words, in such instances, the act of voyeuristically, and often compulsively, viewing such imagery over the Internet would appear to be an end in and of itself; rather than a means to some other end – such as actual sexual contact. (2012, pp. 31–32)

The sexual voyeurism category is compatible with Elliott and Beech's (2009) *periodically prurient* typology, which refers to those who access child pornography sporadically out of a general sense of curiosity or impulsivity, and whose behaviour is likely to be linked to a broader interest in pornography (including 'extreme' pornography) rather than because they have a specific sexual interest in children (see also Elliott et al. 2009). Here, the offender may understand themselves to be taking a more passive role in the viewing of child pornography. The behaviour of these individuals may even be symptomatic of Young's (2001) *Internet-enabled pathology* which manifests itself as deviant online experimentation and forms part of a more general cyber-sex addiction. This category may also include Elliot and Beech's *fantasy-only* typology, which describes those who much more actively seek out, access and often trade images owing to their sexual interest in children, but who do not engage in (and have no known history of) hands-on sexual contact with children.

Research by Seto and Eke (2005) found that child pornography offenders with a history of hands-on sexual abuse were more likely to reoffend, whereas those guilty exclusively of child pornography violations did not go on to contact offend (at least during the study's follow-up period; see Eke et al. 2011, for a further follow-up study). Webb et al. (2007), in turn, found that although there were evidence of reoffending among those convicted of child pornography offences (that were not also child molesters) their reoffending was confined to the accessing of child pornography; there were no evidence that they went on to contact offend. In fact, McCarthy considers the lack of contact offending (including the Internet grooming) to be a factor instrumental in distinguishing between these two deviant groups. To illustrate further, among a sample ($n = 290$) of child pornography offenders with a sexual preference for boys, Riegel (2004) found that 84% reported that the image acted as a substitute for an actual child, with 84.5% of these stating that viewing such imagery did not increase their desire to engage in hands-on abuse with boys.

Importantly, though, although a demarcation based on those who engage in hands-on child abuse and those who only view child pornography may be legitimate in some cases, it nevertheless remains evident that those charged with child solicitation and/or molestation are often caught in possession of child pornography (Kingston et al. 2008; Riegel 2004). There is also the matter of escalation (and whether this reflects an unsubstantiated fear on society's part or constitutes a genuine risk), whereby the offender's deviant sexual excitation could prompt them to seek out

increasingly shocking depictions (Niveau 2010) and/or engage in hands-on child abuse (Seto and Eke 2005). Quayle and Taylor (2002) and Sullivan and Beech (2003), for example, argue that a proclivity for child pornography increases the likelihood of committing a contact offense against a minor either in the form of sexual molestation or sexual solicitation (see also Seto et al. 2012).

What is left unresolved by this research is the extent to which an initial exclusive interest in child pornography is significantly likely to result in perpetrators going on to contact offend, or whether they will remain content with viewing sexualized images of children. An uncertainty echoed by McCarthy (2010):

> The idea of a causal relationship between possessing child pornography and the sexual abuse of minors is not only the basis for child pornography legislation, but also . . . the "harm thesis" wherein it is perceived as "common sense" that viewing pornography causes men to commit sex crimes. The idea of a causal relationship between both phenomena, however, flies in the face of decades of literature on child sexual abuse – literature which contends that child sexual abuse is a complex phenomenon that is best explained by considering various factors . . . (p. 183)

As already expressed, virtual child pornography and actual child pornography are not the same (something I will discuss in more detail in Section 4.3). Therefore, caution must be exercised when considering the findings presented above. It is also worth recalling the view of the US Supreme Court on virtual child pornography (Section 1.3.1): that it is not intrinsically related to child sexual abuse and that any connection with actual child sexual abuse is indirect and contingent (see also, Veber 2004).

Given the paucity of research on virtual child pornography/paedophilia and its relation to actual paedophilia, and given the existence of a category of child sex offender who, evidence suggests, does not go on to engage in hands-on molestation or any form of online sexual contact, and given the need to be cautious when extrapolating from evidence showing an escalation from child pornography to molestation (among certain groups) to the claim that this supports the view that engaging in virtual paedophilia is significantly more likely to result in actual paedophilia or even some other harm (e.g. child pornography), one must be sceptical over the merits of the 'significant likelihood' argument. Given the lack of substantive

evidence in support of a causal link or even correlation between engaging in virtual child pornography and actual child sexual assault, even with the caveat of accepting that more research needs to be done, perhaps we have no alternative but to accept Malamuth and Huppin's (2007) conclusion that current data (regarding such a relationship) do not support a "blanket prohibition against the use of virtual child pornography" (p. 827) and, as a consequence, endorse William's (2004) view that to control an activity on this basis of a *belief* in a relationship that current data do not support "is not logical" (p. 253).[3]

For Williams (and as noted earlier), arguments about the protection of children from sexual harm seem, instead, to have at their heart concerns about maintaining the proscribed moral standards of decency. Ryder (2003) goes even further, stating with reference to virtual child pornography:

> [The US child pornography] law causes harm to society by suppressing thoughts and expression concerning child and youth sexuality that involved no harm in production, fall short of advocating harm and that have at best a tenuous connection to the commission of harmful acts. The child pornography offence criminalizes a range of creative expression in the absence of any persuasive evidence of a risk of *harm*. (Ryder 2003, p. 103; emphasis added)

Despite what our intuition may be guiding us to believe, then, at present, there is insufficient evidence to support the view that engaging in virtual paedophilia is significantly likely to result in harm (actual paedophilia or otherwise), This in itself is enough to undermine the 'significant likelihood' argument. Nevertheless, Luck's challenge to this approach continues; although I am less convinced of the merits of this move, as we will see in the next section.

2.5 Argument in Favour of Virtual Paedophilia

> If this were enough to feed and satisfy their sexual desire, then pseudo-images might be seen as having social utility even if most of us would be wholly disgusted by their existence and the use made of them by the paedophile. (Williams 2004, p. 253)

As part of Luck's rejection of an appeal to significant likelihoods, he claims the following: "this argument [the 'significant likelihood' argument]

allows for situations where not committing acts of virtual paedophilia might be immoral" (2009, p. 33). Luck defends the claim with the following example:

> [I]magine a person who has overwhelming urges to commit paedophilia. She can satisfy these urges if she were to indulge in virtual paedophilia. However, if she does not, she will be driven to commit acts of actual paedophilia. (2009, p. 33)

I agree with Luck that an appeal to significant likelihoods should be rejected; nevertheless, I disagree with his claim that the 'significant likelihood' argument *allows* the possibility that not permitting acts of virtual paedophilia may in fact be immoral (immorality through omission, presumably), and therefore *allows* that one may be morally obliged to permit the enactment of child abuse.

Still in keeping with a broadly consequentialist approach, Luck's argument in favour of the moral worth of virtual paedophilia can be presented as follows:

1) S has an urge to do P;
2) Doing P is harmful;
3) If S does VP then S will not do P;
4) S should do VP.

If one's motivation is to avoid that which is harmful then, if one *ought* to do VP as (4) indicates, VP should be something that is either not harmful or not as harmful as doing P. Returning to the 'significant likelihood' argument: the basis for this argument is that engaging in virtual paedophilia (in effect, doing VP) is significantly likely to result in harm (actual paedophilia), and this 'fact' is what makes VP immoral (in accordance with a consequentialist approach); although we have seen that, at present, there is scant empirical evidence supporting this claim. Given what the 'significant likelihood' argument is promoting, it is difficult to see how it would allow (*qua* morally oblige us to permit) virtual paedophilia. To understand why not, consider the following:

   I. S has an urge to do P;
  II. P is immoral;
 III. Doing VP is immoral;

    IV. If S does VP then S will not do P;
    V. S should do VP.

When presented in this way, it is difficult to see why S *should* do VP, as doing VP simply involves substituting one immoral act for another. Yet, apparently, the reason why we should allow S to do VP is because it will prevent S from doing P (actual paedophilia). But if this is the case then an inconsistency is revealed; and it is because of this inconsistency that the 'significant likelihood' argument cannot allow (*qua* morally oblige us to permit) virtual paedophilia. For on the one hand, and according to premise (ii), VP is significantly likely to result in harm (actual paedophilia); but, on the other hand, we are told that the argument allows that we should permit virtual paedophilia on those occasions when it is significantly likely to *prevent* actual paedophilia by acting as a substitute. Irrespective of the truth of this claim (i.e. irrespective of whether virtual paedophilia is or is not effective at preventing or significantly reducing the likelihood of actual paedophilia), given premise (ii), my point is this: what is being claimed – namely, the significant likelihood that actual paedophilia will be prevented through the enactment of virtual paedophilia – could not be allowed by the 'significant likelihood' *argument* because it fundamentally undermines premise (ii).

    In the guise of a hypothetical coping strategy for paedophiles, Luck suggests that virtual paedophilia could be used as a more acceptable outlet for unacceptable desires. Presumably, this hypothetical coping strategy is something that the paedophile chooses to engage in when not incarcerated. (If incarcerated, why should we permit an individual to satisfy vicariously the very desire that led – let us assume – to their incarceration in the first place?) What is implied within Luck's description of the paedophile enacting paedophilic activities virtually is that this is something they agree to do (perhaps as part of a treatment programme). We are therefore asked to accept that the paedophile is able to engage in virtual paedophilia as a means of preventing themselves from abusing an actual child. This being the case, what is to prevent a person with no previous history of engaging in paedophilia from likewise restraining from such an activity? What makes this person (the non-paedophile) significantly likely to engage in actual paedophilia (in accordance with premise (ii)) as a result of virtually enacting it while the actual paedophile is not? If one's response is that the actual paedophile is trained to use the virtual enactment as a means of coping then this argument needs to be fleshed out in more detail.

In particular, it needs to be explained what this training would involve such that the non-paedophile would (again, in accordance with premise (ii)) find it very difficult to resist engaging in actual paedophilia as a result of dabbling with virtual paedophilia in the absence of this training.

It may be that Luck is making the point that a consequentialist approach, which grounds the 'significant likelihood' argument, would allow virtual paedophilia if it could be shown (empirically) to reduce the likelihood of those with paedophilic desires committing actual paedophilia. If this is the case then where is the *harm* in enacting virtual paedophilia? Essential this move is susceptible to the same challenge just presented. If allowing virtual paedophilia can prevent or reduce the likelihood of those with paedophilic desires committing actual paedophilia then why would the same virtual act result in an increased risk for those without a history of paedophilic desires? To say that the virtual act could result in different outcomes for different groups (reduce the risk for paedophiles, increase the risk for non-paedophiles) is not contradictory, of course; but it would require a clear rationale for why this should be; ideally, with empirical support.

Even if we accepted that there is no significant likelihood of committing actual paedophilia as a result of engaging in virtual paedophilia, perhaps there is some other harm (broadly construed) – say, in terms of offence – that results from the general permissibility of virtual paedophilia which outweighs the reduction in harm that (hypothetically) might be brought about through its use as a coping mechanism for those with paedophilic desires. Such a possibility would be taking us away from an appeal to significant likelihoods, of course; it is therefore not a possibility I intend to discuss further here. I will, however, return to it in Section 3.3.2 when discussing Patridge's (2011) argument based on incorrigible social meaning.

In concluding this chapter, what I hope to have shown, as indeed Luck does, is that an appeal to social convention is unable to resolve the gamer's dilemma. In addition, despite my disagreement with Luck over what the 'significant likelihood' argument will allow, I do agree with him that an appeal to significant likelihoods as a means of differentiating, morally, between virtual murder and virtual paedophilia is unsatisfactory and therefore should be rejected. It is unsatisfactory because, as things stand, there is insufficient empirical support for premise (ii) (that virtual paedophilia is significantly likely to result in harm). There is therefore insufficient empirical evidence to support the view that virtual paedophilia is more harmful than virtual murder

and (in accordance with the consequentialist approach adopted throughout this discussion) subsequently less moral. Time, then, to move on.

NOTES

1. *Yaoi* (or 'boy love') refers to a particular genre of manga depicting male homosexual (sexual) relations. This is often produced by and marketed at heterosexual women (McLelland 2005; McLelland and Yoo 2007).
2. Responding in this way does not mean that some further reason could not be found.
3. While Malamuth and Huppins are against a blanket prohibition, they do propose selective prohibition targeted at convicted sex offenders. Evidence indicates that child pornography is a significant risk factor for reoffending in those convicted of a child sex offence. They therefore suggest, as a precaution, that the possession of *virtual* child pornography should be outlawed for this group. As interesting a discussion point as this is, regarding legislation, it does not provide a means of resolving the gamer's dilemma and so will not be pursued here.

Motivation, Discrimination and Special Status: Luck's Further Attempts at Resolving the Dilemma

Abstract This chapter assesses Luck's remaining attempts at resolving the dilemma. It begins with an examination of player motivation, and dismisses the claim that those who engage in (hypothetical) virtual paedophilia necessarily do so because they enjoy the idea of actual paedophilia; offering alternative motivations as a rebuttal. It then challenges the claim that virtual paedophilia should be prohibited because it amounts to unfairly singling out a particular group for harm (i.e. children). While this argument has some merit, given that virtual child murder is permitted, it is not clear why virtual paedophilia should receive differential treatment. Finally, an argument appealing to the special status of children is rejected for similar reasons to those presented against the previous attempt at proffering a solution.

Keywords Prurient appeal · Simulating vice · Differing motivations · Random versus targeted murder · Incorrigible social meaning

In this chapter, I discuss Luck's three remaining attempts at resolving the gamer's dilemma, and the problems each faces. Like Luck, I too conclude that each is unconvincing.

3.1   Finding Pleasure in the *Idea* of Paedophilia

Luck's next attempt at resolving the gamer's dilemma focuses on player motivation, particularly with regard to enjoyment, but also examines the putative connection between motivation and harm; only, this time, self-harm rather than harm done to others. The initial argument Luck presents is based on an assumption – albeit one that he will eventually challenge – that someone who enjoys engaging in virtual paedophilia must do so because they find something pleasurable about the idea of actual paedophilia. *Mutatis mutandis*, this is not assumed to be the case with those who engage in virtual murder.

A similar argument is presented by Goldblatt (2012) when arguing the legal case against virtual child pornography.

> [A]s opposed to other forms of media that depict lawless action, but are distributed to the general public, virtual child pornography is a depiction and encouragement of lawless action that is presented to and sought out by a very narrow, specific audience that is likely to be stimulated to react to it. (p. 37)

Essentially, Goldblatt's argument rests on the idea (the assumption) that, mostly, it will be paedophiles who will engage with virtual paedophilia; they will make up the audience. As such, and this is Goldblatt's point, given the audience (who it is assumed already break the law) and given the content and its depiction of illegal activity, the producers of this material are in effect inciting their audience to *immanent lawless action*; something that would not be the case with a more general audience, and something that would not be the case, therefore, with enactments of virtual murder.

Contrary to the assumption Goldblatt's argument is built on, however, finding something pleasurable about the *idea* of paedophilia does not require that such a person has ever engaged in paedophilia, nor (*a priori*) does it necessitate, nor (*a posteriori*) does it make it significantly likely, that they will (recall discussion on this issue in Section 2.3).[1] Instead, the assumption is that one's motivation to engage in virtual paedophilia should be understood with reference to the pleasure the idea of actual paedophilia elicits. Arguably, this assumptions grounds the incredulity with which someone might ask (with reference to virtual paedophilia): why would anyone want to do *that*? In short, feeding the intuition that fuels the

incredulity is the lesser assumption (compared to that presented by Goldblatt) that anyone who enjoys engaging in virtual paedophilia must find the idea of actual paedophilia pleasurable. Should this (lesser) assumption find empirical support, alongside the contrary assumption regarding virtual murder (that this is not the case), then the distinction created could form the basis for selective *moral* prohibition and thus be a means of resolving the gamer's dilemma.

Even if we accept that an assumption of the lesser kind is being made, regarding the motivation for engaging in virtual paedophilia, what remains unclear – given this assumption – is how engaging in virtual paedophilia, specifically, is potentially harmful to oneself. To explain, if I already find the idea of actual paedophilia pleasurable, such that the enjoyment I get from engaging in virtual paedophilia stems from this pleasure, then, as things stand, it is unclear how virtual paedophilia could be said to be harming me above and beyond the harm an appeal to intuition (i.e. our assumption) would suggest I am already inflicting on myself by experiencing pleasure at the thought of actual paedophilia. In short, if one accepts the premise that deriving pleasure from the idea of paedophilia is harmful to oneself then what harm, *in addition to this*, would I be inflicting on myself by virtually enacting paedophilia? To address this question, Luck draws on the work of McCormick (2001) who, in turn, takes inspiration from Aristotle's virtue ethic.

3.1.1 Simulating Vice

According to Aristotle (1976), "Moral goodness . . . is the result of habit" (NE II I, 1103a12-14). We "become just by performing just acts, temperate by performing temperate ones, brave by performing brave ones" (1103b1-2). So, in understanding how to become a good person, we are directed to understand the relationship between our actions and our dispositions. It is through repeated performance of just or unjust acts that we likewise become disposed to be just or unjust people by indulging virtue or vice.

Applying Aristotle's virtue ethic to virtual paedophilia, a potential problem is revealed which McCormick tries to overcome. Aristotle holds that one becomes disposed to virtue or vice through repetition of virtue or vice. The more one repeats a virtuous act, the more one develops the habit of behaving in this virtuous way. *Mutatis mutandis*, the same can be said of vice. But it is yet to be established whether engaging in virtual paedophilia *is* immoral and therefore a vice. Arguably, virtual paedophilia is not a

vice – indeed, this is what we are trying to establish – and so, arguably, by enacting paedophilia one is not repeating and immoral act. Given this, it is unclear how repeating something that is arguably not immoral can lead one to be disposed to vice. As a way around this potential objection, McCormick (2001) argues that merely *simulating* vice contributes to the repetition that Aristotle refers to and therefore aids in the formation of a disposition towards the thing the enactment represents: in this case, paedophilia. By "participating in simulations of excessive, indulgent and wrongful acts ... you do harm to yourself in that you erode your virtue" (McCormick 2001, p. 285; cited in Luck 2009, pp. 33–34).

There are a number of issues raised by the discussion so far:

1. The initial assumption that one must derive some form of pleasure from engaging in virtual paedophilia, because one finds the idea of actual paedophilia pleasurable, seems unnecessary to McCormick's claim that simulating vice leads to a corrosion of virtue. It would appear that such corrosion is the product of repetition irrespective of motive.

I concede that one would need to explain why an individual would continue to engage in simulated vice if they did not gain some form of enjoyment from it. But even *if* enjoyment is necessary, this does not entail that one's enjoyment must be derived from the pleasure one receives from the idea of paedophilia. Indeed, this point will be addressed below when considering one of Luck challenges. Nevertheless, it remains the case that one's motivation and any (alleged) corrosion through repetition are independent of each other, such that repetition of simulated vice is said to be sufficient for harm to occur, irrespective of motive.

2. Given #1, is someone who engages in virtual paedophilia because they derive pleasure from the idea of paedophilia *more* susceptible to harm, or to a greater degree of harm, than someone who engages in virtual paedophilia for reasons other than they enjoy the idea of actual paedophilia?

As we have seen, an implication of McCormick's position is that both of these hypothetical individuals are vulnerable to harm through a corrosion of virtue because they repeatedly engage in virtual paedophilia. Simulating vice, for McCormick, is sufficient for harm, irrespective of

motive. If, ultimately, both individuals are harmed (*qua* both undergo a corrosion of virtue) then, arguably, the difference between them in terms of their respective motivations is something of a moot point. Nevertheless, one might consider it reasonable to wonder what a further consequence of this corrosion of virtue might be, besides harm to oneself. What, for example, might the harm one has inflicted on oneself lead to?

Again, one might conjecture that it could lead to someone deriving pleasure from the idea of actual paedophilia when previously they had not. Should this be the case then this person would be in a similar position to the individual who derives pleasure from the idea of paedophilia and engages in virtual paedophilia because of this. As such, there seems to be no additional consequence for the individual who already derives pleasure from the idea of actual paedophilia. One might therefore speculate further and proffer that as a direct consequence of the harm inflicted on this individual through simulated vice, things might escalate to the point of seeking sexual contact with a minor. Perhaps, ultimately, this is the outcome that awaits both individuals as a consequence of the repetition of virtual paedophilia, irrespective of the original motivation for doing so. Here, the harm to oneself eventually finds expression as a desire to harm another – a minor – which one seeks to satisfy. Of course, what we have here is conjecture heaped upon conjecture; and necessarily so given #3.

3. McCormick's claim is essentially an empirical matter.

I am not aware of any research supporting the connection between virtual paedophilia and harm to oneself (*qua* corrosion of virtue), nor for the claim that harm would be done to others (see Section 2.4.2). Therefore, at present, McCormick's assertion remains unsubstantiated and is, at best, indicative of an intuitive fear. Of more concern to resolving the gamer's dilemma, however, is this:

4. The enactment of virtual murder appears to be vulnerable to the same arguments that have been presented against virtual paedophilia.

The basis for this attempt to resolve the gamer's dilemma is that the assumption regarding the connection between virtual paedophilia and actual paedophilia, *mutatis mutandis*, is not assumed to be the case with

regard to virtual murder and actual murder. Arguments that have been directed against virtual paedophilia were assumed not to be applicable to virtual murder. If this assumption is shown to be false then, irrespective of the quality of the arguments directed against virtual paedophilia, where these same arguments are in fact proven to be applicable to virtual murder, a means of differentiation has been lost. Consequently, the gamer's dilemma would remain unresolved and we would continue to lack the means of differentiating between virtual paedophilia and virtual murder on moral grounds.

Essentially, #4 is indicative of the approach undertaken by Luck when showing why this particular attempt at resolving the gamer's dilemma is unsuccessful. We will now look at Luck's objection in more detail.

3.2   DIFFERENT MOTIVATIONS: ENJOYING THE COMPETITION RATHER THAN THE KILL

One needs to be cautious when scrutinizing a player's *willingness per se* to engage in the virtual act (Bartel 2015). Instead, the *reason* a person is willing to engage should be considered of greater moral interest than simply the fact that they do (or have) engaged, and should therefore be thought of as a more legitimate and hence productive aspect of one's moral appraisal. As such, following Luck, in this section I challenge the assumption that those who engage in virtual paedophilia must do so because they find the idea of actual paedophilia pleasurable. In Young (2013a), I present three motivations for engaging in a virtual act within a game. These motivations are not intended to be exhaustive, nor are they mutually exclusive; although it is my contention that each is sufficient. They are intended to help us understand what might motivate someone to enact a real-world taboo. The three motivations are as follows:

$M_{(strategic)}$: S engages in the virtual act because it benefits S's overall strategy, which is to win the game. As such, S does not desire to engage in the virtual act because of what it represents but, conversely, neither does S desire not to engage in it for this reason. Ultimately, winning the game is what S desires, and S construes the virtual act simply as a means of achieving this end.

$M_{(enjoyment)}$: S engages in the virtual act because S anticipates that it will be fun/thrilling. S anticipates that it will be fun/thrilling because the virtual act represents something that is taboo. In short, S desires to engage

in the virtual act because the symbolic violation of a real-world taboo, in virtue of it being an enactment of a taboo, is something S anticipates deriving enjoyment from.

$M_{(substitution)}$: S desires to engage in a particular real-world activity which happens to be taboo. This activity is represented by the virtual act. S therefore desires to engage in the virtual act not because it is taboo (as is the case in $M_{(enjoyment)}$) but because it represents the real-world activity S desires to engage in (which happens to be taboo). Enacting the real-world taboo affords S the opportunity to satisfy this desire, vicariously.

The player whose motivation is categorized as $M_{(strategic)}$ is effectively endorsing the amoralist position captured by the declaration, 'it's just a game'. As such, what is being enacted is beyond the realm of moral obligation.[2] There is certainly some truth to this assertion. After all, what is happening within a video game is literally nothing but the manipulation of pixels. As Klimmt et al. (2006) explain:

> Obviously, in violent video games no living creatures are harmed and no real objects are damaged. Dead bodies, blood, and injuries are nothing more than pixels. The non-reality status of video games can therefore be used to explain why moral concerns are not 'necessary', applicable, or rational in their context; there simply seems nothing to be 'real' in a game that moral concerns could arise from. (p. 313)

In the case of enacting virtual murder, there seems little intuitive appeal in the idea that those who engage in virtual murder do so because they derive some kind of pleasure from the idea of actual murder (and certainly there is no empirical support for this as a trend). In fact, empirically, there is support for the claim that those who engage in virtual murder or other violence do so for strategic reasons, as captured by $M_{(strategic)}$, out of a sense of competition (Adachi and Willoughby 2011; Griffiths et al. 2016). Glock and Kneer (2009), for example, when commenting on the findings of a study by Ladas (2003), note how gamers seemed "to focus on competition, success, thrill [indicative of $M_{(enjoyment)}$], and the virtual simulation of power and control rather than damaging other persons" (p. 153). Glock and Kneer consider this way of thinking about the game (notably, *not* in saliently aggressive terms) to be suggestive of the existence of *differentiated knowledge structures* in those with pro-longed violent game exposure when compared to novice gamers. It may

be, they surmise, that novice players associate violent video games with aggression because of media coverage to that effect; however, through "repeated exposure to violent digital games, links to game-specific concepts are strengthened, thereby overrunning [media-related] associations to aggression" (p. 153).

To illustrate at least the possibility of equivalence in the case of virtual paedophilia, Luck creates a scenario based on a fictitious video game whereby, for strategic reasons – and therefore, for reasons in keeping with $M_{(strategic)}$ – one might decide to engage in an act of virtual paedophilia. In the words of Luck:

> [I]magine you are playing a computer game, the object of which is to steal the Crown Jewels from the Tower of London. One way to achieve this goal is to seduce and sleep with a Beefeater's daughter, who just so happens to be 15. A player who commits this act of virtual paedophilia may do so, not because he enjoys the notion of having sex with a child, but because he wishes to complete the game. (2009, p. 34)

If a player's motivation best fits the category $M_{(strategic)}$ then this undermines the assumption that, in the case of virtual paedophilia, to engage in such an act, one *must* find the idea of actual paedophilia pleasurably. Contrasting virtual murder with virtual paedophilia, in the context of $M_{(strategic)}$, we get:

(a) S engages in virtual murder as a means to an end; it helps him/her progress through the game.
(b) S engages in virtual paedophilia as a means to an end; it helps him/her progress through the game.

Statements (a) and (b) provide equivalent motivations for engaging in each respective activity: motivations compatible with $M_{(strategic)}$. Such an outcome challenges the legitimacy of the assumption presented at the start of this chapter regarding the motivation for engaging in virtual paedophilia. Suppose, however, that the gamer admits that the reason they engage in virtual paedophilia is because it is fun/thrilling. Still contrasting with virtual murder, the following possibilities present themselves:

(c) S engages in virtual murder because it is fun/thrilling, irrespective of whether it helps S progress through the game.

(d) S engages in virtual paedophilia because it is fun/thrilling, irrespective of whether it helps S progress through the game.

Is there a sense in which engaging in virtual paedophilia might be deemed pleasurable that does not bolster the assumption that this is because one must derive pleasure from the idea of actual paedophilia? Before answering this question directly, let us redirect the question towards virtual violence, including murder. If one wished to argue against the idea that enacting virtual murder because it is fun necessitates that one derives pleasure from the idea of actual murder then how might one do this? In other words, what would such an argument look like? Should such an argument be forthcoming, could the same argument be applied to understand better statement (d) and therefore counter the assumption that enjoying virtual paedophilia means one must enjoy the idea of actual paedophilia?

3.2.1   The Thrill of Virtual Violence

When considering the appeal of violent video games and why people are drawn to them, Nys (2010) has the following to say, "Knowing that it is wrong is part of the fun and games. The thrill of such virtual actions is precisely that they transgress ethical boundaries" (p. 81). In keeping with Nys' comments, it is not inconceivable that enacting virtual violence holds a certain allure for some people; it is gratifying and pleasurable, such that many "identify with bad characters and enjoy committing or observing simulated immoral action" (Schulzke 2011, p. 63; see also Konijn and Hoorn 2005).

In fact, Juul (2005) holds that video games "are playgrounds where players can experiment with doing things they . . . would not normally do" (p. 193) which, in the context we are discussing, may well involve virtual murder. Jansz (2005) likewise describes video games as "private laboratories" (p. 231) within which gamers can engage with different emotions and identities in relative safety – relative to the actual world, that is – and invest in their own form of psychological exploration (see also Konijn et al. 2011). Such exploration might result in the player being both disgusted and thrilled by the virtual violence they enact (Rubenking and Lang 2014); all of which adds to their enjoyment and motivation to continue.

In essence, under the guidance of $M_{(enjoyment)}$, where one's goal is simply to have fun, irrespective of whether what one holds as fun is

congruent with facilitating one's progression through the game, if 'fun' constitutes doing (a), (b), (c) then one ought to do (in a practical rather than moral sense) (a), (b), (c). In the case of $M_{(enjoyment)}$, and with reference to virtual murder (but not exclusively so), the activity has symbolic transcendence insofar as it represents in one space that which is taboo in another. Moreover, the symbolic connection which transcends these two spaces presupposes a different psychological connection to that evident in $M_{(strategic)}$. The action is psychologically meaningful not only in terms of understanding what it represents but also as a motivation to engage in the activity in the first place: because it is fun in virtue of what it represents, or at least that is what one anticipates. In the case of virtual murder, or indeed the enactment of any real-world taboo, "an inquiry into [its] appeal will reveal that [the] enjoyment presupposes a moral awareness, and therefore that morality is included from the start" (Nys 2010, p. 81; emphasis in original). In accordance with $M_{(enjoyment)}$, then, for some, simulating virtual violence is appealing precisely because it involves enacting taboos and therefore violating an offline moral code.

If one can engage in virtual murder in accordance with $M_{(enjoyment)}$ – whereby the object of one's desire and reason for enjoyment is the enactment of a transgression (a real-world taboo) – then one is left to ask why this motivation can be employed in the case of virtual murder but not in virtual paedophilia. In other words, if one is willing to accept that, in the case of virtual murder, one can enjoy enacting this transgression precisely because it represents a *transgression* – without deriving pleasure from the idea of engaging in the transgression for real, then how can any unwillingness to accept the same possibility (that is, the same reason for enjoyment) be justified in the case of virtual paedophilia?

One may wish to appeal to $M_{(substitution)}$ in the case of virtual paedophilia and declare that, in such an instance, the individual's real motivation must in fact be to satisfy vicariously their desire to engage in actual paedophilia. The problem with this approach, as I am sure the reader has anticipated, is: (A) *a priori*, why *must* this be the case for virtual paedophilia and not for virtual murder? And (B), *a posteriori*, is there any support for this motivational differentiation? We cannot rely on an appeal to intuition, as such an appeal is hardly infallible; rather, any examination of the grounds for differentiation must look beyond this. When we do look, at least in relation to motivation, we find that

one cannot differentiate between virtual murder and virtual paedophilia based on (A) and (B), and therefore #4 – which states: the enactment of virtual murder appears to be vulnerable to the same arguments that have been presented against virtual paedophilia – holds in the absence of a convincing rebuttal.

To qualify the position regarding (B), there is at present no empirical support for such a motivational difference because there is a paucity of research on this issue, thereby making any empirically based conclusion impossible to draw. Of course, if one wished to pursue empirical research in this area then one would need to provide some sort of rationale – especially given (A) – for why a difference in the motivation underlying each virtual act is something future research would be expected to discover.

3.3   Unfairly Singling Out a Group for Harm

The focus of Luck's next attempt at resolving the gamer's dilemma is the idea that virtual paedophilia, like the act it represents, involves unfairly targeting for harm a single group or category of individuals: in this case, children. This is contrasted with random virtual murder which does not involve unfairly singling out a specific group for harm.

Luck is quick to point out that it is not at all apparent that paedophilia is worse (in term of the harm inflicted) than murder, even the murder of a child. As he declares, "given that most parents hope to minimize the amount of harm that might befall their children, it is not clear that they would prefer their child to be murdered rather than molested" (2009, p. 34). Given this, if we wish to minimize representations of intentional harm, it is not immediately obvious why virtual murder, which may *include* the murder of a child, should currently be permitted (e.g. *Fallout 1 & 2, Dying Light* and *No More Room in Hell* (where the children are zombies), *Deus Ex* and *Deus Ex: Invisible War* in which you can kill children in a school) but not virtual paedophilia; unless one considers the intentional targeting of a specific group or category of individuals to be an additional harm that should not be permitted. This is the view expressed by Luck:

> . . . although computer games which entail virtual murder may be socially acceptable, it is doubtful that a game involving, for example, only murdering Jews or homosexuals, would be tolerated. It seems therefore, that unfairly singling out a group for harm is, in itself, additionally harmful.

Subsequently, since virtual paedophilia not only represents a harmful act, but also singles out children as the recipients of this harm, it could be seen as more harmful than virtual murder (since virtual murder does not necessarily single out any particular group). (2009, p. 34)

I will address Luck's example of a video game involving the murder of a specific category of people in Sections 3.3.1 and 5.1.1. Before doing so, it is worth pointing out that there is a danger here of endorsing, or at least seeming to endorse, some form of harm calculus: the opposite, perhaps, of the hedonistic calculus suggested by Bentham (1830) as part of his utilitarian ethic. To illustrate, suppose one's initial view is that murder is more harmful than sexual assault. Here, A > B. In the case of *child* sexual assault, however (also known as molestation), let us say that one also believes that such an act necessitates the singling out (unfairly) of a specific group or category of individuals (children, in this case) for harm, thereby incurring an extra harm. With regard to molestation, then, one's belief concerning the amount of harm inflicted becomes: B (sexual assault) + C (the act of singling out, in this case, children) > A (murder).

In order to make sense of this move, we are left to ponder how much less harm is caused by sexual assault compared to murder, such that one would adopt the belief A > B in a manner inspired by some form of Bentham-style calculus. We would also be left to wonder how much more harm must be caused by the addition of what is, in effect, discrimination, in the context of molestation; at least if such a union is to amount to the accumulation of harm beyond that of murder, and therefore if one is to be justified in holding the view (B + C) > A.

Such an approach, which one might call *indeterminate calculus*, is unsatisfactory precisely because the putative values calculating harm are indeterminate. Nevertheless, the *abstract nature* of the calculus is able to account for Luck's intuition regarding the parents' preference for molestation over murder, at least where the murder of a child is *targeted*. In the case of paedophilia and targeted child murder, both of which involve the singling out of a particular group, paedophilia is known to us abstractly already as (B + C) and child murder becomes (A + C). In the case of the targeted murder of a child (effectively, A + C), where one endorses the view that A > B, then (A + C) > (B + C): a position in keeping with Luck's intuition regarding parental preference. Of course, where the child's murder is not targeted then we should conclude (B + C) > A. In other words, as a normative position, one would be forced to conclude that the

untargeted murder of a child is better than molestation because less harm is inflicted. This means that, morally, parents should prefer their child (or *a* child in the case of disinterested parties) to be molested rather than be the victim of a targeted murder, but prefer them to be the victim of an untargeted murder than molested, owing to alleged differences in harm inflicted. This is not a position I find coherent (both in regard to the parents' or the disinterested party's moral preference).

Moreover, it is legitimate to ask how *any* act of sexual assault can be seen as anything other than harmful. From the point of view of the victim, then, this means that the harm caused through an act of sexual assault cannot be tallied or in any useful way quantified so as to determine whether the act (of sexual assault) was more or less harmful depending on the context; depending, that is, on whether one targets an individual at random or directs one's interest exclusively towards members of a minority group. While I accept that this is a legitimate point to make, I still feel a case can be made for articulating the nature of additional harm beyond that incurred directly by the victim.

To do this, I will begin by considering the difference between random murder and targeted murder (i.e. at a particular minority group), by examining the suggestion made by Luck that a video game in which one can murder specific minority groups would not be tolerated. Now, it may be that Luck is correct when it comes to his description of social convention but, in terms of a normative position, what would justify the claim that one *ought* to be less tolerant of a video game in which one murdered members of a minority group compared to one in which random persons are targeted for violent assault, including murder? Conclusions drawn in the case of virtual murder will then be applied to virtual paedophilia.

3.3.1   Random Versus Targeted Virtual Murder

In Young (2013b), I present the following fictitious example of a video game:

> Suppose I . . . play a game in which I am able to target, harass and eventually kill individuals categorized in terms of their race/ethnicity, or even their gender, sexual preference, or religious beliefs: a game I will call R.A.C.I.S.T. (which stands for *Rage Against Community: Intercept, Segregate, Terminate*). (p. 76)

Let us contrast this fiction with another called *S.H.: Random Attack* (S.H. stands for *Sh#t Happens*). In *S.H.: Random Attack*, I am able to kill virtual characters at random – ordinary citizens from all walks of life; a fiction not dissimilar to any number of violent video games currently available, and therefore indicative of the current state of play (as noted in Section 1.1)

When playing *R.A.C.I.S.T.*, I select from the minority group menu *African-American*. After all, I have to select something in order to play the game. As part of the gameplay associated with the selection of this group, I am able to enslave my victims before hanging them from a tree whilst still in their leg irons and manacles, or chase them down with a pack of dogs before setting the dogs on the exhausted victim(s), and so on. In fact, the gameplays of both *S.H.: Random Attack* and *R.A.C.I.S.T.* make possible ever more elaborate, cruel and unusual ways to target, harass and eventually dispatch the respective virtual victims. The only difference of note between these two games is that the victims in *S.H.: Random Attack* are selected at random, whereas in *R.A.C.I.S.T.* they all belong to a targeted minority group. In essence, *R.A.C.I.S.T.* permits the enactment of two actions prohibited in the real world, murder and discrimination, unlike *S.H.: Random Attack* which 'permits' only murder. (To qualify this last remark, someone may play *S.H.: Random Attack* with the intention of targeting only a certain minority group, against the purpose of the gameplay. Such a possibility will be dealt with in Section 5.3.2.) The enactment of discrimination in the context of murder is therefore the key feature which differentiates these two video games. Let us consider each enactment in turn, based on what I will call random murder (RM) and targeted murder (TM).

There are two ways in which I would like to examine both of these forms of virtual murder: first, with regard to what the act *represents*; second, in terms of player *motivation*. Before proceeding, however, a note of clarification: RM_v refers to the *virtual* enactment of actual random murder (RM_a), whereas TM_v refers to the virtual enactment of actual targeted murder (TM_a).

Starting with player motivation:

(a) Does RM_v entail that the player is motivated to play the game because they delight in the idea of RM_a irrespective of what the gameplay promotes? No, it does not *entail* this.

As discussed in Section 3.2, the player may be motivated by $M_{(strategic)}$ or $M_{(enjoyment)}$. In terms of what the act represents:

(b) Does RM_v entail that the virtual content is promoting actual random murder (RM_a)? No, it does not *entail* this.

Given (a) and (b), let us consider targeted murder (TM_v) using African-Americans as the target group. Again, starting with player motivation:

(c) Does TM_v entail that the gamer is motivated to play the game because they delight in the idea of TM_a irrespective of what the gameplay promotes? No, it does not *entail* this; although, I accept that, in the case of TM_v, one may intuitively feel this to be so.

And, again, in terms of what the enactment represents:

(d) Does TM_v entail that the virtual content is promoting actual targeted murder (TM_a)? No, it does not *entail* this; although, again, I accept that, in the case of TM_v, one may intuitively feel this to be so, or more so than in the case of RM_v.

One might object, of course, to the reasoning shown above; declaring that it demonstrates only that in the case of virtual murder, irrespective of whether it is targeted or random, there is no *logical* connection between representing these acts and promoting what the enactment represents. Nor is there any logical connection between the enactment itself (and even enjoying the enactment) and being motivated to engage in this activity because one delights in the idea of carrying out murder for real. Declaring that one event does not necessarily follow from the other does little, therefore, to alleviate the intuition that Luck was alluding to when suggesting that a game like *R.A.C.I.S.T.* would be objected to much more vehemently than a game like *S.H.: Random Attack*.

One might also argue that, in targeting African-Americans (for example), the game designer/publishing company, even if not intentionally seeking to promote racial hatred, may well be misconstrued as doing this, and that such an accusation would not be an unreasonable one to make, even if not factually correct; or that the gamer is intentionally or

inadvertently supporting this view by playing the game, and in the latter case may even risk coming to delight in the idea of racial hatred through continued enactments (we touched on this last point in Section 3.1.1). Intuitively, one may feel that this is the case with targeted murder much more than in the example of random murder. Or one may simply view the targeting of minorities in this way as offensive and therefore in poor taste. To make the moral case more forceful, however, we need more than intuition; otherwise we are back where we started when discussing social convention in Section 2.1. A possible way forward is presented by Stephanie Patridge.

3.3.2   *Incorrigible Social Meaning*

Patridge (2011) argues that the meaning of representations, and whether these are or should be deemed offensive and, from this, morally reprehensible, is contingent on whether they have *incorrigible social meaning*. That is, on whether the content represents an association that has deep-rooted (actual) social meaning to members of a particular society, which may therefore be deemed offensive to certain members of that society, and even be morally and legally proscribed. She illustrates this with a fictitious example of a cartoon image of the US president Barack Obama eating a watermelon. The association of an African-American with a watermelon (and similar imagery), we are told, has "been used as a mechanism to insult and dehumanize African-Americans, and to bind racist Americans together through the practice of telling racially demeaning jokes" (Patridge 2011, p. 308; see also Brenick et al. 2007, for a discussion on perceived stereotypes in video games).

The representations and virtual enactments targeted by Patridge are those which were once held to be something of a social norm (e.g. institutionalized racism) within the USA (for example) but which are no longer viewed in the same way. What she seems less concerned with are actual morally/legally prohibited actions that have never been a social norm. This is alluded to by Patridge (2011) with reference to the game, *Mafia Wars*. Thus, she says, "The fact that we enjoy playing this game seems to say nothing at all by itself about our attitude towards organized crime" (p. 307). Organized crime, as far as I am aware, has never been established as an acceptable social norm in the USA. Therefore, what I take Patridge to be saying here is that if we enjoy playing a game that features organized crime, our enjoyment is not necessarily a sign of our

approval of organized crime. The same must be said of random murder, one presumes, owing to its lack of incorrigible social meaning.

To be fair, I do not take Patridge to be claiming that someone who plays a video game like *R.A.C.I.S.T.* is necessarily racist; rather, she is offering a reason why such a game would likely offend members of minority groups and even others, regardless of one's motivation for playing the game. She is therefore proffering a reason, *a priori*, why one *ought* not to play such a game.

In the case of random sexual assault (meaning not targeted at a specific age group or gender) and molestation (which is, by definition, specific to minors), the same lack of a logical connection evident in (a)–(d) when applied to targeted or random murder applies here, along with the same criticism that a lack of logical connection does not satisfy any moral intuition we may have regarding the inappropriateness of enacting paedophilia. In presenting her case for incorrigible social meaning, as a reason to avoid the virtual targeting of minority or specific groups, Patridge discusses the enactment of sexual assault (specifically rape) on women. She considers virtual rape to have incorrigible social meaning because of the "global history and current reality of women's oppression" (2011, p. 312). Therefore, anyone who knowingly ignores the incorrigible social meaning of certain video game content – such as the rape in games like *RapeLay* or *Custer's Revenge* – shows "an obvious vice of character" (2011, p. 310). Moreover:

> To insist that one's imagination is one's own private affair, detached from one's own actual commitments and similarly detached from the contextualized moral facts on the ground, amounts minimally, in this case, to a thumbing of one's nose at a requirement of solidarity with the victims of oppression. (2011)

We can see this as a response to the position adopted by Ryder in Section 2.4.2 in defence of one's freedom to imagine and fantasize (with the aim of creatively expressing oneself), and therefore against US child pornography law. For Patridge, minimally, a player who engages in virtual rape metaphorically thumbs his/her nose up at the requirement of solidarity with the victims of oppression; but more than this, there is the possibility that such a person exposes a *flaw* in their character. Through the idea of incorrigible social meaning, Patridge provides an argument against unfairly singling out for harm minority groups. Her argument is

compelling in the context of enactments of racial discrimination and the rape/sexual assault of women. What it does not seem able to challenge, however, is the acceptability of enacting the rape/sexual assault of, say, white, middle class, heterosexual men or, importantly, in the context of trying to resolve the gamer's dilemma, virtual paedophilia (however, see Section 5.1 for a continuation of Patridge's argument). The reason for this is that neither activity in the real world has ever been presented historically as a social norm, nor has there ever been a global history of oppression towards the male demographic just described or children. In short, neither enactment is of a real-world act that has incorrigible social meaning.

If the gamer's dilemma involved trying to differentiate, morally, between virtual rape and virtual murder or even between murder and targeted murder then maybe incorrigible social meaning would provide a way to differentiate between the two, at least in terms of accounting for our intuitions or the social convention that playing a rape game or a game like *R.A.C.I.S.T.* is morally wrong. It would seek to differentiate based on offence caused (or at least the likelihood of this), and Patridge's argument based on incorrigible social meaning articulates the reasons for the offence well. Of course, one might be more cautious about wanting to establish a normative position based on offence (or likelihood of offence) caused, however legitimate the reason for being offended might be.

In short, the argument for differentiation between murder and molestation, based on unfairly singling out a group for harm, does not work. In terms of increased harm, not only is this difficult to determine but, even when considered more abstractly, it leads to the incoherent conclusion that, morally, parents should prefer their child (or *a* child in the case of disinterested parties) to be molested rather than be the victim of a targeted murder, but prefer them to be the victim of an untargeted murder than be molested, owing to alleged differences in harm inflicted. And, finally, as we have just seen, incorrigible social meaning cannot explain why we should not enact paedophilia.

3.4   THE SPECIAL STATUS OF CHILDREN

As a means of overcoming problems identified with unfairly singling out a group for harm, Luck's next move is to consider children as a distinct category with a special status. After all, the argument goes, "children possess properties such as innocence, defencelessness, etc. . . .

which . . . make harming a child worse than harming an adult" (2009, p. 35). Transferring this idea to virtual enactments, one might argue that, owing to the special status of children, representations of molestation would cause more moral outrage than representations of adult murder. Therefore, as a means of curbing this moral outrage, one should differentiate between virtual murder and virtual paedophilia, prohibiting the latter even if we allow the former.

There are problems with this argument. As Luck points out, even allowing for the special status of children: "it does not follow that child molestation is more harmful than adult murder" (2009, p. 35). What index of harm should we use to establish this? (Recall, we have already encountered problems with the idea of a harm calculus.) Given this lack of clearly defined measurement, how might we justify the alleged difference in moral outrage between enactments of adult murder and molestation? Furthermore, and as touched on already: *prima facie*, child murder is more harmful than molestation (recall Luck's assumption regarding parental preference). Given this, and in accordance with the special status of children, enactments of child murder and molestation should be prohibited. Yet, as mentioned earlier, the current state of play in the UK and USA includes video games depicting child murder (e.g. *Fallout 1 & 2, Dying Light, No More Room in Hell, Deus Ex* and *Deus Ex: Invisible War*), although not exclusively so. It would seem that this discrepancy cannot be accounted for by the special status of children argument (although, it may be explained in a more practically sense by legislation regarding the legal status of virtual paedophilia). Therefore, once again, an attempt to resolve the gamer's dilemma has fallen short of its goal.

In the next chapter, I switch authors and consider an argument proposed by Christopher Bartel. Bartel (2012) presents three premises or propositions which ground his argument for differentiating between virtual paedophilia and virtual murder. He also argues for a moral difference based on indirect harm; interestingly, not to children but to women. It is towards an examination of this position that I now turn.

Notes

1. Bader (2003) argues (from a psychoanalytic perspective) that sexual fantasies do not necessarily indicate a desire for the object of one's fantasy (in this case, sex with a child); they could be motivated by other factors such as a sense of helplessness or guilt, among others.

2. Saying this does not negate the possibility that one will play in a way that conforms to certain moral principles. Doing so, however, is ultimately for strategic rather than moral reasons (e.g. one may benefit from adopting a particular moral approach by not incurring certain penalties that may hinder one's progression).

Virtual Paedophilia as Child Pornography, and Harm Done to Women: Bartel's Attempt at Resolving the Dilemma

Abstract This chapter examines Christopher Bartel's attempt at resolving the gamer's dilemma, in which he equates virtual paedophilia with child pornography; arguing that such virtual enactments are immoral, not because they necessarily harm children but because they indirectly harm women by eroticizing inequality. None of this can be said of virtual murder. A systematic appraisal of the premises on which Bartel grounds his argument finds them all problematic. In particular, the claim that virtual paedophilia is child pornography is contested because of a lack ontological equivalence (the former not being synonymous with child abuse). Moreover, attempting to resolve the dilemma by appealing to the indirect harm caused to women is criticized for failing to focus on the appropriate object of moral concern: children rather than women.

Keywords Virtual paedophilia and child abuse · Ontological equivalence · Eroticization of inequality

4.1   BARTEL'S THREE PROPOSITIONS

In his 2012 paper, Christopher Bartel makes the following claims (Bartel 2012, p. 14):

(C1) Virtual paedophilia amounts to child pornography as it necessarily involves the depiction of sexual acts involving children.

(C2) Virtual paedophilia is morally objectionable insofar as child por-
      nography is morally objectionable.
(C3) Virtual murder is distinct from virtual paedophilia as the latter
      necessarily involves child pornography while the former does not.

Bartel believes that the truth of each proposition provides the basis for
resolving the gamer's dilemma, although he acknowledges that a full
resolution will likely involve addressing a number of wider issues which
he does not discuss in detail. That issue aside, let us consider whether
Bartel's approach is at least pointing us in the right direction.

   In essence, Bartel's argument is this: virtual paedophilia is child porno-
graphy and child pornography is morally objectionable. As virtual murder
is not necessarily a form of child pornography, we have the basis for a
relevant moral distinction between virtual murder and virtual paedophilia:
the latter is morally objectionable, the former is not or is less so. Hence,
we have a way of (potentially) resolving the gamer's dilemma. There are,
of course, several aspects to Bartel's argument that need to be unpacked as
part of a more considered appraisal of his position. This I will do in the
sections to follow. First, however, a point of clarification is required. When
discussing virtual paedophilia and child pornography, Bartel refers to
depictions of sexual acts involving children. Luck and Ellerby (2013)
describe this as a narrow approach to resolving the gamer's dilemma: for
acts of virtual paedophilia do not need to be depicted within the gameplay
to have knowingly occurred. As they explain:

> . . . suppose a game allows players to approach virtual children, and after
> progressing through various bits of suggestive dialogue, they have a chance
> to initiate an instance of child molestation, upon which the game screen
> would fade to black and the game would recommence in such a way as to
> make clear that the act had occurred. Such a game might count as one in
> which players commit the act of virtual paedophilia, despite the fact that the
> act itself is never depicted. (Luck and Ellerby 2013, p. 231)

Luck and Ellerby's point is: even under the circumstances described, it is
likely (they claim) that many gamers would object to this part of the
gameplay, despite the fact that virtual paedophilia is not actually depicted.
Thus, they conjecture, a 'deeper reason' (Luck and Ellerby 2013) exists
for why one would object to virtual paedophilia but not virtual murder,
irrespective of any actual depiction. Before one can fully resolve the

gamer's dilemma, then, this 'deeper' reason needs to be articulated and evaluated. For now, however, I will restrict discussion to what Luck and Ellerby refer to as a narrower approach by focusing on *depictions* of virtual paedophilia (which is consistent with my use of the term in previous chapters). In Section 6.5, I will broaden the approach in order to tackle (briefly) the issue of implied paedophilia and related issues. By 'related issues' I mean such things as the virtual grooming of a minor, which seems to be occurring in Luck and Ellerby's example even if the virtual enactment of sexual assault is only implied. The virtual grooming of a minor is a further example of the potential enactment of a legally and morally proscribed action within a video game and so will need to be including within the wider debate on the gamer's dilemma.

In the meantime, within the sections to come, I will consider each of Bartel's three claims (C1–C3). It is my contention that each is problematic in its own right and therefore the argument Bartel presents to resolve the gamer's dilemma, which is based on these claims, is unsuccessful.

4.2   *Virtual* Paedophilia as Pornography

In an attempt to support the first of his claims (that virtual paedophilia amounts to child pornography as it necessarily involves the depiction of sexual acts involving children), Bartel presents us with the following hypothetical video game:

> . . . imagine a video game in which the gamer is allowed to voluntarily commit an act of virtual paedophilia and the act is graphically depicted. In such a case, the graphic depiction of a character – who is clearly depicted as an adult – engaging in sexual acts with another character – who is clearly depicted as a child – would count as an instance of child pornography. While these may be *virtual* instances of paedophilia, they are still *actual* instances of child pornography. (2012, p. 13; emphasis in original)

Bartel's definition of pornography is taken from Rea (2001, p. 134). Following Rea, an object acquires the ontological status of *being pornography* – if (a) the object is put to pornographic use (more on what this entails later), and (b) it is reasonable to believe that the object will be used as pornography, in accordance with point (a), by most of the audience for which it was produced (see Bartel 2012, p. 14). Condition (b) is important: for although something may be treated as pornography by an individual or even a group of people (in

accordance with condition (a)), it should not be considered pornography unless it is treated in this way by the majority of the object's intended audience. Thus, although a nude image published in a naturist magazine may be treated as pornography by some, it should not be labelled 'pornography' because (I assume) the majority of its intended audience do not treat it as such. Contrast this with an equivalent image published in, say, *Penthouse* or *Hustler* magazine.

Bartel also has the following to say about the motivation underlying one's willingness to engage in virtual paedophilia: "If gamers commit voluntary acts of virtual paedophilia, then presumably they do so because there is something about it that they like intrinsically" (2012, p. 14). One may be forgiven for thinking that Bartel's presumptive claim is similar to the assumption we addressed in Section 3.1 when examining the assertion that those who engage in virtual paedophilia do so because they enjoy the idea of actual paedophilia. Importantly, though, Bartel does not say this explicitly, and in fact does not need to make this connection at all. He may well accept that there could be some other intrinsic aspect of the virtual enactment that the gamer enjoys: related to the fact that it involves simulating a *taboo*, for example. Seeking enjoyment of this kind (even in the context of virtual paedophilia) would be in keeping with the motivation expressed by $M_{(enjoyment)}$, which does not require that one derive pleasure from the idea of actual paedophilia. What is important for Bartel, is that the gamer's motivation is compatible with Rea's definition of pornography; and it is arguably the case that, by adopting $M_{(enjoyment)}$, one finds some intrinsic quality of the *representation itself* sexually arousing without necessarily being aroused by the idea of engaging in what the representation is *of* (i.e. actual paedophilia). But even this possibility is not necessary. After all, and as Patridge (2013b) notes, it may be that my interest in some intrinsic quality of, say, a sex scene depicted within a film stems from its cinematic quality and/or the acting abilities of the performers, and so is not sexual in nature. *Mutatis mutandis*, the same could be said of a sequence within a video game depicting a sexual act, irrespective of the 'age' of the avatars within the gameplay.

Equally, one may derive enjoyment from the enactment because it amounts to the realization of a particular strategy one favours to progress through the game which, in this case, just so happens to involve virtual paedophilia (it signifies, for example, that one's strategy is working and one is achieving one's goal). Such an approach would be in keeping with both $M_{(enjoyment)}$ and $M_{(strategic)}$ (they are not, after all, mutually exclusive). Here, what is intrinsically enjoyable may well be

the fact that one is able to adopt *this* particular strategy in *this* particular context, as permitted by the game mechanics. Recall Luck's (2009) example of a fictitious game in which it is contrived that one can obtain one's goal of stealing the Crown Jewels by seducing and sleeping with the Beefeaters's 15-year-old daughter (see Section 3.2). With reference to this, Luck and Ellerby (2013) argue that it is not at all clear that such an act is being treated pornographically by the player; rather, and echoing my own argument, it could be that the gamer adopts this strategy simply as a means of achieving the goal of stealing the jewels, irrespective of whether they derive enjoyment from this. In short, then, engaging in virtual paedophilia does not *necessitate* that one is deriving enjoyment from some intrinsic feature of the representation (as M$_{(strategic)}$ attests) or, even if one is, that one is treating the act pornographically (M$_{(enjoyment)}$.permits many other reasons to enjoy the enactment).

Bartel can accept Luck and Ellerby's (as well as mine and Patridge's) point about different motivations for engaging in virtual paedophilia. Bartel simply has to emphasize the fact that in order to satisfy Rea's definition (regarding what it is for something to *be* pornography), one simply has to find it reasonable to believe that the majority of gamers who engage in the act of virtual paedophilia treat the enactment as a form of pornography, even if not all do, and therefore even if it does not logically follow that this must be the case. More specifically, in treating the object as pornography, the majority must satisfy the following criteria, *where applicable* (i.e. criterion (iii) is not applicable to single-player video games):

(i) *x* (the virtual object/event) is a token of some sort of communicative material.
(ii) S (the gamer) desires to be sexually aroused or gratified by the content of the communicative material.
(iii) Even if S believes that the content of the communicative material is intended to foster intimacy between S and the subject(s) of the communicative content, this belief is not among the reasons for S attending to the content of the communicative material.
(iv) If the desire described in (ii) was no longer among the reasons for S to attend to *x*'s content then S would have at most a weak desire to attend to it (criteria adapted from Rea 2001, p. 134).[1]

When deciding whether the virtual enactment is pornographic, the fact that it does not involve actual children is something of a moot point, for

the point is this: where the function of the virtual enactment satisfies Rea's (and therefore Bartel's) definition of pornography, the virtual act of paedophilia *is* pornography. Even if, commercially, such a video game is not presently available, such content is theoretically possible and that is all Bartel needs in order to *proceed* with his argument for resolving the gamer's dilemma.

4.3   CHILD PORNOGRAPHY IS CHILD ABUSE

Let us, for the sake of argument, accept that within a particular video game an act of virtual paedophilia occurs that satisfies Rea's definition of pornography (it is treated as pornography by the *majority* of its intended audience). In fact, for the sake of argument, let us allow, at least for now, that all hypothetical video games involving the enactment of paedophilia satisfy Rea's definition (see Section 5.1 for an example of where this is not necessarily the case). Under these circumstances, the following applies:

a) Virtual paedophilia is pornography;
b) *x* is an act of virtual paedophilia;
c) Therefore, *x* is pornography.

Notice how I refer to virtual paedophilia as pornography rather than *child* pornography. I do this first because Rea is interested in defining pornography in the absence of specific content, but also for a reason that will become clear as we progress. Now, one might respond to this, not unreasonably, by pointing out that if virtual paedophilia is pornography then it is pornography which necessarily involves the depiction of children (or at least one child), and that this depiction is intended to elicit sexual arousal from its audience. *Prima facie*, pornography involving children is child pornography. Indeed, this is part of Bartel's first claim (virtual paedophilia is child pornography because it necessarily involves the depiction of sexual acts involving *children*).

This is not an unreasonable position to adopt. Nevertheless, I have the following response in mind. I am willing to accept that child pornography typically involves the *depiction* of a sexual or sexualized act, but wish to consider more carefully Bartel's claim that it necessarily involves the depiction of *children*.[2] On the question of children, I am willing to accept that the depiction need not involve actual children in order for it to satisfy Rae's

definition of *pornography* (i.e. it could involve virtual entities in the form of children which are intended to elicit sexual arousal); but, importantly, it is my contention that while 'not involving *actual* children, only virtual ones' is not a barrier to the depiction being classified as pornography, it is a barrier to it being classified as *child* pornography. The reason for this is that I consider child pornography to be synonymous with child abuse.

Recall from Section 1.3.1 the view that child pornography can be and often is a record of serious sexual assault on young children (Adams 2010; Edwards 2000; Tate 1992). The 2007 *Convention on the Protection of Children Against Sexual Exploitation and Sexual Abuse*[3] agrees, stating in no uncertain terms that child pornography *is* sexual abuse (Kalim 2013). Likewise, Eneman et al. (2009) note how: "it is inextricably harmful to children...[because the] production of child pornography *requires* a child to be abused" (p. 5; emphasis added). Similarly, Mal Shervill, Assistant Commissioner of the Western Australia Police, has this to say, "for every [pornographic] image they download there's a child somewhere in the world who's defenceless and without a choice, being abused and degraded, so it is a form of child abuse" (cited in Simpson 2009, p. 255).

To my mind, child pornography and child abuse are *ontologically equivalent*. Given this, the following should apply:

d) Child pornography is child abuse;
e) *x* is child pornography;
f) Therefore, *x* is child abuse.

As further support for their ontological equivalence, while also alluding to child pornography's broader impact, Russell (2008) holds that "sexually explicit photographs of minors...document the abuse; contribute to the abuse; and are the purpose of the abuse" (p. 1484). This view is shared by the US Supreme Court who, in 1982, ruled that child pornography is intrinsically related to the sexual abuse of children both in terms of the abuse suffered by the victim during its creation and the continued injury inflicted on the victim through its publication and every subsequent viewing (Rogers 2009). Child pornography should therefore be thought of as synonymous with child abuse not just in a direct and narrow sense but, rather, in an all-encompassing sense, insofar as the manufacture and/or distribution or sharing and/or collecting and/or viewing of these images (in accordance with Rea's definition of pornography) means that the

children represented by these images have not only been abused directly but continue to suffer secondary harm in the form of psychological distress (Gillespie 2008; Palmer 2005). Such imagery is therefore "a crime not only against a *particular* child, but against *all* children" (Oswell 2006, p. 252; emphasis in original). Given this, as King (2008) makes clear, "the harm is so obvious that there seems to be little to say, and little need to say it" (p. 331).

Even if the image is not of a child being sexually abused (i.e. a photograph of a naked child in a bath or even a fully clothed child in a park), this does not detract from the fact that images of children used for pornographic purposes are still abusive to these children. Superimposing the image of a child onto another image, so that it appears to be engaged in a sexual act, is likewise exploitative: for although it is not a record of actual abuse, and in a sense is no more 'true' than a painting (Williams 2003), the child whose image is superimposed is still being exploited by the fact that the image purportedly shows them engaged in some form of sexual activity (Eneman et al. 2009).

Child pornography, in virtue of the fact that it involves actual children, is therefore *sufficient* for child abuse to have occurred, and for it to still be occurring to the children represented. In the case of virtual paedophilia, of course, no actual children are involved in the depiction, and so no actual child is abused. This being the case, if we equate virtual paedophilia with child pornography then we have a situation in which child pornography both does and does not involve the abuse of actual children. This would mean that, ontologically, child pornography both is and is not equivalent to child abuse.

Irrespective of current categorizations within legislation – in which virtual paedophilia is typically classified as child pornography – and given that my interest is in the ontological and moral status of virtual paedophilia and not its legal classification, it is my contention that the occurrence or not of child abuse should count as a pertinent ontological and, importantly, moral distinction, and therefore constitute a relevant means of moral discrimination in the cases we are discussing. Given this, we have a means of morally discriminating between child pornography and virtual paedophilia in virtue of the latter's lack of ontological equivalence to child abuse. It should not be difficult to accommodate this distinction into our discussion while accepting that virtual paedophilia is capable of satisfying Rea's definition of pornography. Somewhat unremarkably, we simply distinguish between child

pornography and virtual child pornography and equate virtual paedophilia with the latter and not the former; either that or we are forced to use a phrase like 'non-abusive child pornography' which does not seem appropriate, and in fact seems ripe for misinterpretation. In light of my proposal, consider the following:

(1) Child pornography is child abuse.
(2) Non-child pornography is not child abuse.
(3) Virtual child pornography is a form of non-child pornography.
(4) Given (2) and (3), virtual child pornography is not child abuse.

I accept that the claim "virtual child pornography is a form of non-child pornography" may seem like an odd thing to say; it does however makes sense if one thinks of it as part of a broader classification of pornography said to be homogenous *only* insofar as none of it involves images of actual children, even if, in the case of virtual pornography, it depicts (*inter alia*) child sexual activities. With virtual child pornography, the depiction can be *of* child sexual abuse (insofar as that is what the image is meant to depict) without the depiction itself being a record of actual abuse. In fact, one might liken it to what the French philosopher, Baudrillard (1983), refers to as *simulacrum*, which he takes to mean a copy or a representation of a thing that has no original. The photograph is a copy of some original event or object. The computer-generated image can depict something that does not exist, other than as the depiction; yet we take it to represent something beyond itself.

To ease the awkwardness of the labelling I am using here (and its somewhat clunky fit), let us think of all forms of pornography, other than pornography involving actual children, as non-child pornography, and refer to this simply as pornography. If we do this then any labelling or ontological confusing dissipates (at least when contrasting this much broader category with child pornography, specifically), as we can see:

(5) Child pornography is child abuse (it is a sufficient condition for abuse to have occurred or still be occurring).
(6) Pornography does not amount to child abuse.
(7) Virtual child pornography is a form of pornography (and not a form of child pornography).
(8) Given (6) and (7), virtual child pornography is not child abuse.

The conclusion found in (8) – that virtual child pornography does not amount to child abuse – is a point not lost on Bartel. As he acknowledges:

> The worry is that if virtual paedophilia is to count as child pornography, then we must admit that it is virtual child pornography – that is, it is not the depiction of actual children engaging in sexual acts, rather it is the depiction of computer-generated virtual children. One might think that, as no actual child is involved in virtual child pornography, then at least the immoral status of such pornography is diminished, and perhaps may even be morally permissible. (2012, p. 15)

If we accept that virtual paedophilia can amount to a form of pornography (but not necessarily so), whose content necessarily involves computer-generated children engaged in sexual activities, and further accept that virtual paedophilia does not involve child abuse and, for this reason and contra Bartel, accept that it is not a form of *child* pornography (which necessarily involves child abuse), then Bartel's first claim (C1) can be rejected. It also means that Bartel's second claim (C2) – that virtual paedophilia is morally objectionable insofar as child pornography is morally objectionable – is problematic and in need of closer examination.

Before discussing (C2), another point of clarification is required. A possible response to my argument – that virtual child pornography while being a form of pornography is not child pornography – is this: it does not seem unreasonable to suggest that many people would find the idea of engaging with virtual child pornography more morally objectionable than engaging with mainstream pornography (involving consenting adults, for example, whether virtual or actual).[4] In response to this, I have the following to say: my claim that virtual child pornography is morally distinguishable from child pornography is not to say that the former cannot be distinguished, morally, from mainstream pornography. Making the ontological claim that virtual child pornography is equivalent to mainstream pornography therefore needs to be understood (and is possibly only true) in the context of child abuse and what counts as a sufficient condition for this. What I am saying, then, is simply that mainstream pornography and virtual child pornography share the fact that they do not constitute child abuse (and so are not sufficient conditions for child abuse to occur). Accepting this, however, does not negate the possibility that ontological and moral differences

can be found between mainstream pornography and virtual child pornography when discussed in different contexts: ones that do not involve child abuse, for example.

4.3.1 Contesting the Moral Equivalence of Virtual and Actual Child Pornography

Returning to (C2), the problematic status of this claim is made apparent when one considers that Bartel accepts that "there is no reason to think that virtual child pornography harms actual children" (2012, p. 15). In contrast, we have plenty of reason to think that actual child pornography *does* harm actual children (which is why it is said to be sufficient for child abuse to have occurred or still be occurring). Highlighted, again, is a moral discrepancy between two forms of pornography Bartel wishes to classify as the same. Actual child pornography is morally repugnant, first and foremost, because it involves child abuse. Virtual paedophilia (*qua* virtual child pornography) does not involve actual child abuse, nor at present is there evidence indicating a direct link between it and actual molestation (again, recall the judgement of the US Supreme Court presented in Section 1.3.1). Given these facts, the primary moral objection to virtual child pornography (yet to be established) cannot be based on the same primary reason for one's moral objection to child pornography (that it amounts to child abuse), irrespective of whether they share additional moral reasons to object to their occurrence.

In C2, Bartel uses the term "insofar as", meaning "to the extent that". Saying that *x* is morally objectionable to the extent that *y* is does not necessarily mean that they are morally objectionable for the same *reason(s)*, of course; and certainly, such a claim would be problematic in the case of virtual and actual child pornography, as already noted. Instead, we could take Bartel to be saying simply that virtual child pornography and actual child pornography are morally objectionable to the same extent: that is, equally repugnant. If this is the case, then one is left to wonder why Bartel needs to establish his first claim: that virtual paedophilia is a form of child pornography. The need for the same classification (ontological equivalence) would make more sense if one wished to posit the same moral objection to each type of depiction. In other words, if one wished to maintain that they are morally repugnant for the same reason(s) because they amount to the same

thing. Certainly, this would fit with the idea that the phrases "insofar as" and "to the extent that" imply some kind of connection between the two points of comparison, whereby one is in some sense *en par* with the other. Wishing merely to show that their level of moral repugnance is equivalent, however, does not necessitate that the two forms of pornography are ontologically the same.

Bartel seems to be adopting this latter strategy (although his position is ambiguous). I say this because I find it hard to believe that the reason he gives for the moral objection to virtual paedophilia (*qua* virtual child pornography) – which we will discuss in the next section – could be exhaustively the same as the reason he would give (if he were to discuss it) for why we should morally object to actual child pornography. That said, I would find it equally difficult to accept that his reason for objecting to virtual paedophilia, which (as noted) I imagine is different to his objection to actual child pornography, could nevertheless bestow on the latter the same level of moral repugnance as he would (and indeed we should) bestow on actual child pornography. In short, do we really want a form of pornography that does not involve child sexual abuse, or in fact any abuse, to merit the same level of moral objection, and therefore to be judged morally equivalent, to pornography that necessarily does? If we would not object, morally, to mainstream pornography (involving consenting adults) to the *same degree* (if at all) as child pornography (a not unreasonable assertion to make) then why should we object, morally, to an instance of virtual paedophilia to the same degree as we would (should) child pornography?

Perhaps, the best Bartel can hope for is that we accept his argument (discussed below) is able to establish *a* legitimate moral objection to virtual paedophilia without it being for the same reason as (I take to be) the primary objection to child pornography (i.e. child abuse), or without the need to accept that their immoral status is equivalent. After all, Bartel's aim is to find a legitimate means of discriminating morally between virtual paedophilia and virtual murder. *Pace* Bartel, it remains to be seen whether a resolution to the gamer's dilemma has to involve virtual paedophilia being classified as "*actual* instances of child pornography" (Bartel 2012, p. 13; emphasis in original); certainly, this identity relation has been challenged and attempts have been made to resolve the dilemma without classifying virtual paedophilia in this way, as we will see in the section below and in the next chapter.

4.4   THE EROTICIZATION OF INEQUALITY: BARTEL'S MORAL OBJECTION TO VIRTUAL PAEDOPHILIA

Consider the following deduction:

(g) Child pornography is morally objectionable;
(h) x (virtual paedophilia) is not child pornography;
(i) Therefore, x is not morally objectionable.

The conclusion in (i) does not necessarily follow from premises (g) and (h). As such, and reiterating the point made at the end of the last section, Bartel does not need to equate virtual paedophilia with child pornography in order to argue that it is morally objectionable or even that there is some relevant moral distinction to be made between it and virtual murder. In fact, Patridge (2013b) goes a step further and argues that the moral objection Bartel raises regarding the eroticization of inequality (see below) does not require that virtual paedophilia be classified as pornography at all, let alone child pornography. Patridge holds that we cannot classify potential instances of virtual paedophilia as pornography unless we know more about their intended audience and how this audience treats the virtual content (in keeping with Rea's definition). I accept Patridge's point but I am prepared to continue with my hypothetical position (introduced in Section 4.3) whereby instances of virtual paedophilia do (for the sake of the present argument) satisfy Rea's definition. This small difference between Patridge and myself should not detract from what I consider to be a shared view regarding certain problems with Bartel's moral argument against the permissibility of virtual paedophilia.

Before continuing, a further point of clarification. When discussing Bartel's position, primarily through the theorist he draws from (namely, Neil Levy), in order to be consistent with their shared position, I will refer to virtual paedophilia as a form of child pornography, even though I have rejected this identity relation. To be clear, then, when referring to virtual paedophilia as child pornography in the discussion to follow, I do so only to remain consistent with Bartel's (and Levy's) terminology and only as a means of showing how Levy develops his argument. Doing this should not be interpreted as some kind of tacit endorsement of his view. What I intend to show is that the objection Bartel raises (through an endorsement of Levy's argument) can in fact be made without classifying virtual paedophilia as a form of child pornography.

4.4.1   Levy's Argument for the Eroticization of Inequality

In order to maintain that virtual paedophilia, in virtue of being a form of child pornography, is morally impermissible, even though it does not harm actual children either in its manufacture or (given the evidence to date) in terms of increasing the risk of actual child abuse, Bartel adopts a position put forward by Levy (2002). In essence, Levy presents the following argument:

> [Feminists] have criticized pornography...on the grounds that it is the *eroticization* of inequality...It encourages both men and women to think of women as naturally inferior...But *child* pornography, actual or virtual, cannot depict children as equal participants in sexual activity with adults, nor can it establish a relation of equality between the adult viewer and the viewed child. Children are not equal; this is not a contingent fact about our social relations but a reflection of their physical, mental and psychological immaturity. For that reason, sexualizing children for adult viewers is *necessarily* sexualizing inequality. Child pornography is an extension of mainstream sexual relations, which are contingently unequal, into new arenas...But since child pornography is necessarily an eroticization of inequality, allowing it undermines efforts to forge this new sexuality [the eroticization of *equality* between men and women]. Perhaps, then, it is because of harm to actual women, and not children, that virtual child pornography is objectionable. (2002, p. 322; emphasis in original)

What is important to note within Levy's argument is that both virtual and actual child pornography promote the eroticization of inequality. This is necessarily so in the case of pornography involving children (whether actual or virtual) because children are necessarily unequal to adults. In contrast, the unequal status promoted within much mainstream adult pornography (which depicts women as dominated, and where sexual fulfilment for both males and females can only be achieved if women adopt a position of subjugation) is a contingent (not a necessary) fact about sexual satisfaction and the status of women more generally. Virtual and actual child pornography is therefore complicit in maintaining this contingent unequal relation between men and women by further eroticizing inequality and so helping to maintain the current and contingent status quo.

It is also important to note that Levy presents the following conjecture at the end of the passage quoted above: it is because of harm to actual

women, and not children, that *virtual child pornography* is objectionable. When discussing the eroticization of inequality, Levy does not differentiate between virtual and actual child pornography. In the last sentence of the passage, however, he does. What Levy is implying is that there are other objections one can raise against actual child pornography, as I have discussed, which do not apply to virtual child pornography: namely, that it necessarily harms children and is a form of abuse. Given that this objection cannot be presented against virtual child pornography, Levy proffers what might be considered by some, including Bartel, to be an unorthodox, even surprising, approach: that virtual child pornography is morally objectionable because it harms women by further eroticizing inequality, thereby maintaining their unequal status among men.

Patridge (2013b) shows some sympathy for this view, holding that imagery involving sexual inequality, including virtual paedophilia, harms women because:

> ...it is deployed in a cultural climate in which women are systematically treated as unequal, and this inequality is achieved in large part by treating women as sexually unequal...[Moreover,]...any imagery that sexualizes inequality more generally [i.e. virtual paedophilia] will contribute to the larger cultural assumption that inequality is sexy and so is as things should be. (p. 29)

She does not share Bartel's view that Levy's argument is 'surprising', however, given that it finds support through the historical subjugation of women, including higher instances of sexual assault or sexual aggression on woman than men.[5] Indeed, Patridge's view echoes the re-occurring criticism of adult pornography: that it "serves to disseminate an untrue and damaging view of women, and ..., in doing so [, ...] supports sexist attitudes, reinforcing the oppression and exploitation of women" (King 2008, p. 335; see also MacKinnon 1991; Wilkinson 2011; Wright et al. 2016). The essence of this long-standing critique is captured by Longino (1995): "Because it is simply being female that, in the pornographic vision, justifies being violated, the lies of pornography are the lies about all women" (p. 39). Yet in the case of children, Patridge adds:

> ...children are not generally subjected to representations that sexualize them; in fact, it is quite the opposite. In the United States, for example, we have very little cultural tolerance for images of children that are

sexualized . . . I am not saying that this does not happen to individual children, it does. What I am saying is that children in general are not harmed in this way because in general we have very little tolerance for such treatment. In contrast, we seem to have quite a bit of tolerance for sexualizing women in ways that contribute to their oppression . . . It is for this reason, that it is very difficult to make the case that virtual sexualized images of children harm actual children in a way that would parallel the case that Levy makes about women . . . So, if making the moral case relies on making the case for harm, then it seems more promising to rely on a more remote harm, namely the harm to women. (2013b, pp. 29–30)

As far as Patridge is concerned, the case for virtual paedophilia harming children cannot be made because, culturally, we have little tolerance for sexualized images of children. That said, Elliot (1992) warns that we are becoming desensitized to, and therefore more tolerant of, inappropriate (sexualized) images of children through their increased use in advertising. Likewise Russell (2008) make the point that our culture, and indeed a number of others, is overflowing with images of sexualized youth (recall Britney Spears sexy school girl look, circa 2000); what Hartley (1998) calls *juvenation*[6] (see also Jewkes and Wykes 2005). Consequently, "[i]s it really so strange that these same images feature in people's fantasies?" (Russell 2008, p. 1499). There is therefore a danger that we may come to think of the sexualization of children, at least in the context of advertising or pop music, as normal. This, in turn, may support the paedophile in his belief that children are 'asking for sex' (Goode 2010; King 2008).

Such a view (increased tolerance for sexualized children) is not universally accepted, of course. Leaving that debate aside, one way to advocate a case for harm that avoids describing certain advertising campaigns as soft-core child pornography, as Elliot does, or treating child beauty pageants as similarly sexualized and exploitative, is to make a case for indirect harm: namely, as a further example of harm towards women (just as Levy and Bartel claim). It is indirect harm because the representations of abuse are not *of* women but children; yet, such representations act to reinforce the continued subjugation of women. Moreover, as alluded to earlier, while Patridge may accept that mainstream pornography, child pornography and virtual paedophilia eroticize inequality (even holding that these images necessarily do this),[7] she also holds that virtual paedophilia is able to do this without acquiring the ontological status 'pornography' (contra Levy and Bartel). In other words, while virtual paedophilia necessarily eroticizes

inequality, something can eroticize inequality (i.e. virtual paedophilia) without necessarily being classified as pornography. (Recall, I have no problem with this position, in principle, but can also envisage instances where virtual paedophilia within video games does satisfy the criteria for pornography. I do not anticipate Patridge objecting to this caveat.)

In light of Levy's objection to virtual child pornography, which Bartel adopts and Patridge is sympathetic to (at least in part; see below), what are we to make of Bartel's second claim (C2) (that virtual paedophilia is morally objectionable insofar as child pornography is morally objectionable)? Levy's moral argument against virtual paedophilia – that it eroticizes inequality – can be levelled against child pornography, but this does not mean, in my view (which I would say is an orthodox one), that Levy is claiming that this is the primary reason one should object to child pornography, although he is saying precisely this in the case of virtual paedophilia (because Levy's argument in set in the context of presenting a case for harm). Consequently, I do not believe it is controversial to say that the primary moral reason for objecting to child pornography is different to virtual paedophilia; I would even go so far as to say that it is necessarily different.

Interpreted in this way (whereby both forms of depiction are morally objectionable, only for different reasons), (C2) is sustainable, even when (C1) is false (as I have argued). It does, however, have the effect of making (C2) somewhat weaker and therefore less bold. Effectively, (C2) asserts simply that child pornography is morally objectionable, and so is virtual paedophilia. As such, it is much less of an asset when trying to resolve the gamer's dilemma, as it cannot co-opt the strength of moral repugnance typically directed at child pornography because the reason for this moral repugnance is not applicable.

In the next section, I examine Bartel's third claim (C3): that virtual murder is distinct from virtual paedophilia as the latter necessarily involves child pornography while the former does not. What is true about (C3), perhaps somewhat unremarkably, is that virtual murder does not necessarily involve child pornography. It is also the case that virtual murder does not necessarily involve pornography more generally (including virtual child pornography); although it could of course be the case that some enactments of murder within video games are able to satisfy Rea's definition of pornography (whether currently available or merely fictitious). The point is: it is not a necessary condition of virtual murder that they do. That said, the idea that virtual paedophilia necessarily involves child pornography has

been challenged. I have argued that it does not constitute child pornography at all; although I have conceded that it could satisfy Rea's definition or pornography more generally. Thus, while I do not hold that virtual paedophilia is *necessarily* pornography, I am prepared to allow that hypothetical examples can constitute pornography. I qualify my position in this way, not only as a nod to Patridge (2013b) but also because it is in keeping with my own discussion on different player motives.

Given that there is convincing argument against the assertion that virtual paedophilia amounts to child pornography, we could (indeed should) simply dismiss (C3). In order not to do this, it is necessary to make an adjustment which, I believe, still preserves Bartel's aim of identifying a morally relevant distinction between virtual paedophilia and virtual murder. Having made this adjustment, the question becomes: is the distinction between virtual paedophilia and virtual murder of a kind that is able to resolve the gamer's dilemma?

4.5   Is There a Morally Relevant Difference Between Virtual Paedophilia and Virtual Murder?

Adjusting (C3) so that it is not vulnerable to the arguments presented against virtual paedophilia as a form of child pornography, while still making it amenable to Levy's argument for eroticizing inequality, we get:

(C3*) Virtual murder is distinct from virtual paedophilia as the latter necessarily involves eroticizing inequality (irrespective of whether it is classified as pornography) while the former does not.

Underlying Bartel's third claim is the following assertion (taken from Luck and Ellerby 2013, p. 230):

(C3a) If an action is wrong for some reason, and another action is not wrong for this same reason, then there is a relevant moral distinction between the actions.

Given what (C3a) is claiming, suppose we accept, for the sake of argument, that Levy's objection, based on the eroticization of inequality, is something that applies to virtual paedophilia and not to virtual murder (as noted in C3*). This difference could then be presented *for consideration* as a morally relevant means of distinguishing between the

two virtual enactments, and therefore as a way of resolving the gamer's dilemma.

Is the distinction identified within (C3*) of moral relevance insofar as it provides the means of resolving the gamer's dilemma? No, I do not believe so. I say this for two reasons. First, consider each of the deductions below:

Deduction 1

  A) The eroticization of inequality is morally objectionable.
  B) *x* (*qua* virtual paedophilia) leads to the eroticization of inequality.
  C) Therefore, *x* is morally objectionable.

Deduction 2

  D) The eroticization of inequality is morally objectionable.
  E) *x* (*qua* virtual murder) does not leads to the eroticization of inequality.
  F) Therefore, *x* is not morally objectionable.

In accordance with (C3a), a distinction has been made between virtual murder and virtual paedophilia. However, in deduction 2, the conclusion (F) does not necessarily follow from (D) and (E). Where *x* equates to virtual murder, even if virtual murder does not lead to the eroticization of inequality and so cannot be said to be morally objectionable for *this* reason, it does not mean that virtual murder cannot be held as morally objectionable for some other reason. The eroticization of inequality is presented as sufficient but not necessary for a moral objection. To resolve the gamer's dilemma, Bartel needs to do more than show that virtual murder cannot be judged immoral for the same reason as virtual paedophilia (see Luck and Ellerby 2013, p. 233 for a similar argument). At the very least, he needs to show that it cannot be said to be *as* immoral as virtual paedophilia – for some other reason yet to be discussed – or, better still, that it cannot be judged immoral at all.

The second objection to (C3*) is found in Patridge (2013b). As we have seen, Patridge offers some support to Levy's and therefore Bartel's argument against virtual paedophilia based on the eroticization of

inequality. Nevertheless, she is not altogether convinced it is able to resolve the gamer's dilemma. As she explains:

> I do not think that those of us who are interested in resolving the gamer's dilemma as it is posed by Luck will be entirely satisfied with . . . Bartel's resolution. This is so because, rather than telling us what is distinctively wrong with . . . virtual child sexual assault, Bartel points us in the direction of an indirect harm, the harm that such images cause to some other kind of entity, namely women. It is precisely this move that makes Bartel's resolution less than satisfying . . . [T]hose of us who are interested in Luck's version of the gamer's dilemma feel its pull because we think that there is something particularly egregious about it specifically because it involves our virtually sexually assaulting children. Since, Bartel's analysis does not make essential reference to the role that children play in our moral assessment, his resolution seems to rely on the wrong kind of moral reason. (Patridge 2013b, p. 30)

While Patridge accepts that the eroticization of inequality provides a means of morally distinguishing between virtual paedophilia (or child sexual assault, as she prefers to call it) and virtual murder, she hesitates over whether it provides the right kind of moral reason to resolve the gamer's dilemma, precisely because it does not take as its object of moral concern children. Importantly, then, it is not enough simply for there to be a legitimate moral objection to virtual paedophilia, even if this objection cannot be applied to virtual murder; rather, the objection must have as its object of moral concern the right kind of object: in this case, children. Only, then, can the objection be proffered as a means of resolving the gamer's dilemma, and only if the reason for the moral difference between virtual paedophilia and virtual murder is able to show *why* the former is morally worse than the latter.

What is left unresolved, of course, is the question: what marks out the right sort of moral difference from the wrong sort in the context we are discussing? While Patridge does not refer to intuition in the quotation above, certainly she is relying on a shared sense of something being 'not quite right' about Bartel's use of Luck's argument to resolve the gamer's dilemma. Patridge does seem to be appealing (not unreasonably) to the idea (the shared intuition, perhaps?) that our moral scrutiny should have as its focus the fact that virtual paedophilia necessarily involves images and even enactments that depict child abuse; and that we should find this fact morally objectionable irrespective of whether the audience treats these

depictions/enactments as a form of pornography, and therefore irrespective of whether they amount to pornography (in accordance with Rea's definition). But as well as justifying virtual depictions of child abuse as the primary reason for our moral objection (which the eroticization of inequality fails to do), we have to show that this is not only a means of distinguishing between virtual paedophilia and virtual murder (a straightforward enough task) but justify why this difference is relevant to resolving the gamer's dilemma (as less straightforward task), especially if we concede that virtual murder is not immune to its own moral objections.

In sum, it is far from clear that the eroticization of inequality provides a suitably means of distinguishing, morally, between virtual paedophilia and virtual murder. This is because we have not (as yet) identified a suitable marker or means of measuring *relevant* moral differences. It is therefore difficult to assess with confidence whether the moral objection to virtual paedophilia presented in the form of the eroticization of inequality is sufficiently distinct or strong enough to differentiate it from any separate moral objection to virtual murder we may care to present, or to justify the claim that it is an objection that is targeted at the appropriate object. Certainly, there is reason to find a moral distinction based on the eroticization of inequality unconvincing for at least one if not both of the reason just given. But if we had to select just one then I would say that positing the eroticization of inequality as a primary moral objection to virtual paedophilia fails to convince because it misses the point, in that it does not have as its object of moral inquiry the fact that virtual paedophilia necessarily involves the depiction of (computer generated) children being sexually abused. Wishing to direct one's moral inquiry towards such depictions does not mean that a way of morally distinguishing between virtual paedophilia and virtual murder will be found, of course. That requires further critical discussion, but perhaps it is pointing us in the direction we need to go. Perhaps, but as things stand, the dilemma remains unresolved.

NOTES

1. Patridge (2013b) challenges (iv) with reference to her (2013a) work, *Exclusivism and evaluation: Art, erotica, and pornography.* This challenge need not concern us here, however.
2. To be clear, I am ignoring other modes of representation (e.g. audio) as, from the outset, my focus has been on visual depictions.

3. Available at http://www.coe.int/t/dg3/children/1in5/Source/Lanzarote%20Convention_EN.pdf. Accessed 9 August 2016.

4. I appreciate that in the case of computer-generated avatars, consent cannot be given. What I mean by consent in this context is therefore the appearance of consent within the gameplay.

5. Rape Crisis England and Wales http://rapecrisis.org.uk/statistics.php. Accessed 28 July 2016; RAINN https://www.rainn.org/statistics/victims-sexual-violence. Accessed 28 July 2016.

6. According to Hartley, juvenation is the practice of communicating with an audience through the medium of youthfulness.

7. To be clear, I am saying that she would argue that mainstream pornography, as it presents itself traditionally (which is a contingent fact), necessarily promotes the eroticization of inequality. This does not negate the possibility that future mainstream pornography may be more egalitarian in the way it presents the sexual act.

Targeting Morally Irrelevant Characteristics and the Need for Context: Further Attempts at Resolving the Dilemma

Abstract This chapter presents two recent and more promising attempts at resolving the dilemma proffered by Staphanie Patridge and Rami Ali. Patridge seeks to differentiate morally between virtual murder and paedophilia by arguing that the latter involves the targeting of individuals for harm based on non-morally relevant criteria (being children). This approach is, however, vulnerable to an objection in the form of game which allows one to sexually assault randomly (including children) rather than murder. A possible response to this can be found in Ali's argument for context: a mitigating narrative that could (depending on context) either permit or prohibit virtual murder and paedophilia. Ali, however, fails to articulate clearly why simulating murder or paedophilia for its own sake (without a narrative) should be morally objectionable.

Towards the end of Chapter 4 we discussed Patridge's (2013b) objection to Bartel's (2012) attempt at resolving the gamer's dilemma. Patridge argues that a morally relevant distinction based on the eroticization of inequality, while having some merit, ultimately misses the point. It is unconvincing because it has as its focus the wrong object of moral

concern: namely, women and not children. Does Patridge offer an alternative means of resolving the dilemma?

In this chapter, I will examine Patridge's approach to the gamer's dilemma: not only its merits but also why it fails. Following this, I will consider an attempt by Rami Ali to dissolve the dilemma by undermining a fundamental premise on which it is built. While admiring the originality of this approach, and even accepting his argument for the importance of context, I nevertheless remain unconvinced by Ali's solution, for reasons I discuss.

5.1   Patridge's Non-harm-Based Approach to Resolving the Dilemma

Patridge is dissatisfied with Bartel's attempt at a solution to the gamer's dilemma because it proffers only an indirect moral difference between virtual paedophilia and virtual murder. Patridge's aim is therefore to identify a moral difference that is directly relevant insofar as it concerns itself with pertinent characteristics of each of the contrasting virtual events, but, importantly, examines whether these characteristics are construed by the subject as representative of our "lived moral reality" (2013b, p. 32) or whether, instead, they are seen as a departure from it.

What is noticeable about Patridge's approach is that it does not concern itself directly with harm; that is, she does not try to resolve the gamer's dilemma by attempting to convince us that virtual paedophilia is in some way more harmful than virtual murder. This is not because she fails to see the relevance of harm within moral reasoning; rather, it is because she believes that non-harm-based moral resources offer not only a new and independent way to approach the gamer's dilemma but also subsequently a more germane and therefore direct means of resolving it. As such, one might reasonably take Patridge to be positing virtual paedophilia as a *harmless wrongdoing* (Feinberg 1988) or as *materially innocent* but *morally non-innocent* (McMahan 2006)[1] insofar as she accepts that the act is devoid of direct harm but still considers it to be morally wrong. I would say that this view is partly correct. Her approach does, however, allude to harm when, for example, she asks whether playing video games with certain content will send the wrong moral message (see also, Powers' 2003, discussion on *socially significant expression*). Her question implies concern over what one might come to think of as indirect, perhaps even

longer-term, *cultural harm* (McGlynn & Rackley 2009). Cultural harm, in the context we are discussing, might occur as a result of engendering within a society a trivializing attitude towards actual child abuse.[2] Barrowing from Oswell (2006), this is because, "[t]he ethical intensity of the virtual image lies precisely in its capacity to refer to a scene beyond itself" (p. 258). While the virtual image may not be a record of harm, for Patridge, how we respond to this image, including our willingness to engage with it in the first place, is meaningful. Patridge's ethical concern is therefore related to the virtual image's capacity to refer to something beyond itself (i.e. actual child sexual abuse). How should our willingness to enact paedophilia be understood? What should our (*qua* society's) moral attitude towards child sexual abuse be taken to be *if* we permit video games to include its enactment as part of the gameplay?

To help illustrate and contextualize her concern, Patridge presents a fictitious game called *Child Sexual Assault*.

> You find yourself at a party where a group of individuals is playing a fictitious game called *Child Sexual Assault*...In *Child Sexual Assault*, gamers are incentivized to virtually hunt down and sexually assault what appear to be very young children, both male and female...For their part, the group members do not seem to treat the video game as pornography in the sense that Rea uses this term. That is, they do not seem to be remotely turned on by the depictions. Instead, they are laughing and joking, and most of them seem to think that...the game is hilarious precisely because it is so morally transgressive. (Patridge 2013b, p. 31; emphasis in original)

As part of the scenario, she also contrives that the person looking on is asked to play. Through the use of the fictitious game and a motivation to play that is compatible with $M_{(enjoyment)}$ – insofar as the group's motivation appears to be based on the appeal of engaging in a taboo activity[3] rather than because its members find the idea of actual paedophilia appealing – she sets out to examine whether there are any direct moral resources that could be used to support one's refusal to play the game, or presumably rebuke someone for choosing to play.

Patridge dismisses the potential objection that playing *Child Sexual Assault* in this way might be misconstrued as directly promoting the idea that sexually assaulting *actual* children is fun or less of a moral concern than we have been led to believe. The reason for this (as we saw in Section 4.4.1) is because she holds that committing such acts against

actual children is a culturally entrenched taboo; therefore, it is unlikely that anyone (or at least the majority) would misconstrue the group's play behaviour as promoting this message. But Patridge still believes that there is something morally troubling about the game. Let us examine what she considers this to be.

5.1.1   Targeting Individuals Based on Morally Irrelevant Criteria

The basis for Patridge's non-harm-based objection is as follows. In *Child Sexual Assault*, or any game of that ilk, one is invited (insofar as it is a key aspect of the gameplay) not simply to assault sexually a character, but to do so *because* that character is (*qua* represents) a child. As such, one's actions within the game emulate, albeit virtually, immoral acts that occur in the real world. Given that *Child Sexual Assault* simulates real-world actions, it is said to reflect our moral reality. Of course, lots of video game content can be said to reflect our moral reality. A game in which one's character is part of a team of soldiers fighting and killing the 'enemy' would be one example, as would playing the part of an interrogator trying to obtain (through torture, if required) vital information about a terrorist plot that will endanger the lives of innocent civilians. In the first example, one is expected to target and *kill* characters because they are enemy soldiers; in the latter case, an individual is targeted and possibly *tortured* because they are a terrorist 'known' to be withholding vital information.

Young and Whitty (2011) argue that simulation of this kind is guided by – indeed, reflects the moral principle of – *sanctioned equivalence*. As they explain:

> In judging what constitutes a suitable topic for the gameplay, or at least in judging what is not totally inappropriate, one might be guided by the principle of *sanctioned equivalence*. Killing, for example, can occur in legitimate or illegitimate ways. A sanctioned equivalent of killing is state-authorized execution, or the death of combatants during a war. Torture has been justified in the past by legitimate authorities..., and in some cases still is; or at least its legitimate use is debated (in the ticking bomb scenario...). The unofficial "Law of the Sea" maintains that cannibalism is acceptable, or is at least tolerated, when one's life depends on it and the victim is already dead, or was selected through the mutually agreed drawing of lots. However, it is difficult to think of a sanctioned equivalent in the case of rape or necrophilia, or of cases in which one's life might depend on an act of incest or bestiality. Sanctioned equivalence differentiates between equivalent outcomes that are

either legitimate or illegitimate. All legitimate outcomes are viewed as *instrumental* – a means to an end. Actions that do not have sanctioned equivalence appear *pathological*, an end in themselves. (2011, p. 807; emphasis in original)

Targeting a child in order to engage them in sexual activity is a form of assault that has no sanctioned equivalence. For Patridge, the moral reality reflected within the game, *Child Sexual Assault*, is therefore one in which the individual is targeted based on some non-moral characteristic, and it is in virtue of this non-moral characteristic that immoral acts are performed on them. In contrast, targeting someone to be killed has a sanctioned equivalence if the individual targeted is, say, an enemy combatant. In this context, the combatant is targeted because of the morally relevant characteristic of being a lethal threat (Walzer 1977); and in dispatching such a threat – in virtue of this characteristic and in the context of conflict – the action is deemed to be justified. Where such moral reality forms part of the reasons for a player's action within a game (i.e. killing a combatant who represents the enemy), then the player is said to be *morally managing* their potentially quite 'bloody' and violent enactments (Klimmt et al. 2006, 2008; Whitty et al. 2011). In the case of *Child Sexual Assault*, however, one's moral position would be more difficult to manage (or so the argument goes), such that, if one were to continue playing the game, one would likely have to disengage morally from the activity (Hartmann et al. 2014; Hartmann and Vorderer 2010) rather than manage one's moral approach. The most common means of doing this (disengaging, morally) is to hold that what one is doing "is just a game" and therefore of no moral concern. But, for Patridge, by doing this, there is a risk that one would be guilty of conveying the wrong message. In the context of *Child Sexual Assault*, as noted, the message one might be interpreted as sending is unlikely to be "delight in the idea of actual paedophilia"; instead, it may well be something less direct, like: "It's okay to target individuals based on non-morally relevant characteristics". In the case of *Child Sexual Assault*, the non-morally relevant characteristic happens to be their age; but the appeal of Patridge's argument is broader than this. Instead, it has the potential to impact on the virtual targeting of all minority groups (and others), where targeting is based on non-moral characteristics of the person.

It is important to note that Patridge is attempting to do more than describe a psychological connection between the enactment and the

extent to which it reflects our moral reality, such that someone who interprets *Child Sexual Assault* as reflecting our moral reality will be less likely to want to play the game or find the enactments in any way amusing. Instead, she is proffering a *normative* account in which "not finding the gameplay amusing or in any way appropriate" is how one *ought* to react. Thus, the targeting of virtual characters/objects/events based on non-morally relevant criteria means "that there may be moral reasons to avoid responding positively to at least some putatively imaginative representations that do not rely on making the case for the harm of such responses" (2013b, p. 32). This means that we will lose a certain "interpretive flexibility" (2013b, p. 31) regarding certain virtual content, and so be in a position to rely much less on the amoralist retort that it is only a game.

As a means of resolving the gamer's dilemma, virtual paedophilia or child sexual assault (to use Patridge's preferred term) can be distinguished from virtual murder on the basis of a morally relevant characteristic. It is immoral to target, harass and in any way injure someone based on a non-morally relevant characteristic. This form of immoral discrimination is occurring (*qua* being enacted) through the virtual paedophilia portrayed within *Child Sexual Assault* (a point I will return to shortly) but not in the case of virtual murder, at least where the murderer targets his/her victims at random.

On this last point, it could be argued that victims targeted at random are in fact being targeted on the basis of some non-morally relevant criterion: namely, being conveniently located. If I were to play a game in which I murdered passers-by at random then, in a sense, these victims are being *targeted* (compared to non-passers-by). They are being targeted because of their location: the fact that they happen to be *there*; 'there' being the place where I chose to murder people. Typically, and certainly in this context, location is a non-morally relevant criterion by which to target someone for harm. While this is true; importantly, 'location' it is not a characteristic of the *person* being targeted. Where the criterion relates to some feature or characteristic of the person that is not morally-relevant (e.g. age or skin colour or sexual orientation; compared, say, to lethal threat in the case of the soldier) then we can make a morally relevant distinction.

But even if the non-morally relevant criterion constitutes some characteristic or feature of the individual, the extent to which targeting someone for harm because of this, *within a computer game*, conveys the wrong message and is therefore something one ought not to do is

itself open to debate. To illustrate, suppose I were to play a (fictitious) game in which I targeted for ridicule balding men. Morally wrong? What about people who have ginger hair? Or how about a game in which as part of the gameplay I am able to target women with blonde hair and exclude them from certain occupations because of their lack of intelligence (at least as depicted by the gameplay), or force them to take certain jobs for the same reason. What about a game in which it is not just blonde women discriminated against but all women?

Interestingly, and as we saw earlier, Patridge offers some insight into how we may judge virtual content to be conveying the wrong message in certain circumstances, and as less so or not at all in others. Recall from Section 3.3.2 her argument based on incorrigible social meaning. Patridge (2011) holds certain representations to be immoral if they are related to historical (and contingent) socially entrenched views concerning particular groups and/or practices: practices that were once the norm but are now outlawed. In the case of the (fictitious) game in which one discriminates against bald men or those with ginger hair, given (as far as I am aware) the lack of historically entrenched negative social attitude and behaviour towards these groups, the argument from incorrigible social meaning would hold that such enactment should not be morally prohibited, or at least should be thought of as less of a moral concern than targeting gender or sexual orientation, for example. I concede, however, that the example of discriminating against blonde women is less clear-cut; but accept that discriminating against all women within a video game is vulnerable to the same argument Patridge uses against the morality of virtual rape games.

Patridge's argument for the immorality of targeting groups based on morally irrelevant criteria, even when understood within the context of her earlier argument-based incorrigible social meaning, does offer some promise: both for the selective prohibition of video game content, in general, and as a means of resolving the gamer's dilemma, specifically. It is not, however, without its problems, as I shall now demonstrate.

5.2   Targeted (Child) Sexual Assault versus Ubiquitous Sexual Assault

When describing the fictitious video game, *Child Sexual Assault*, Patridge draws our attention (in a suggestive rather than explicit way) to an important aspect of the gameplay: that its *primary* purpose is to

hunt down and sexually assault children. She also describes the players as seemingly not treating the virtual content as pornography (as defined by Rea). In this way we have an example that differs from the hypothetical scenario discussed in earlier chapters which, for the sake of argument, allowed that the virtual enactment satisfied Rea's definition of pornography.

Let us therefore consider the implications for the gamer's dilemma of a video game in which child sexual assault *occurs* but this is not the primary purpose of the gameplay. If we accept Patridge's argument articulating why it is immoral within a video game to target individuals based on non-morally relevant characteristics (which, I admit, has a certain appeal) then we have before us a sufficient condition for the moral prohibition of virtual content. At first glance, at least if referring solely to the content of *Child Sexual Assault*, it would appear that Patridge's position shares certain similarities with an argument Luck presents in his 2009 paper (discussed in Section 3.3) which concerns unfairly targeting a particular group for harm. Patridge's approach is vulnerable to the same objection Luck himself raises against this kind of argument. I explored aspects of this objection when contrasting targeted and random murder (Section 3.3.1), and will develop it further here in relation to Patridge's example, specifically.

Suppose we amend *Child Sexual Assault* so that the purpose of the gameplay is to assault sexually as many people as possible. Let us call this game *Sexual Assault*, and allow that, typically, when gamers play the game they randomly assault as many characters of various ages as they can, from the very young to the very old, males and females alike, in order to accumulate as high a score as possible. We could even allow that extra points are awarded for more elaborate forms of sexual assault but, importantly, insist that this is not necessarily tied to the age or sex/gender of the virtual character being assaulted. We might even allow that the person being sexually assaulted does not have to be alive! In *Sexual Assault*, the only non-morally relevant criterion determining who is targeted is 'location', insofar as victims are targeted because they happen to be available to be targeted. Let us also make it the case that the gamer's virtual character is able to overpower easily even the strongest and fittest of the victims, so targeting 'vulnerability' does not offer the gamer any advantage.

The group of friends playing this game do not treat the content as pornography and can be heard laughing and joking and even applauding the ingenuity with which sexual assaults unfold within the gameplay.

Or course, one could target specific characters within the game, if one wished – based, say, on age, race, sex or sexual orientation– and if this were to occur then the actions of that particular player could be morally scrutinized in accordance with Patridge's argument. But what about the actions of the group of friends in the example we are now discussing? They target at random. So how should we judge the morality of enacting the random sexual assault of someone who happens to be a child and who was not therefore targeted because they are a child? We cannot argue that the child is being targeted for non-morally relevant reasons based on the features or characteristic of being a child, because this is not the case. Neither could we make the argument if the victim happened to be a woman, or a man, or dead for that matter.

By way of a response, consider what Patridge has to say about (what she calls) run-of-the-mill first-person shooter games in which individuals are targeted and killed at random, including children. Because the gamer in these games is not targeting individuals based on non-morally relevant characteristics of the person, she holds that it is reasonable to view such games "as a departure from rather than a reflection of real world moral concerns" (2013b, p. 33). Such gaming content seems far removed from what goes on in the real world and therefore what preoccupies us, morally (I say this while recognizing that random mass shooting and killings are not unheard of). This move away from the representation of our perceived moral reality is perhaps made all the more apparent in cases of first-person shooter games that target the undead or mutants, or aliens (for example).

Patridge contrasts a game involving random murder with one that targets particular minority groups based, say, on ethnicity, or sexual preference and so on. A game in which it is possible to target specific minority groups would, according to Patridge, be more morally problematic than one in which all citizens are potential targets for murder. To illustrate, recall the fictitious game *R.A.C.I.S.T.*, introduced in Section 3.3.1. When playing the game, on this occasion I select 'Jewish' from the menu. I then target, harass and eventually kill Jewish citizens, and only receive points for doing so (we could even allow that points are deducted if the 'wrong' – in this case, non-Jewish – individuals are killed).

In contrast to *R.A.C.I.S.T.*, suppose the first-person shooter game I decide to play enables me to target any citizen. This is reminiscent of the other fictitious video game I introduced alongside *R.A.C.I.S.T.* called *S.H.: Random Attack*. As we have seen, for Patridge, such a

game is not morally problematic, or not as morally problematic as *R.A.C. I.S.T.* Given this, let us compare *S.H.: Random Attack* with *Sexual Assault*. Other than the nature of the harm done to the virtual victims, there seems to be no difference between the two respective games. With this in mind, should there be any *moral* reason preventing the possibility of run-of-the-mill (*qua* random) sexual assault games being made available which cannot be applied to run-of-the-mill first-person shooter games? To be clear, each game would include the possibility of targeting exactly the same random group of individuals. If one were to object to the idea of a run-of-the-mill sexual assault game, then what would form the basis for this objection? It cannot be Patridge's argument for the immorality of targeting victims based on non-moral relevance (as is the case with *Child Sexual Assault*), as this does not apply here.

To be fair, Patridge is neither claiming to have resolved completely the gamer's dilemma, nor is that her aim; rather, and more modestly, she claims merely to have "sketched [a position that] is at least getting close to what explains the moral responses of those who find *Child Sexual Assault* morally disturbing independent of the harm that enjoying such a game might cause" (Patridge 2013b, p. 33; emphasis in original). Moreover, her position makes it easier for us to understand why those who see games like *Child Sexual Assault* as reflecting our moral reality "are likely to conceive of those who are capable of enjoying such representational content as at the very least morally distasteful and/or morally immature" (Patridge 2013b). As far as this has been her aim, I would say that Patridge has succeeded. She is able to provide a useful *sketch* of at least some of the descriptive (or psychological) elements that fuel the gamer's dilemma (perhaps even the most important one), at least in relation to the difference between virtual paedophilia and virtual murder, specifically. Whether the same argument is strong enough to form the basis for a *normative* position in relation to these two virtual events, as well as other more general enactments that contribute to a broader version of the gamer's dilemma (e.g. sexual assault), is, however, another matter. In short, conceding that Patridge helps explain why someone would find another's (alleged) laissez-faire approach to virtual paedophilia in poor moral taste is one thing, and not without its uses, but it is still a far cry from establishing a normative position that is able to resolve the gamer's dilemma, as illustrated by the fictitious example where sexual assault of a ubiquitous kind is virtually enacted (*which includes paedophilia* alongside all manner of other sexual assaults) compared to enactments of random murder.

In the next section, I consider the argument presented by Rami Ali in his 2015 paper, *A new solution to the gamer's dilemma*. Rami does not try to resolve the dilemma as much as he tries to dissolve it be undermining a fundamental assumption on which the dilemma is built. In doing so, he introduces a factor not yet examined (except perhaps fleetingly and indirectly in the case of manga images discussed in Chapter 2): namely, *context*. As with Patridge's argument, Ali's position is promising but is not without its problems, both in terms of resolving the gamer's dilemma and proffering a normative account of virtual content, as I shall demonstrate.

5.3   Ali: The Importance of Context

Ali (2015) challenges a part of the gamer's dilemma that hitherto has avoided direct examination: namely, the claim (assumption, in fact) that people *do* intuitively believe that virtual murder is morally permissible while virtual paedophilia is not (an assumption I questioned briefly in Section 2.1). Ali holds that such a claim amounts to little more than a generic, decontextualized comparison between two enactments of real-world prohibited behaviour and so offers little insight into the manner in which gamers morally judge virtual enactments or the perceived moral status of the content. In contrast, where context is provided, Ali argues, "it is neither the case that all acts of virtual murder are acceptable, nor that all acts of virtual pedophilia are unacceptable" (2015, p. 268). We have already seen how Patridge has morally criticized enactments of targeted murder, or at least given a reason for why (she claims) they are more vulnerable to moral criticism than enactments of random murder. The reader may also feel that certain depictions of murder, or other violence, even if random, may be a cause for moral concern: if say, the violence is particularly graphic and/or realistic and/or prolonged. Others may disagree, of course, and continue to claim that even this depends on the context in which it occurs.

5.3.1   Differentiating Between In-Game and Gamer's Contexts

According to Ali, in order to understand and therefore interpret correctly, a particular virtual enactment, we need to take account of two contexts. The first he calls the *in-game context*. Here, whether the killing of another virtual character is (*inter alia*) murder or self-defence is dependent on the context in which it occurs *within the game*. This, in turn, may be partly or

wholly dependent on the nature of the gameplay and the constraints placed on the narrative by the game mechanics. If, for example, one is playing a game in which one takes on the role of a combatant who is being attacked by enemy soldiers, and it is 'kill or be killed', then the in-game context demands that the killing of an enemy combatant is legitimized (as self-defence) under the rules of war (*jus in bello*). In contrast, if the same soldier intentionally killed a civilian who was not a threat then this would (should) be classified as murder within the game.

De Vane and Squire (2008, p. 267), suggest that experienced players "develop metacognitive understandings of how violence is represented" within the game: namely as instrumental to the success of the game, or even (for example) as immersed within a narrative that extols the principle of sanctioned equivalence (see Section 5.1.1). De Vane and Squire go on to note that the meaning that players derive from interaction with various media (such as violent video games) must therefore be contextualized.

On its own, the in-game context is incomplete and so is capable of only partially informing our moral inquiry concerning the virtual enactment. In addition, Ali tells us, we need to take into account the *gamer's context*. As he explains:

> While it is true that the gamer's contribution to the virtual world depends on what the contribution amounts to in that world, it is also true that the what the act amounts to in that world may be entirely irrelevant to the gamer's virtual performance. A fuller picture requires that we also attend to the context of the gamer performing the virtual acts. (2015, p. 269)

Ali makes the valid point that a gamer may be playing the game without a full, or perhaps only a very basic, understanding of the in-game context, or may have chosen to ignore this context completely. Ali illustrates the potential incongruence between in-game and gamer contexts with an example of a character called Drake (taken from the *Uncharted* video game series) who is a modern-day treasure hunter. Often, within the narrative/gameplay, he will be attacked. If he kills his attacker under these circumstances then this is judged to be self-defence (again, in accordance with the principle of sanctioned equivalence). Now, it may be that the gamer, when playing the game, does so as a means of enacting fantasies of murder. The gamer fantasizes that he (*qua* Drake) is murdering his (Drake's) assailants. The implications of this for our interpretation of the virtual killing are outlined by Ali:

On the one hand it is clear that Drake continues to perform the very same act in the game's fiction. If the act was one of self-defense, it continues to be so. However, since the gamer is not aware of the in-game context, and anyway would choose to disregard it if he was aware of it, it seems implausible to attribute virtual self-defense to him. Instead his act is plausibly one of virtual murder. What he is doing is virtually murdering, but the way he commits this act is through Drake's act of self-defense. (2015, p. 269)

5.3.2   *Appropriate Engagement*

Essentially, Ali is distinguishing between the (intended) narrative within the game and the motivation underlying the manner in which the gamer interacts with the gameplay, which (broadly construed) will be either congruent or incongruent with the narrative. Given the independence of these two contexts and the potential for a discrepancy between them, Ali proposes that we adopt an *appropriate engagement* view, whereby virtual acts "are individuated by *the gamer's appropriate engagement with the in-game context*" (2015, p. 270; emphasis in original). Whether an act is judged to be an appropriate form of engagement will therefore be determined by the extent to which the gamer's actions accord with the in-game context. Under such a condition, the gamer is required to engage with the game in order to meet the aims of the game designer(s) and not to satisfy their own ends should they be different. Of course, this alignment does not necessarily make the enactment morally permissible; it simply means that the gamer is acting in accordance with the narrative. One may still wish to challenge the morality of the enactment itself regardless of the extent to which the two contexts are aligned. The advantage of the 'appropriate engagement' requirement is that it allows any moral ruling to be applied to both in-game and gamer contexts, as these are now congruent. In short, it removes the ambiguity surrounding who or what is the object of our moral concern where these two are potentially different.

5.3.3   *Morally Objectionable Intrinsic Properties and Questionable Viewpoints*

Focusing, for now, on the in-game context, Ali, following Tavinor (2009), proposes two ways in which a virtual representation can be morally objectionable. The first is the representation itself. I find this claim troubling

as it requires that there is some intrinsic property of the virtual object that is morally objectionable (perhaps in virtue of what the object is a representation *of*). In response to my worry, one may wish to claim that *prima facie* representations of child rape have just such an objectionable intrinsic property, even if they do not involve images of actual children (leaving aside issues relating to how abstract or realistic the depiction is). This being the case, it seems reasonable to add that one could make the same claim about child murder or adult rape (among other things). Yet child murder is famously depicted on two occasions in works of art by Rubens (for example) when illustrating the biblical account of the *Massacre of the Innocents.* Likewise, a number of artists have captured in paint or marble the historical rape of the Sabine women. Moreover, the Ancient Greek custom of *paiderastia* (meaning *boy love*) which was represented in artwork and/or decorative pieces at the time is still readily available to view today, even though this custom is no longer considered morally acceptable or indeed lawful.

Is each of these representations *necessarily* objectionable on account of some intrinsic property? I would say no; and the fact that each example is or has been available to view lends support to this claim. Nevertheless, I concede that some of the examples may strike the reader as more controversial than others. To understand why, it is again reasonable to surmise that what might appease or fuel our moral objection is the viewpoint we are invited to adopt when viewing these depictions (whether this is made explicit or merely implied, or even left ambiguous, intentionally or otherwise). This seems particularly pertinent to the example of the naked 10-year-old Brook Shields (introduced in Section 1.3.1; and perhaps also to the other examples discussed in that section) which some have interpreted as eliciting only prurient appeal and condemned as nothing more than a magnet for paedophiles (Singh 2009; Young and Whitty 2012). (For further discussion on the difference between art, erotic art and pornography, see Levinson 2005; Mag Uidhir 2009.) Ali's second in-game objection is therefore based on the viewpoint the representation serves to express (or may be reasonably construed as expressing, I would add).

5.3.4   Differentiating Between Storytelling and Simulation Games

Of course, it is likely that a moral objection is raised in response to each of these suggestions combined: that is, to the viewpoint one is invited to adopt (whether implied or otherwise) in conjunction with the nature of

the representation itself. Together, these aspects or attributes of the representation constitute a suitable target for moral scrutiny, irrespective of whether our aim is to exonerate or condemn. So how does such an approach challenge the premise on which the gamer's dilemma is built: that, intuitively, virtual murder is permissible and virtual paedophilia is not? To see how, consider the following example used by Ali with reference to the *Silent Hill* survival horror series:

> In *Silent Hill 2*...the gamer controls a character who has murdered his own wife. The gamer controls this character as he uncovers the repressed truth about what he has done. Consider now the possibility of a *Silent Hill* game that takes on an equivalent scenario involving pedophilia...[I]t is not clear that a virtual pedophilic act in that game would be impermissible. (2015, p. 272; emphasis in original)

Among other things, the aim of the *Silent Hill* video games, Ali informs us, is to evoke within the gamer a sense of the psychologically disturbed state of the protagonist who, in *Silent Hill 2*, has murdered his wife or, in the case of the hypothetical game, has committed child abuse. The *Silent Hill* series is an example of what Ali calls storytelling games, which he contrasts with simulation games (discussed below) in which there is no *in-game* story to tell but, rather, one creates one's own narrative (e.g. *The Sims* series).[4] Where the video game involves storytelling, Ali draws the following conclusion:

> [W]hen it comes to storytelling games, acts of virtual murder and virtual pedophilia can be equally acceptable/unacceptable. This is because their unacceptability hinges on the very same features of the game, namely, the moral viewpoint of the story, and the use of objectionable or non-objectionable representations. (2015, p. 273)

Let us apply Ali's approach to Luck's original example of a (fictitious) game featuring virtual paedophilia: the jewel thief sleeping with the Beefeater's 15-year-old daughter (introduced in Section 3.2). If one plays the game in accordance with Ali's 'appropriate engagement' requirement then the aim is to align one's gaming objective(s) with that of the game designer(s) intentions (so that the gamer and in-game contexts match). One would therefore engage with the gameplay in a manner permitted in order to steal the jewels, which means sleeping with a character which, according to the in-

game context, has the legal status of a child. The use of this strategy is, of course, a contingent fact relating to the virtual theft. A different version of the game could have been designed in which one is required to kill the daughter instead of seducing her, or kill the Beefeater rather than the daughter, or where no death is required in order to acquire the jewels or where all of these options are available. In the case of the examples involving killing or paedophilia (which of course are the ones we are interested in), Ali's point is this: that (a) the act of paedophilia or killing would be classified as unlawful and immoral within the game and (b) the act can and should be understood within the narrative provided by the game designer(s), perhaps illustrating further the immoral nature of the protagonist, thereby making a case for the act itself not being gratuitous, and therefore being permissible *in this context*. Given this, whether the thief sleeps with or kills the daughter (or kills the Beefeater, instead), each enactment should be viewed as something that is morally equivalent. This means that if one version of the scenario is allowed then so should the others, and vice versa in the case of impermissibility. What cannot be justified, Ali would argue, at least in this and similar storytelling scenarios, is *selective* prohibition. Where selective prohibition targets virtual paedophilia, Ali's point is that such selectivity in the context described is morally unjustified.

What we have so far is an argument against *selective* prohibition. But this does not provide us with a clear indication or line *not to cross* in the case of the perspective the narrative presents. In other words, could there ever be a context or intended narrative within a gameplay or other work of fiction in which the depiction and therefore the perspective we are invited to adopt should be impermissible? This is a contentious issue. After all, the fiction is necessarily make-believe. As such, perhaps what we are invited to *imagine* is the world through the protagonist's eyes (as the *Silent Hill* example illustrates) without wishing to endorse the view that 'to explain, or perhaps merely to experience, is to condone'. Even in cases where one is invited to make-believe a moral inverse (i.e. play a video game in which it is virtuous to murder and take what you want, including sexual conquests), what one is enacting is, by its very design, not true (in fact, more than this, it cannot be true because it lacks truth-aptness). On the other hand, where one is invited to play a video game in which one is not invited to *make-believe* some deviant moral position (e.g. pretend that murder is acceptable) but, rather, enact something that corresponds to an actual moral attitude/belief (simulate *x* within the game and in doing so endorse the actual *belief* that *x* is acceptable) then one has the grounds for an argument

in favour of moral prohibition (see Young 2015a, for further discussion on this point). Of course, it is unlikely that such a game (which explicitly promotes/invites us to adopt some form of inverse morality as an actual world view) would become commercially available. Instead, what is more likely, and therefore what has been used as an argument for moral prohibition, is the example of an individual who engages with the in-game narrative (which is perhaps inviting the gamer to *make-believe* a world with some form of inverse morality) as a medium through which they can enact/express their own beliefs and desires (it is possible, of course that one may believe that murder is wrong but still desire it). The problem, of course, as we have already discussed, is in knowing which is of these is the case. If I enjoy playing the game is this because I enjoy the make-believe (*qua* make-belief) or because it provides an outlet for my own moral attitude and/or desire (which the make-belief corresponds to) which I repeatedly endorse through this virtual fantasy (Gaut 1998). Here, we see a shift in the object of moral concern: from the content itself (or in-game context) to the motivation of the gamer (gamer's context) It is a switch we have encountered already when examining earlier arguments by Morgan Luck, and it is one that Ali makes in relation to simulation games.

With simulation games, there is no in-game context because there is no pre-designed narrative. For Ali, what the gamer wishes to do (the gamer's context) – in conjunction with what the game-mechanics permit (ignoring cheats) – defines the in-game context. At the very least, then, what unfolds within the game reflects what the gamer finds desirable. A game in which there is no in-built narrative, where one simply goes around murdering innocent people or committing other acts of violence, Ali refers to as a *violence simulator*. Regarding such games, he proffers the following opinion: morally, they should not be permitted; not if the only reason for carrying out an act of virtual violence is the fact that one is free to do so. In essence, having the freedom to enact violence *for its own sake* does not mean one should.

5.3.5   *Accounting for Our Intuitions*

Ali does not provide a clear argument for why one should not engage in virtual violence for its own sake, although towards the end of his paper he does offer some insight into the form his argument might take (a point I will return to). What he does present is a plausible explanation for the intuition that virtual murder is permissible and virtual paedophilia is not

(at least with regard to those who endorse this intuition). He explains this by contrasting storytelling games with simulation games and argues that in the case of virtual murder our default position is to think of this act within the context of the storytelling game. As such, we play and therefore virtually murder in order to engage with the narrative. That is our motivation. In contrast, our default position regarding virtual paedophilia is to envisage such acts within a (hypothetical) simulation game. This means that the gamer creates the narrative which (according to Ali) is suggestive of an immoral desire because it is a means of expressing that desire. To reiterate, I find this a reasonable description of what causes the difference in what we are said to intuit. An implication of his repositioning of our moral gaze away from the virtual content *per se* towards the type of context in which the content appears (i.e. storytelling or simulation) means that what we *should* distinguish between (*qua* a normative position) is not virtual murder and virtual paedophilia (in and of themselves) but, rather, the context in which the enactment of either occurs; that is, whether it forms part of the narrative (in-game context) or is purely a reflection of the player's beliefs and/or desires and therefore motivation (gamer's context).

5.4   What Is Wrong with Enacting Taboos for Their Own Sake?

Ali's move from a description of how we think about these acts (based on the context in which they occur) to an account of what ought to be permissible is, however, less convincing. It is less convincing primarily because Ali does not present a clear argument for why engaging in virtual murder or paedophilia, for its own sake, is morally wrong. Given this, one could retort that those who play simulation games of the kind described are simply engaging in an activity that (to borrow from Patridge) is best understood as an intentional departure from the moral reality of our lives, and that is all there is to it. Consequently, such virtual events should be of limited moral concern. One might even argue that a simulation game in the mould of, say, *Sexual Assault*, which permits random sexual attacks on any character within the game, is in fact more divorced from our lived morality than a game like *S.H.: Random Attack* (involving random murder) given the not infrequent occurrence of random shootings, particularly in the USA, but increasingly in Europe.

While Ali might be willing to concede that a game like *Sexual Assault* is far more removed from our reality than *S.H.: Random Attack*, equally, he may wish to fall back on the intuitive pull that the gamer who chooses to enact sexual assaults, including assaults on children, even within the confines of what is meant to be a playful space, must be motivated to do so because there is something about the enactment that they find enjoyable. Otherwise, why else would someone play such a game *for its own sake*? But equally, Ali seems to be suggesting that someone who plays any form of simulation game in which they inflict harm on another for its own sake is doing more than just engaging their imagination. What they are doing is "materializing the fantasy, enacting it virtually, in a way that is perceptible to the gamer" (2015, p. 274). Such a view is similar to Žižek's (1997) claim, albeit in a slightly different virtual context, that "fantasies are increasingly immediately externalized in the public symbolic space; the sphere of intimacy is more and more directly socialized" (p. 164). The normative aspect of Ali's argument is then made apparent when he states, "It is in having this desire and seeking to actualize it that virtual murder is unacceptable" (2015, p. 274), although he would say the same of virtual paedophilia.

Importantly, the expression of one's fantasy in a more public (game) space, in the absence of context (in the form of an acceptable, legitimate and therefore mitigating narrative), makes Ali's argument as much applicable to virtual paedophilia as it is to virtual murder, thereby challenging the assumption that we (*qua* gamers and whoever takes an interest) intuitively differentiate, morally, between these two types of enactment. What our putative moral intuition relies on, Ali suggests, is context and not simply content. (I say '*putative* moral intuition' in anticipation of my move in the next chapter towards talk of moral attitude and reasoning rather than intuition.)

As interesting as Ali's position is, what remains unclear, given the decontextualized and narrative-free nature of the virtual murder and paedophilia we are discussing, is what the gamer's desire is exactly. In other words, given the freedom afforded by simulation games, what is the gamer seeking to actualize? Is it the desire to murder for real that is actualized within the gamespace (as a means of achieving vicarious satisfaction), or the desire to *enact* murder because it represents that which is unlawful, immoral: a taboo (thereby enabling the satisfaction of this desire to be achieved directly)? In the latter case, the desire to enact a taboo, because what one is enacting is a *taboo*, changes the object of one's desire and moral

gaze from desiring something that is illegal (actual murder), and delighting in the idea of this, to desiring to do something that is not (i.e. *enact* murder). It is a trivial truth that one cannot actualize murder through simulation. One can, however, actualize fantasy through simulation, which no doubt is Ali's point. And so we return to the crucial question: Is the virtual enactment of murder, the actualization of one's fantasy regarding the desire to engage in *actual* murder, or the actualization of one's desire to *play* at simulating something that is taboo? *Mutatis mutandis*, the same question applies to virtual paedophilia?

Recall from Section 5.3.1, Ali's example of a gamer enacting fantasies of murder through the in-game context of killing as a means of self-defence (and therefore not engaging appropriately with the in-game context). What Ali seems to be describing here is a gamer fantasizing about actual murder. But even if this is the case, it does not negate the possibility that some other gamer could enact a murder within gamespace, whether in the form of appropriate engagement or not, as a means of satisfying their desire to enact a taboo because it is a taboo. Recall also how, in Chapter 3, I discussed and dismissed the argument that the gamer is necessarily motivated to engage in virtual paedophilia because they find the idea of actual paedophilia appealing. This seems to be the move Ali is making here. Yet, one could argue that trying vicariously to satisfy one's desire for actual murder or actual paedophilia is not engaging in the virtual act *for its own sake*; rather, the act has instrumental value only. As such, where virtual murder or child abuse is carried out for its own sake then this has to be in order to satisfy directly one's desire to *enact* a taboo, and not to satisfy vicariously one's desire to engage in the activity for real. Therefore, for the reasons expressed above and in Chapter 3, Ali's move is unconvincing because he cannot say for certain what is motivating the gamer to play a 'violence simulator' or the equivalent in the case of virtual paedophilia (e.g. *Child Sexual Assault*).

In sum, Ali's approach has promise, and proffers an original argument for undermining the assumption grounding the gamer's dilemma. That said, I find it difficult to accept that a hypothetical simulation games like *S.H.: Random Attack* would be as morally condemned as a game like *Sexual Assault*, let alone *Child Sexual Assault*. In his defence, Ali could claim not to be accounting for what would be the moral position of gamers (in a descriptive, psychological sense) but what it *should* be (in a normative ethic sense). Even so, given Ali's failure to make explicit the type of fantasy the gamer is seeking to materialize within gamespace, as well as the underlying

motivation for its materialization, especially given the possible motivations available with their potentially different moral implications, his account ultimately fails to convince as a *normative* approach to video game content and therefore as a means of resolving the gamer's dilemma.

NOTES

1. McMahan used these terms originally in relation to his critique of just war theory.
2. To be clear, although Patridge is presenting an argument for 'non-harm-based' moral reasons not to engage in virtual paedophilia, this does mean that such enactments (if one were to engage in them) will not contribute to harm, such as the cultural harm mentioned (see Section 1.3.2).
3. $M_{(enjoyment)}$ does not negate that one is motivated to play the game because one finds virtual paedophilia sexually arousing in and of itself. Here, however, Patridge has explicitly ruled out the possibility that the group treat the depictions as pornography.
4. Ali also includes sporting games (as a separate category), but the addition of this extra category is unnecessary here, as it adds nothing to the discussion, nor does its removal diminish our understanding of the point Ali is making.

A New Approach to Resolving the Gamer's Dilemma: Applying Constructive Ecumenical Expressivism

Abstract In this final chapter, a new approach to understanding the gamer's dilemma is presented which seeks not to identify a single morally relevant factors which differentiates virtual murder from virtual paedophilia but, rather, aims to articulate the means by which (a) we acquire the moral attitude we do and (b) how this attitude is elevated to the status of a social norm. Constructive ecumenical expressivism is posited as the means of accounting for this and therefore explaining the intuition that is said to form the basis for the gamer's dilemma. The new approach's ability to resist objections raised against an appeal to social convention is also discussed, as is the form a normative ethic would take if one were to endorse constructive ecumenical expressivism.

Keywords Moral attitude · Moral realism · Anti-realism · Meta-ethical approach · *De re* and *de dicto* attitude

So far, I have presented a critical review of the various attempts proffered to resolve the gamer's dilemma. To be fair, a number of these attempts were rejected by the author at the time of their original publication (e.g. Luck) or have been challenged by other author's since (e.g. Luck & Ellerby's and Patridge's responses to Bartel's proposed resolution). Some recent attempts offer promise, however (e.g. Patridge and Ali),

although for reasons discussed fail to resolve the dilemma and lack the resources to be co-opted as a normative ethic.

In this final chapter, I present my own thoughts on how the gamer's dilemma might be resolved. My strategy is to consider the nature of the moral discrepancy on which the dilemma is built: the intuition that virtual murder is permissible and virtual paedophilia is not. What does it mean to morally approve or disapprove of something? My answer is that it means that one has either a positive or negative attitude towards the target of one's moral inquiry. Given this, how does one arrive at this attitude? Rather than simply falling back on the idea of an intuition, I intend to examine the relationship between the object of moral concern, our moral attitude towards it and the process by which we arrive at this attitude. Once this has been established, I intend to apply this approach to resolving the gamer's dilemma, with a view to broadening its application to virtual gaming content more generally.

6.1   Now *That* Is Immoral, Isn't It?

There is no objective understanding of what an image might represent – it is in the mind of the viewer.

(Simpson 2009, p. 260).

Suppose I agree to play *Child Sexual Assault* with a friend. Shortly after commencing the game, my friend turns to me and pointing at something within gameplay says: "*That* is immoral". For my part, I disagree and tell him that *that* is not immoral. What has brought about this moral disagreement? Before responding, a passage from Patridge should prove informative:

[I]t seems that the gamer who cannot help but see *Child Sexual Assault* as a reflection of or extension of our moral reality should be unable to find this content enjoyable (or, again, should find it very difficult to do so). This is so because the object of her amusement is a different object altogether from the object as interpreted by the gamer who sees it only as a bit of harmless fun. The objects are interpreted quite differently and as a result the instances of amusement involved have different intentional objects. (Patridge 2013b, p. 32)

Here, Patridge is contrasting the two gamers (above) based on how amusing they find aspects of the gameplay. Given the discussion on

Patridge in Chapter 5, the difference in their amusement is meant to reflect a difference in their moral attitude towards certain enactments: moral disapproval in the case of not being amused and, at the very least, failure to disapprove in the case of amusement.[1] Returning to my example, if we take the demonstrative pronoun ('that') to be referring to an instance of virtual paedophilia then, following Patridge, the moral disagreement my friend and I express is not symptomatic of a difference in moral attitude towards *actual* child rape (or not necessarily so); instead, it is that in the case of *this* token enactment, my friend and I have interpreted the intentional object differently. Again, following Patridge, the enactment is seen as either an extension of (in my friend's case), or a departure from (in my case), our moral reality.

This difference in interpretation is important and deserving of further attention. Patridge's comments are in fact in keeping with a meta-ethical position I have previously discussed called *constructive ecumenical expressivism* (CEE) (Young 2014, 2015b). CEE offers a new way of thinking about the gamer's dilemma. It does this not by identifying what the morally relevant difference is between virtual murder and virtual paedophilia, in some moral-realist sense, but by explaining why a difference in moral *attitude* occurs. But more than this, CEE provides the means of understanding what is required for a normative ethic to be established; not only in regard to those virtual enactments involved in the gamer's dilemma, but all video game content.

6.2   CONSTRUCTIVE ECUMENICAL EXPRESSIVISM

CEE, like its forerunner – ecumenical expressivism (Ridge 2006) – holds that, when uttered, the proposition "*That* is immoral" reveals two inter-related facts about the mental states of the subject. The first concerns an attitude. The second relates to a particular belief that is said to make anaphoric reference to this attitude.[2] Importantly, though, the subject does not hold (in this case) a negative attitude towards the particular act referred to by the demonstrative pronoun. Instead, in the case of "*That* is immoral", the subject disapproves of some property – call it P – and believes that x (which represents what 'that' refers to) realizes P. Thus, in declaring that murder is immoral, the subject holds a negative attitude towards P (some property yet to be described) and believes that an act of murder realizes P. Moreover, in stating that murder is immoral, the subject is not (should not be) declaring only that *this* token act of murder

is immoral but, rather, that the *type* of action, of which *this* particular act is a token, – in virtue of realizing property P – is immoral. Where a different type of act realizes the same property, then one should disapprove of any token act of this action type. If kidnapping, for example, realizes property P then, if one disapproves of murder (in virtue of property P), one should disapprove of kidnapping, also.

So what is property P? Property P can and does amount to different things for different people. S_1 may view P in terms of negative utility – for example, the realizing of more displeasure than pleasure (say, in the form of increased harm) – while S_2 may hold it to be a violation of God's law, or constitutive of a failure in one's duty to others. S_3, in turn, may characterize P as a vice rather than a virtue, and so on. Declaring that "x is immoral" – where x equates to murder – reveals the following:

(CEE) S disapproves of P and believes that x realizes P (thus making anaphoric reference to that of which S disapproves).

To state that x is immoral does not denote the truth of the proposition "x is immoral"; rather, it denotes the truth of (a) S's disapproval of P, and (b) S's belief that x realizes P.

In the context of *Child Sexual Assault*, where A (*qua* my friend) declares "*That* is immoral" and B (*qua* myself) denies this, both A and B are expressing a moral attitude. A disapproves of *that* (whatever 'that' happens to be) and B does not. More specifically, A disapproves of *p* and believes that *that* (whatever 'that' happens to be) realizes *p* (thus making anaphoric reference to that of which A disapproves). Before moving on, note that I have used the lowercase *p* in italics to denote the specific property of which A in particular disapproves. This should be contrasted with the uppercase P used earlier, which refers to some unspecified property of which a generalized subject disapproves.

Why does B not disapprove of *that*? It could be that both A and B disapprove of property P (where property P refers to the same thing). Nevertheless, it could also be that they are employing different interpretations of x: the event within the game. How each interprets the virtual event will shape whether they come to believe that x realizes P, and based on this belief whether they disapprove of x. Alternatively, it could be that both A and B consider x to be immoral. In other words, both agree that "*That* is immoral" in relation to the same virtual event. How might this be achieved?

In this case, moral agreement (or shared moral attitude) is achieved because both believe that x realizes some property (P) of which they each disapprove. It could be that P is the same for both players. However, it could also be that what counts as P is different in each case. Suppose A believes that x realizes p: where p equates to commending one to delight in that which is immoral (actual paedophilia, in this case). B, on the other hand, does not believe x realizes the property just described but still considers x to be immoral because B interprets x as realizing q, where q equates to an increase in harm, either to oneself and/or others (based on increased negative affect/ attitude/behaviour), which is something B disapproves of.

According to this explanation, both A and B consider x to be immoral but for different reasons. It is not that A does not disapprove of something which causes increased harm, or that B does not disapprove of something which commends us to delight in the immoral; rather, it is that A does not believe that increased harm is a property realized by x, or does not prioritize it above a different property (commending one to delight in the immoral) which is held to be the main reason for A's moral disapproval. *Mutatis mutandis*, the same goes for B. After all, in the case of B, it is possible to believe that x does not commend one to delight in the immoral but still believe it can lead to increased harm. Conversely, one can believe that x commends us to delight in the immoral, and that this is reason enough for disapproval, irrespective of whether it causes any increased harm. In short, in this scenario, both A and B express a negative attitude towards x but for different reasons: that is, in virtue of believing that x realizes some property (P) which equates to something different in each case (p or q), but nevertheless serves the same reason-giving function. To illustrate, in the case of "x (*qua* virtual paedophilia) is morally wrong":

(CEE a)   A disapproves of p (where p equates to commending one to delight in the immoral) and believes that x realizes p (thus making anaphoric reference to that of which A disapproves).

(CEE b)   B disapproves of q (where q equates to increased harm) and believes that x realizes q (thus making anaphoric reference to that of which B disapproves).

In each case, the moral attitude towards x is the same: namely, "it is wrong". This is because some property (P), of which A (*qua* property p) and B (*qua* property q) disapprove, is believed by A and B, respectively, to be realized by

x. There can be any number of reasons for one's attitude towards something. What CEE teaches us is that moral attitude is no different.

6.2.1   *Constructing a Moral Norm*

Where a shared moral attitude occurs with regard to some object or event, as a society we are able to create or *construct* a social norm that then acquires its own objectified moral standard. As Prinz (2007) states, "Things that we construct or build come from us, but, once there, they are real entities that we perceive" (p. 168). With the force of social consensus, and the moral norm this creates, we can adopt a normative position whereby a particular (agreed) attitude is the one we *ought* to have, at least with regard to *this* object of moral inquiry. Copp (2011) likens this to what he calls *realist expressivism* (see also Copp 2001). Where S shares this attitude, we can commend her for doing so. Where S does not, it is appropriate (given the constructed moral norm's objectified status) to rebuke her for her alternative (some might even say deviant) moral attitude. This is because both the rebuke and a change of attitude on the part of S are deemed to be warranted (Nichols 2008).

In the case of "murder is immoral", I may share this attitude with a neighbour. Our shared attitude in turn aligns with the objectified moral norm of our society. Unlike my neighbour, though, I am not interested in what constitutes a violation of God's law and so cannot be said to have a negative moral attitude towards murder *because* it violates God's law. Despite these differences, we (my neighbour, myself and wider society) still express a shared negative attitude towards the act. These similarities and differences CEE is able to accommodate and explain. In this instance, how (for example) my neighbour and I interpret the act (the object of our moral inquiry) leads us to draw the same conclusion about whether we approve or disapprove of it. But, importantly, not because of the fact that we agree on what the act is (a token example of a type of intentional, illegal killing: namely, murder), nor because we agree on what properties the act realizes (because in this instance we do not, at least not completely) but, importantly, because we each disapprove of at least one property we believe the act realizes, even though this property is different for each of us. In other words, even though we both interpret the act in the same way (as murder), we nevertheless differ in terms of our beliefs about a certain property or properties it realizes, or how we prioritize these properties (i.e. a violation of God's law or something else: say, violating Kant's categorical

imperative or increased negative utility, a vice rather than a virtue, and so on). But the fact that we each believe that murder realizes *some* property of which we disapprove, even where we fail to agree on what this property is, means that our attitude towards murder is the same. Where enough people share the same attitude, but not necessarily for the same reason (*qua* belief about a property or properties realized), a social norm is constructed resulting in an objectified moral standard.

With run-of-the-mill first-person shooter games, however, in which one can enact random murder, while moral consensus is forthcoming in the case of actual murder – because such an event realizes a selection of properties at least one of which, but quite possibly more, the majority of people are willing to condemn (e.g. negative utility, a violation of God's law or Kant categorical imperative, a vice rather than a virtue) – this is less the case with regard to virtual murder. In essence, with immoral acts involving actual persons (i.e. any kind of sexual assault and murder), if one endorses CEE, such is the array of properties to disapprove of that it is simply a case of taking one's pick. In contrast, it is far less clear which properties are realized by the virtual enactment of an immoral act because this depends much more on how one *interprets* the event which, in turn, affects one's belief about the properties it realizes. Recall, for example, my discussion in Section 4.3 on whether virtual paedophilia is child pornography. Therefore, whether I interpret a token act of virtual paedophilia as *child* pornography or just pornography, or neither, will likely affect the properties I believe the virtual event realizes and, depending on whether I approve or disapprove of these properties, my moral attitude towards this type of enactment. Conversely, it may be that how I interpret and therefore categorize a virtual event depends on the properties I believe it realizes. Again, using the example of virtual paedophilia, I may not categorize this type of enactment as child pornography because I do not believe it realizes the property of child abuse, and therefore do not have a negative attitude towards it for that reason at least.

When observing or interacting with a virtual event, if one cannot agree on what *that* is – and by this I mean what the virtual enactment is meant to represent and/or one's belief about the properties it is said to realize – then it becomes easier to see why we might find it harder to agree on whether *that* should be judged morally good or bad. Having said that, where different people do hold different beliefs about which properties a virtual event realizes then as long as these different properties are disapproved of (by each respective person holding the belief about their

realization), there is a good chance that moral consensus will be achieved and a shared moral attitude established. What I hope to show in the discussion to follow, which amounts to the application of CEE to the gamer's dilemma, is why there is a greater consensus of negative moral attitude towards virtual paedophilia than virtual murder. In addressing the 'why' question I hope to establish a normative ethic that is able to justify the selective prohibition of virtual content not only in relation to the gamer's dilemma but also beyond.

6.3   Social Convention by Another Name?

One anticipated objection to CEE is that, in reality, it amounts to the same, albeit slightly more sophisticated, argument based on social convention rejected earlier (see Chapter 2). In response to this anticipated objection, I would say that it is precisely this added sophistication that enables CEE to overcome the problems raised against the social convention argument. Each, it must be said, adopts an anti-realist approach to moral utterances. While it is not my intention to defend my anti-realist stance in detail here (for further discussion, see Young 2014), a cursory exposition of my reasoning is required in order to show how CEE can overcome the challenges levelled at an appeal to social convention and therefore moral subjectivism more broadly construed.

When considering the proposition "*x* is morally wrong", one should not understand the sentence to be a *description* of some state of the world which captures a moral reality: namely, that *x*, so described, is literally picking out a moral wrongdoing. The problem with descriptivism in the context of moral realism is that moral utterances seem to be both descriptive and *evaluative*. To illustrate, the proposition "S is a paedophile" purports to describe some fact about S (that he is sexually attracted to children). This statement is truth-apt and is therefore either true of false. In addition, there is also an implied evaluative component: that being a paedophile is something one ought not to be because it is morally wrong. Moral utterances therefore contain both an 'is' (descriptive) component and an 'ought' (prescriptive) component, with the former being on a more secure metaphysical footing than the latter in terms of purporting facts about the world. Consequently, while it is true that S being a paedophile is either true or false based on how one defines paedophile and the sexual preferences of S, what is less clear is what makes it *true* (as in some independent fact about the world) that being a paedophile is morally

wrong? Recall, for example, the Ancient Greek custom of *paiderastia* (meaning *boy love*), mentioned in Section 5.3.3. Certainly, in its day, this practice was not judged to be immoral. Contra moral realism, then, the evaluative component of the proposition "*x* is morally wrong" should not be thought of as a description of some mind-independent moral property; instead, it should be thought of as a direct *expression* of one's attitude to *x* which functions in an evaluative way. In Ancient Greece, the attitude to adult men having sex with adolescent boys was certainly different to our *conventional* moral attitude; but in neither case – that is, either then or now – could the utterance "paiderastia is morally wrong" be said to be true or false: for it lacks truth-aptness. What it amounts to is either something one agrees with (*qua* approves of) or does not.

Embracing anti-realism with regard to moral utterances does not, in and of itself, help us differentiate between social convention and CEE. As we saw in Section 6.2, the constructive component of this form of expressivism describes the social elevation of a shared attitude to the point where it achieves, through consensus, an objectified moral status; at which point it acts as an independent (of any individual) measure of morality and, in doing so, delineates what is morally warranted within a given society (again, independent of any individual moral attitude), and therefore what constitutes *that society's* moral reality. How this differs from the social convention argument presented in Chapter 2 is outlined below:

> To state that A and B have a shared attitude towards *x*, such that they both hold that *x* is immoral, is to declare that they have the same *de re* attitude. When considering the act that A and B's attitude is directed towards (the intentional object), their attitude towards that act (the thing in itself) is the same. But this shared *de re* attitude exits in virtue of the belief that *x* realizes some property (P) which they both disapprove of, but which can be (and is) different for A and B: A believes that *x* realizes *p* and B believes it realizes *q*. Their differing belief about which property is realized by *x* means that they have different reasons for their shared *de re* attitude. One could say that they have different *de dicto* attitudes regarding *x*… (namely, different beliefs about why it is immoral). (Young 2015b, pp. 317–318)

What A and B have in common is their negative attitude towards *x*. However, this singular attitude (it is singular because it is held by both A and B: hence, *de re*) is adopted by A and B for different reasons (they have different *de dicto* attitudes). It is therefore my contention that, in the

absence of the truth of moral-realism (which is the position I am endorsing here), a *de re* attitude shared by the majority of people within a given society (such that it becomes the constructed moral norm or, if you like, convention of that society) is more robust if it is the product of a number of different *de dicto* attitudes (i.e. if it is based on a number of different reasons for having the moral attitude). This should not be taken as evidence of inconsistency, and therefore as a reason to undermine the normative authority of the moral attitude; rather, and to reiterate, it should be taken as evidence of its robustness, insofar as there are purportedly many reasons for why *this* (whatever 'this' happens to be) is morally wrong. It just so happens that different people have different views on what these reasons are or how they prioritize them. Therefore, to undermine the moral (*de re*) attitude, one would have to undermine the various reasons (*de dicto* attitudes) justifying its (objectified) normative status. Such a position does not rule out a change of *de re* attitude – there is therefore a degree of fluidity inherent within the position – but it does make any such change less capricious.

An appeal to social convention is therefore vulnerable to capriciousness in a way that CEE is not: because the appeal to social convention presented in Chapter 2 was said to be intertwined with the more elusive idea of moral intuition which is not the case with CEE. But what if one particular society has a shared negative attitude towards people of a particular race or sexual orientation or who hold certain religious beliefs? Given the lack of truth-aptness in regard to moral utterances, is the proposition "*These* people are morally inferior" just as valid as the converse utterance? In response, I would say that while the proposition itself lacks truth-aptness, what is capable of being true (or false) is (i) whether S *believes* that a particular racial type (for example) realizes a certain property, (ii) whether S disapproves of this property, and (iii) whether this particular racial type actually realize this property, thereby confirming or disconfirming the belief held by S. After all, it is possible for S and S's society to ground their moral attitude on a false belief. Given this, one could challenge the reason (*de dicto* attitude) for the negative (*de re*) attitude should one believe that this reason stems from a false belief.

But suppose S believes that Afro-Caribbeans realize the property of darker skin (darker than S's pale skin, at least) and disapproves of this property. S's negative moral attitude towards Afro-Caribbeans is based on a belief that is in fact true and so S's attitude cannot be challenged, in this instance, for being based on a false belief. Of course, one would simply ask

why S disapproves of darker skin – for example, what belief is the disapproval of darker skin based on? (etc.) – until, ideally, one arrives at a demonstrably false belief. I say 'ideally' because the (alleged) truth or falsity of certain beliefs may depend on how one interprets the evidence available. Some beliefs may therefore lack any definitive verification (i.e. one may fall back on a belief in the word of God, or what is written in some (held to be) sacred text). Such a situation makes it likely (inevitable, some might say) that moral disagreements will arise. This is indeed the state of our moral reality today. I see no reason to advocate moral realism because of this, however.

CEE therefore provides a more fine-grained explanation for why social conventions differ across different cultures/societies (e.g. Japanese attitudes to what in the West we would judge to be virtual child pornography; see Section 2.1.1). In the case of differences across societies or even disagreements within one's own society, given that a society's shared *de re* attitude is likely to be based on different reasons (*de dicto* attitudes), where these reasons are few(er) or perhaps less entrenched, there is more chance that one could challenge the social norm (*de re* attitude) by undermining the different beliefs (*de dicto* attitudes) held by different members of that society (assuming one believes the *de re* attitude to be grounded, at some point, on a false belief). If, for example, the only reason for a society's negative moral attitude towards murder is that it violates God's law then, if one wished, one could challenge the merits of this reason by trying to undermine the belief in the existence of God and, with it, the putative authority of God's law. Fortunately (as already noted), there are many reasons (*de dicto* attitudes) for why one should have a negative (*de re*) attitude towards murder such that undermining them all would prove difficult: thereby accounting for why a negative attitude towards murder is universally expressed and entrenched within different cultures/societies.

In the case of virtual enactments, establishing whether a certain property is realized by a particular enactment is harder to do because of the greater scope for different interpretations compared to actual events. Consequently, it is harder to challenge the *belief* on which the moral attitude is grounded (whether it is the belief that P is realized by *x* or the belief that P is not realized by *x*). Does virtual murder realize the property of eliciting delight in the idea of actual murder, for example? Likewise, does it elicit the property of being a vice or negative utility, and so on? It is also the case that actual events will realize different properties compared to virtual enactment. These former properties are easier to

establish in terms of the truth of their occurrence (i.e. the negative utility of actual murder compared to virtual murder), thereby making it likewise easier to validate a *belief* in their realization (or refute it, depending on the belief).

6.4 APPLYING CONSTRUCTIVE ECUMENICAL EXPRESSIVISM TO ALI'S ACCOUNT

Pace Ali (whose argument was discussed in Chapter 5), I do not consider it controversial to accept the claim that gamers (or even the wider community) typically hold virtual murder to be less objectionable than virtual paedophilia. As we have seen, there have been a number of suggestions/arguments proposed to account for why this is. Each has posited a single factor (a) in the role of the 'morally relevant means of differentiating between virtual murder and virtual paedophilia' and (b) given (a), as the basis for a normative ethic that guides the selective prohibition of wider video game content. While some have shown more promise than others, in my view, each has failed to deliver on points (a) and (b). CEE, in comparison, does not seek to identify a single morally relevant factor. In accounting for the difference in moral attitude between virtual murder and virtual paedophilia, CEE is able to accommodate a number of different reasons based on differing *de dicto* attitudes (*qua* beliefs) about properties realized by the respective virtual event. Why is it more likely that a gamer will find virtual paedophilia morally objectionable? Because the gamer holds the belief that at least one property of which they disapprove is realized by this type of enactment (the converse epistemic relationship typically occurring in the case of virtual murder). These may include, but are not restricted to, prurient appeal (i.e. delighting in the idea of actual paedophilia), the likelihood of harm (whether to children or women), the targeting of non-morally relevant characteristics and so on. Each one of these beliefs (and others) has been challenged and shown to be problematic with regard to (a) and (b). This is partly because of the different ways the virtual event can be interpreted, which also relates to player motivation (Young 2015b). Importantly, though, the same ambiguity that works against satisfying (a) and (b) helps maintain those beliefs the gamer uses to ground their moral attitude (recall, the belief can function in this role even if it is false or not verified). It also accounts for why, in accordance

with Ali (2015), certain gamers may hold that some token acts of virtual paedophilia are not morally objectionable (or any more so than certain token acts of virtual murder). This is because, according to CEE (which is not incompatible with Ali's view), the gamers do not believe that these tokens virtual events, in the context in which they occur, realize the properties of which they disapprove. This does not rule out the possibility that other token acts of virtual paedophilia will amount to tokens of a certain type that *do* (they believe) realize these properties owing, say, to the context in which they occur (again, as Ali attests). Similarly, where virtual murder may appear to be gratuitous and/or targeted (e.g. racist) and/or excessively violent, in the absence of a morally mitigating context, gamer's may likewise believe that such enactments realize the same or similar properties to those realized by certain enactments of paedophilia, meaning that these should be disapproved of, also.

Gamers (and the wider community) typically hold a more negative moral attitude to virtual paedophilia because, collectively, there are believed to be more *de dicto* reasons for disapproving of this type of enactment compared to virtual murder (irrespective of whether these *de dicto* reasons *qua* beliefs are true). Here, I think Ali's point about the default position of gamer's who are asked about virtual paedophilia is germane. Because there are no commercially available games that enact the types of hypothetical scenarios gamers are often asked to envisage, it is plausible that they struggle to contextualize virtual paedophilia within a mitigating narrative (i.e. there are no existing examples for them to draw on, unlike virtual murder). Such a mitigating narrative could negate the belief that a token act of virtual paedophilia realizes a certain property of which the gamer would otherwise disapprove. In the absence of such a belief, CEE teaches us, the gamer has no reason to adopt a negative moral attitude. But this also means that in the context of a (hypothetical) simulation game, perhaps along the lines of *Child Sexual Assault*, should a gamer have no reason to believe that a certain property they disapprove of is realized through enacting a token paedophilic event (even 'for its own sake', although it remains unclear what Ali means by this) then that gamer would, and indeed should, have no reason to disapprove of the enactment. Should this apply to the wider community of gamers and beyond then CEE would predict the establishment of an objectified moral norm that would permit virtual paedophilia, at least of the type alluded to here.

6.5   APPLYING CONSTRUCTIVE ECUMENICAL EXPRESSIVISM TO PATRIDGE'S ACCOUNT

In the context of *Child Sexual Assault*, recall how Patridge holds that virtual paedophilia is an extension of our lived morality because the enactment is a further example of targeting for harm (albeit virtual harm) based on non-morally relevant personal characteristics (in this case, being a child). One is actually targeting something (an avatar) based on characteristics the avatar is represented as having (child-like features). Even targeting based on *representations* of non-morally relevant personal characteristics is, for Patridge, immoral: because the act of targeting for the reasons described is performed elsewhere and on other occasions (i.e. actual child abuse) as part of our lived morality, and it is immoral to do so. *Child Sexual Assault* is therefore just an extension of that immoral activity under the guise of a game.

Interpreting Patridge through the theoretical lens of CEE, we could say that she disapproves of p (targeting for harm based on a non-morally relevant personal criterion) and believes that x (a token act of virtual paedophilia) realizes p. Therefore, she has a negative moral attitude towards x in virtue of her belief that x realizes p and her disapproval of p. Let us allow that my friend (the one I introduced earlier in this chapter) has the same negative moral attitude towards x for the same reason. I, on the other hand, do not believe that x realizes p. Moreover, I do not believe that x realizes any property I disapprove of. Consequently, I do not have a negative moral attitude towards x.

Why do I not believe that x realizes p? Perhaps it is because I do not believe that targeting a *virtual* character, based on *representations* of non-morally relevant characteristics, in order to engage in *simulated* immoral activity is equivalent, morally, to targeting an *actual* person (individually or as part of a minority group) based on *actual* non-morally relevant characteristics for an *actual* immoral activity (e.g. harming them). Importantly, the representation I have in mind is akin to Baudrillard's (1983) notion of a *simulacrum* (mentioned in Section 4.3; meaning a copy or a representation of a thing that has no original but goes beyond itself) and so is not meant to include an actual photograph or similar recording. Let us say that the reason I do not believe they are equivalent is because although the putative non-morally relevant characteristics of the representation are *characteristics*, they are characteristics of a *representation*, thereby making the characteristics of the representation at the same

time representations of characteristics of an actual person (but not a specific person, in accordance with Baudrillard's simulacrum). As such, whether targeting these characteristics is of concern in any *moral* sense is debatable. In short, what I am contesting is whether representations of non-morally relevant characteristics are (and indeed can be) non-morally relevant characteristics of a representation. I am contesting this insofar as I am contesting whether all representations of this kind have any moral relevance at all.

In making this point, I am reminded of Edward Castrovona's notion of a *closed world*. According to Castrovona (2004), the border between our real world and a closed virtual world is impermeable. Consequently, while the virtual world may represent aspects of the real world, the features of each world are incommensurate. For Castrovona, closed worlds are spaces in which nothing matters; where assets have no real value and therefore losses are unimportant. In such a world, what does it matter if one murders someone for no other reason than one can, or abuses a child for the same reason? After all, in this world, nothing matters. Dunn (2012), however, casts doubt on the plausibility of a world in which nothing matters; where everything is without meaning. Why, Dunn asks, would anyone want to inhabit such a world? This is a pertinent question. If one's enactment was completely devoid of meaning then what would be the attraction? Therefore, for Dunn, the border between the two worlds is necessarily *not* impermeable. The meaning of the enactment is derived from what the enactment represents. Consequently, real-world meaning necessarily transcends the two worlds. If I seek to satisfy vicariously my desire to carry out the act for real then the enactment is meaningful as a substitute for that action. On the other hand, if I seek to enact murder or paedophilia because what I am enacting is a taboo in the real world then, similarly, the enactment's meaningfulness as an act I desire to perform is parasitic on what the enactment is meant to represent: namely a *taboo*. Such *unidirectional* transcendence is necessary for the virtual act to be imbued with any meaning *as an enactment*. What is contested, however, is whether the meaning of the virtual act itself transcends worlds so that it has real-world meaning, particularly moral significance. If the meaning of a virtual act does not transcend gamespace, thereby conveying *real-world meaning* and subsequently moral import, or if only certain actions do or should be said to, then the virtual world could be thought of as *partially closed* (Dunn 2012).[3]

It is the partially closed nature of the virtual world that gives me reason not to believe (contra Patridge and my friend) that *x* realizes *p*. Which of us is correct? That is, does *x* realize *p* or not? As I have argued previously, in the context of virtual enactments, this is difficult to verify because it is dependent on how one interprets *x* (*x* being much more open to interpretation when it constitutes a virtual act within gamespace). Such dependence on interpretation does not prevent the formation of *beliefs* about the realization of *p*, however. Instead, and to reiterate, it makes it difficult to *verify* the belief.

What CEE allows (although it does not require this) is that my friend and I can have a shared negative moral attitude towards actual paedophilia (that it is morally repugnant) while agreeing that what we are enacting within *Child Sexual Assault* is virtual paedophilia. We are simulating the very thing we disapprove of in the real world. Yet it does not require that this real-world accord about actual paedophilia and our agreement on what is being simulated carries with it a further *moral* agreement about the rights and wrongs of what we are doing when enacting virtual paedophilia. As I have outlined, CEE accounts for different moral attitudes, as well as shared attitudes for different reasons. The only consistency required is between one's attitude and one's belief(s). Where I disapprove of *p* and believe that *x* realizes *p* then I should (in both a rational and moral sense) disapprove of *x* (*qua* possess a negative moral attitude) in virtue of my belief that *x* realizes *p*.

If I believe that the virtual targeted murder of homosexuals within a video game (say, when playing *R.A.C.I.S.T.*) is morally wrong then, according to CEE, it is because I believe that this virtual act realizes some property (P) that I disapprove of. If I do not have a negative moral attitude towards *S.H.: Random Attack* then it is because I do not believe that a property of which I disapprove is realized by any of the token virtual engagements within this game. Likewise, I may hold that the (fictitious) video game *Sexual Assault* is morally wrong for similar reason to those given when discussing *R.A.C.I.S.T.*, or hold a negative attitude towards only certain token enactments within this game: say those involving the sexual assault of minors (again, for similar reasons to those just discussed). CEE also accounts for why some gamers may consider non-visual or implied acts of virtual paedophilia, or virtual grooming (etc.) to be morally wrong and, in doing so, is able to accommodate a broader approach to the gamer's dilemma, as suggested by Luck and Ellerby 2013; see Section 4.1).

6.6   ESTABLISHING A NORMATIVE ETHIC

CEE is fundamentally a meta-ethical approach concerned with understanding the nature of moral utterances. As part of this understanding, its focus is primarily on the process by which an individual comes to have the moral attitude they do, followed by how different individuals, possibly with different beliefs, can nevertheless have the same moral attitude. What CEE posits is that a shared *de re* attitude need not be the result of a shared reason (*de dicto* attitude); rather, it is simply the product of a consistency between belief (*de dicto* attitude) and *de re* moral attitude: the *de re* attitude being the appropriate (*qua* rational) expression of one's *de dicto* attitude. What CEE accepts is that one's moral expression is the product of a belief, not necessarily a fact, and so is not itself a moral truth. Given this, is CEE robust enough to establish a normative ethic? In order to address this question, let us consider what CEE as a normative ethic would look like.

According to CEE, what we ought to do is determined by our moral attitude; and when I say 'our' I mean, of course, the attitude indicative of the objectified moral norm within our society. Such a norm is constructed; it is the product of the moral attitude shared by the majority (but, to reiterate, not necessarily for the same reasons). In Section 6.3, I argued for the robustness of this objectified moral norm compared to the type of social convention discussed in Chapter 2. I also described how CEE permits moral change – and so is fluid – without being capricious. What CEE advocates, then, is a morality and, I would argue, a normative ethic based on the shared moral attitude of a given society (or the majority within that society) which evolves into an objectified moral norm. Where an individual's moral attitude deviates from this norm then their attitude deviates from what it *ought to be*. In the case of my own (fictitious) moral attitude towards enactments of virtual paedophilia within *Child Sexual Assault*, I argued that this was based on a particular belief I held regarding x (a token enactment of paedophilia) and a particular property I believed it did not realize, thereby giving me no reason to disapprove of it. If, however, the majority of gamers (and even wider society) believe differently, such that their belief(s) lead them to share the same negative moral attitude towards virtual paedophilia, thereby making it an objectified moral norm, then my attitude is off-kilter with theirs. A change in my attitude is therefore warranted.

What I could try to do is challenge this moral norm by challenging the beliefs on which it is founded; by trying to show either, *a priori*, that the

respective beliefs of those who have contributed to the moral norm are inconsistent or conceptually flawed, or that, *a posteriori*, they are not supported, or perhaps are even refuted, by empirical evidence (much as I have been trying to do throughout this book). Again, recall that a *de re* attitude can be grounded on any number of beliefs *qua de dicto* attitudes and therefore a number of different arguments may need to be forwarded to quash different beliefs. Should I prove to be successful at challenging the beliefs on which a particular moral attitude is based, resulting in the formation of new beliefs, alongside a corresponding change of attitude, then this change of attitude will become the new social norm and therefore the new proscribed way of thinking about *x*. As an aside, I have left unexplored the question of whether one has the *legal* right to have an alternative moral attitude and therefore whether we should defend this right.

6.7   CONCLUDING REMARKS

Where does this leave us with regard to the gamer's dilemma? With qualification, I accept that most gamers within the West and also our wider (Western) society find the idea of virtual paedophilia more morally objectionable than virtual murder. (I say 'with qualification' because some gamer's and non-gaming members of society may differ in their moral attitude depending on the *context* in which the virtual paedophilia and the virtual murder occur, as discussed in relation to Ali, and as allowed by CEE.) When attempting to resolve the dilemma, it is first necessary to understand that the shared moral attitude on which the dilemma is premised is likely to be grounded on different beliefs about virtual paedophilia which are different to corresponding beliefs about virtual murder. This means that there is not just one reason for why a difference in moral attitude exists but potentially many. Consequently, looking for a single morally relevant factor for why the dilemma exists is futile, as is challenging any attempt at resolving the dilemma based on the same single factor. Therefore, in trying to resolve the gamer's dilemma, I have presented CEE as a means of accounting for why this shared moral attitude towards virtual paedophilia exists and why this differs (typically) from our shared moral attitude towards virtual murder.

In conclusion, according to CEE, ultimately, the premise on which the gamer's dilemma is built is the product of different beliefs and not differences in some mind-independent moral fact that differentiate virtual paedophilia

from virtual murder. Beliefs, as has been discussed, do not have to be true; but, once held, one's moral attitude needs to be consistent with them. Attempts at resolving the gamer's dilemma have previously failed because they have targeted single factors (*one* morally relevant difference). What I have argued is that the premise on which the dilemma is built stems from a difference in attitude which itself is not based on a single factor or a single morally relevant difference. To resolve the dilemma, one would need to undermine each or a large number of the different beliefs which ground the single moral attitude (objectified social norm). This, in itself, may be difficult to achieve given that the beliefs are likely to be based on a particular interpretation of the intentional object which, for some, may be less amenable to reinterpretation, particularly in the absence of contradictory evidence.

As a normative ethic, CEE posits an objectified moral norm that is constructed within a given society in virtue of a (majority) shared attitude. What is morally acceptable is therefore based on what the majority consider to be morally acceptable in virtue of their approving moral attitude. A *de re* attitude may be challenged and change over time, but only when a sufficient number of beliefs on which the *de re* attitude is grounded change. Where a number of different beliefs contribute to a particular attitude and where some/all of these are difficult to undermine, the attitude and hence the social norm will be maintained. Of course, this potentially limits the universality of the normative ethic, given that different societies may have different beliefs and therefore attitudes – certainly in the case of virtual enactments – but this closely matches our moral reality and is not therefore reason enough to dismiss CEE.

NOTES

1. I appreciate that there may be occasions when one is both amused and morally disapproves: say, when finding a joke amusing despite disapproving, morally, of the inherent sexisms.
2. An anaphoric reference occurs when a word in a text refers to a previous idea in the text for its meaning. In the sentence "Fred always looked unkempt but this never seemed to bother him", the word 'him' makes anaphoric reference to Fred.
3. My use of Dunn's term is slightly different to his original usage.

REFERENCES

Adachi, P. J. C., & Willoughby, T. (2011). The effect of video game competition and violence on aggressive behavior: Which characteristic has the greatest influence? *Psychology of Violence*, *1*(4), 259–274.

Adams, A. A. (2010). Virtual sex with child avatars. In C. Wankel & S. Malleck (eds.), *Emerging issues in virtual worlds* (pp. 55–72). Charlotte: Information Age Publishing.

Adler, A. (1996). Photography on trial. *Index on Censorship*, *25*(3), 141–146.

Ali, R. (2015). A new solution to the gamer's dilemma. *Ethics and Information Technology*, *17*(4), 267–274.

Alia-Klein, N., Wang, G. J., Preston-Campbell, R. N., Moeller, S. J., Parvaz, M. A., Zhu, W. et al. (2014). Reactions to media violence: It's in the brain of the beholder. *PLoS One*, *9*(9), e107260.

Anderson, C. A., & Bushman, B. J. (2002). The effects of media violence on society. *Science*, *295*, 2377–2379.

Anderson, C. A., Shibuya, A., Ihori, N., Swing, E. L., Bushman, B. J., Sakamoto, A. et al. (2010). Violent video game effects on aggression, empathy, and prosocial behavior in eastern and western countries: A meta-analytic review. *Psychological Bulletin*, *136*(2), 151–173.

Aristotle. (1976). *The Nicomachean ethics*. J.A.K. Thomson (trans.) Harmondsworth: Penguin.

Bader, M. (2003). *Arousal: The secret logic of sexual fantasies*. London: Virgin Books.

Barlett, C. P., & Rodeheffer, C. (2009). Effects of realism on extended violent and nonviolent video game play on aggressive thoughts, feelings, and physiological arousal. *Aggressive Behavior*, *35*, 213–224.

Bartel, C. (2012). Resolving the gamer's dilemma. *Ethics and Information Technology*, *14*(1), 11–16.

Bartel, C. (2015). Free will and moral responsibility in video games. *Ethics and Information Technology*, *17*(4), 285–293.

Bartle, R. (2008). Torture. *The everyday blog of Richard Bartle*, 19 November 2008. Retrieved 14 July 2016 from: http://www.youhaventlived.com/qblog/2008/QBlog191108A.html

Baudrillard, J. (1983). *Simulations*. New York: Semiotext(e).

Beech, A. R., Elliot, I. A., Birgden, A., & Findlater, D. (2008). The internet and child sexual offending: A criminological review. *Aggression and Violent Behaviour*, *13*, 216–228.

Bensley, L., & Van Eenwyk, J. (2001). Video games and real-life aggression: Review of the literature. *Journal of Adolescent Health*, *29*, 244–257.

Bentham, J. (1830). *Rationale of Reward*. London: Robert Heward. http://books.google.co.uk/books?id=L5Q7AAAAYAAJ&printsec=frontcover&source=gbs_ge_summary_r&cad=0#v=onepage&q&f=false. Accessed 3 March 2016.

Berlin, F. S., & Sawyer, D. (2012). Potential consequences of accessing child pornography over the internet and who is accessing it. *Sexual Addiction & Compulsivity: The Journal of Treatment & Prevention*, *19*(1–2), 30–40.

Bird, P. (2011). Virtual child pornography and the constraints imposed by the first amendment. *Barry Law Review*, *16*(1), 161–176.

Brenick, A., Henning, A., Killen, M., O'Connor, A., & Collins, M. (2007). Social evaluations of stereotypic images in video games unfair, legitimate, or 'just entertainment'? *Youth & Society*, *38*(4), 395–419.

Bryant, P., & Linz, D. G. (2008). The effects of exposure to virtual child pornography on viewer cognition and attitudes toward deviant sexual behavior. *Communication Research*, *35*(1), 3–38.

Bushman, B. J., & Huesmann, L. R. (2006). Short-term and long-term effects of violent media on aggression in children and adults. *Archives of Pediatrics and Adolescent Medicine*, *160*, 348–352.

Bushman, B. J., & Huesmann, L. R. (2014). Twenty-five years of research on violence in digital games and aggression revisited: A reply to Elson and Ferguson (2013). *European Psychologist*, *19*(1), 47–55.

Bushman, B. J., Rothstein, H. R., & Anderson, C. A. (2010). Much ado about something: Violent video game effects and a school of red herring: Reply to Ferguson and Kilburn (2010). *Psychological Bulletin*, *136*(2), 182–187.

Bushman, B. J., Gollwitzer, M., & Cruz, C. (2015). There is broad consensus: Media researchers agree that violent media increase aggression in children, and pediatricians and parents concur. *Psychology of Popular Media Culture*, *4*(3), 200–214.

Campbell, C. (2014). The worst trailer of the year revels in slaughtering innocents. *Polygon*, 16 October 2014. Retrieved 15 July 2016, from: http://www.polygon.com/2014/10/16/6988687/the-worst-trailer-of-the-year-revels-in-slaughtering-innocents

Cappuccio, A. F. (2012). Article 10 and restrictions on pornography: Shifting the rationale from protection of morals to cultural harm. *Queen Mary Law Journal*, *2*, 1–14.

Casey, M. (2015). Do violent video games leave to criminal behaviour? *CBS News*, 17 August 2015. Retrieved 25 February 2016 from: http://www.cbsnews.com/news/do-violent-video-games-lead-to-criminal-behavior/

Castronova, E. (2004). The right to play. *New York Law School Law Review*, *49*, 185–210.

Cohen, N. (2016). The vocabulary of manga: Visual morphology in dialects of Japanese visual language. *Journal of Pragmatics*, *92*, 17–29.

Copp, D. (2001). Realist-expressivism: A neglected option for moral realism. *Social Philosophy and Policy*, *18*, 1–43.

Copp, D. (2011). Jesse Prinz, the emotional construction of morals (Oxford: Oxford University Press, 2007): Prinz's subjectivist moral realism. *Noûs*, *45*(3), 577–594.

Crofts, T., & Murray, L. (2013). 'Sexting', children and child pornography. *Sydney Law Review*, *35*, 85–106.

Cunningham, S., Engelstätter, B., & Ward, M. R. (2011). Understanding the effects of violent video games on violent crime. *Ideas*. Retrieved 24 February 2016 from: https://ideas.repec.org/p/zbw/zewdip/11042.html

Cunningham, S., Engelstätter, B., & Ward, M. R. (2016). Violent video games and violent crime. *Southern Economic Journal*, *82*(4), 1247–1265.

Davidson, D (ed.). (1980). Actions, reasons and causes. In D. Davidson (ed.) *Essays on actions and events* (pp. 3–21). Oxford: Clarendon Press.

De Vane, B., & Squire, K. D. (2008). The meaning of race and violence in Grand Theft Auto: San Andreas. *Games and Culture*, *3*, 264–285.

Diamond, M., & Uchiyama, A. (1999). Pornography, rape and sex crimes in Japan. *International Journal of Law and Psychiatry*, *22*(1), 1–22.

Dunn, J. (2012). Virtual worlds and moral evaluation. *Ethics and Information Technology*, *14*(4), 255–265.

Edwards, L., & Waelde, C. (1997). *Law and the internet: Regulating cyberspace*. Oxford: Hart Publishing.

Edwards, S. S. M. (2000). Prosecuting 'child pornography': Possession and taking of indecent photos of children. *Journal of Social Welfare and Family Law*, *22*, 1–21.

Eke, A. W., Seto, M. C., & Williams, J. (2011). Examining the criminal history and future offending of child pornography offenders: An extended prospective follow-up study. *Law and Human Behavior*, *35*, 466–478.

Elliot, M. (1992). Images of children in the media: Soft kiddie porn. In C. Itzin (ed.), *Pornography, women, violence, and civil liberties* (pp. 217–221). Oxford: Oxford University Press.

Elliott, I. A., & Beech, A. R. (2009). Understanding online child pornography use: Applying sexual offense theory to internet offenders. *Aggression and Violent Behavior*, *14*, 180–193.

Elliott, I. A., Beech, A. R., Mandeville-Norden, R., & Hayes, E. (2009). Psychological profiles of internet sexual offenders: Comparisons with contact sexual offenders. *Sexual Abuse: A Journal of Research and Treatment*, *21*, 76–92.

Eneman, M., Gillespie, A. A., & Stahl, B. C. (2009). Criminalising fantasies: The regulation of virtual child pornography (pp. 1–12). Retrieved 7 July 2016 from: http://www.cse.dmu.ac.uk/~bstahl/publications/2009_Criminalising_Fantasies_ECIS.pdf

Feinberg, J. (1988). *The moral limits of criminal law, volume 4: Harmless wrongdoing*. Oxford: Oxford University Press.

Ferguson, C. J. (2007a). Evidence for publication bias in video game violence effects literature: A meta-analytic review. *Aggression and Violent Behavior*, *12*, 470–482.

Ferguson, C. J. (2007b). The good, the bad and the ugly: A metaanalytic review of positive and negative effects of violent video games. *Psychiatric Quarterly*, *78*(4), 309–316.

Ferguson, C. J. (2011). Video games and youth violence: A prospective analysis in adolescents. *Journal of Youth and Adolescence*, *40*, 377–391.

Ferguson, C. J. (2013). Violent video games and the supreme court: Lessons for the scientific community in the wake of Brown v. Entertainment Merchants Association. *American Psychologist*, *68*(2), 57–74.

Ferguson, C. J., & Kilburn, J. (2010). Much ado about nothing: The misestimation and overinterpretation of violent video game effects in eastern and western nations: Comment on Anderson et al. (2010). *Psychological Bulletin*, *136*(2), 174–178.

Galbraith, P. W. (2011). *Lolicon*: The reality of 'Virtual Child Pornography' in Japan. *Image & Narrative*, *12*(1), 83–119.

Gaut, B. (1998). The ethical criticism of art. In J. Levinson (ed.), *Aesthetics and ethics* (pp. 182–203). Cambridge: Cambridge University Press.

Gillespie, A. A. (2008). *Child exploitation and communication technologies*. Dorset: Russell House Publishing.

Glock, S., & Kneer, J. (2009). Game over? The impact of knowledge about violent digital games on the activation of aggression-related concepts. *Journal of Media Psychology*, *21*, 151–160.

Goldblatt, B. (2012). Virtual child pornography: The children aren't real, but the dangers are; Why the Ashcroft Court cot it wrong. *Law School Student Scholarship*, *41*, 1–45.

Goode, S. D. (2010). *Understanding and addressing adult sexual attraction to children: A study of paedophiles in contemporary society*. London: Routledge.

Graham, G. (1999). *The internet: A philosophical inquiry*. London: Routledge.

Greitemeyer, T., & Mügge, O. (2014). Video games do affect social outcomes: A meta-analytic review of the effects of violent and prosocial video game play. Personality and. *Social Psychology Bulletin*, *40*(5), 578–589.

Griffiths, R. P., Eastin, M. S., & Cicchirillo, V. (2016). Competitive video game play: An investigation of identification and competition. *Communication Research*, *43*(4), 468–486.

Gutierrez, R., & Giner-Sorolla, R. (2007). Anger, disgust, and presumption of harm as reactions to taboo-breaking behaviors. *Emotion*, *7*(4), 853–868.

Gwynne, J. (2013). Japan, postfeminism and the consumption of sexual(ised) schoolgirls in male-authored contemporary manga. *Feminist Theory*, *14*(3), 325–343.

Haidt, J., & Hersh, M. A. (2001). Sexual morality: The cultures and emotions of conservatives and liberals. *Journal of Applied Social Psychology*, *31*, 191–221.

Hartley, J. (1998). Juvenation: News, girls and power. In C. Carter, G. Branston, & S. Allan (eds.), *News, gender and power* (pp. 47–70). London: Routledge.

Hartmann, T., & Vorderer, P. (2010). It's okay to shoot a character: Moral disengagement in violent video games". *Journal of Communication*, *60*, 94–119.

Hartmann, T., Toz, E., & Brandon, M. (2010). Just a game? Unjustified virtual violence produces guilt in empathetic players. *Media Psychology*, *13*, 339–363.

Hartmann, T., Krakowiak, K. M., & Tsay-Vogel, M. (2014). How violent video games communicate violence: A literature review and content analysis of moral disengagement factors. *Communication Monographs*, *81*(3), 310–332.

Hayes, J. (2015). 10 Most violent video games of 2015 (and what to play instead). *Common Sense Media*, 26 October 2015. Retrieved 23 February 2016 from: https://www.commonsensemedia.org/blog/10-most-violent-video-games-of-2015-and-what-to-play-instead

Hellmann, M. (2014), Japan finally bans child pornography. *Time*, 18 June 2014. Retrieved 8 July 2016 from: http://time.com/2892728/japan-finally-bans-child-pornography/

Honda, M. (2010). The genealogy of *Hirahira*: Liminality and the girl (T. Aoyama, & B. Hartley, trans.). In T. Aoyama & B. Hartley (eds.), *Girl reading girl in Japan* (pp. 19–37). London: Routledge.

Huesmann, L. R. (2010). Nailing the coffin shut on doubts that violent video games stimulate aggression: Comment on Anderson et al. (2010). *Psychological Bulletin*, *136*(2), 179–181.

Huesmann, L. R., Moise-Titus, J., Podolski, C., & Eron, L. D. (2003). Longitudinal relations between children's exposure to TV violence and their aggressive and violent behavior in young adulthood: 1977–1992. *Developmental Psychology*, *39*, 201–221.

Hunter, I., Saunders, D., & Williamson, D. (1993). *On pornography: Literature, sexuality and obscenity law*. London: Macmillan.

Jansz, J. (2005). The emotional appeal of violent video games for adolescent males. *Communication Theory*, *15*(3), 219–241.

Jewkes, Y., & Wykes, M. (2005). Reconstructing the sexual abuse of children: 'Cyber-paeds', panic and power. *Sexualities*, *15*(8), 934–952.

Juul, J. (2005). *Half-real: Video games between real rules and fictional games.* Cambridge, MA: MIT Press.

Kalim, A. (2013). Addressing the gap in international instruments governing internet child pornography. *CommLaw Conspectus*, *21*, 428–452.

Karnold, S. (2000). The cyber world of child pornography and the Child Pornography Act of 1996: Thoughts on morphing, virtual imaging and the First Amendment. *Journal of Information Ethics*, *2*, 60–65.

Kass, L. R. (2002). *Life, liberty, and the defense of dignity: The challenge for bioethics.* San Francisco: Encounter Books.

King, P. J. (2008). No plaything: Ethical issues concerning child-pornography. *Ethical Theory and Moral Practice*, *11*(3), 327–345.

Kingston, D. A., Fedoroff, P., Firestone, P., Curry, S., & Bradford, J. M. (2008). Pornography use and sexual aggression: The impact of frequency and type of pornography use on recidivism among sexual offenders. *Aggressive Behavior*, *34*(4), 341–351.

Kinsella, S. (1998). Japanese subculture in the 1990s: Otaku and the amateur manga movement. *The Journal of Japanese Studies*, *24*(2), 289–316.

Klimmt, C., Schmid, H., Nosper, A., Hartmann, T., & Vorderer, P. (2006). How players manage moral concerns to make video game violence enjoyable. *Communications*, *31*, 309–328.

Klimmt, C., Schmid, H., Nosper, A., Hartmann, T., & Vorderer, P. (2008). Moral management: Dealing with moral concerns to maintain enjoyment of violent video games. In A. Sudmann-Jahn & R. Stockmann (eds.), *Computer games as a sociocultural phenomenon: Games without frontiers – wars without tears* (pp. 108–118). Hampshire: Palgrave.

Konijn, E. A., & Hoorn, J. F. (2005). Some like it bad: Testing a model for perceiving and experiencing fictional characters. *Media Psychology*, *7*(2), 107–144.

Konijn, E. A., Walma Van Der Molen, J. H., & Hoorn, J. F. (2011). Babies versus bogeys: In-game manipulation of empathy in violent video games. In K. Poels & S. Malliet (eds.), *Vice city virtue: Moral issues in digital game play* (pp. 151–176). Leuven: Acco Academic.

Kontominas, B. (2008). Simpsons cartoon rip-off is child porn: Judge. *The Sydney Morning Herald*, 8 December 2008. Retrieved 9 July 2016 from: http://www.smh.com.au/news/national/simpsonsstyle-cartoon-is-child-porn/2008/12/08/1228584707575.html

Kornegay, J. N. (2006). Protecting our children and the constitution: An analysis of the "virtual" child pornography provisions of the PROTECT Act of 2003. *William & Mary Law Review*, *47*(6), 2129–2167.

Kosse, S. H. (2004). Virtual child pornography – A United States update. *Communications Law*, *9*(2), 39–46.

Krahé, B. (2014). Restoring the spirit of fair play in the debate about violent video games: A comment on Elson and Ferguson (2013). *European Psychologist*, *19*(1), 56–59.

Krcmar, M., Farrar, K., & McGloin, R. (2011). The effects of video game realism on attention, retention and aggressive outcomes. *Computers in Human Behavior*, *27*(1), 432–439.

Ladas, M. (2003). Eine Befragung von 2141 Computerspielern Zuwirkung und Nutzung von Gewalt [A survey of 2,141 computer game players on effect and use of violence]. In F. Rötzer (ed.), *Virtuelle Welten − Reale Gewalt* [Virtual worlds – real violence] (pp. 26–35). Hannover, Germany: Hans Heise Verlag.

Lau, C. Y. (2013). Manga drawings as stereotypes aesthetics. *Visual Arts Research*, *39*(2, issue 77), 42–53.

Levinson, J. (2005). Erotic art and pornographic pictures. *Philosophy and Literature*, *29*, 228–240.

Levy, N. (2002). Virtual child pornography: The eroticization of inequality. *Ethics and Information Technology*, *4*(4), 319–323.

Lichty, P. (2009). The translation of art in virtual worlds. *Leonardo Electronic Almanac*, *16*(4–5), 1–12.

Lodato, V. (1998). Computer-generated child pornography – exposing prejudice in our First Amendment jurisprudence? *Seton Hall Law Review*, *28*(4), 1328–1364.

Longino, H. E. (1995). Pornography, oppression, and freedom: A closer look. In S. Dwyer (ed.), *The problem of pornography* (pp. 34–47). Belmont: Wadsworth.

Luck, M. (2009). The gamer's dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology*, *11*, 31–36.

Luck, M., & Ellerby, N. (2013). Has Bartel resolved the gamer's dilemma? *Ethics and Information Technology*, *15*(3), 229–233.

MacKinnon, C. A. (1991). Pornography as defamation and discrimination. *Boston University Law Review*, *71*, 793–815.

Mag Uidhir, C. (2009). Why pornography can't be art. *Philosophy and Literature*, *33*, 193–203.

Malamuth, N., & Huppin, M. (2007). Drawing the line on virtual child pornography: Bringing the law in line with the research evidence. *New York University Review of Law and Social Change*, *31*(4), 773–827.

Markey, P. M., Markey, C. N., & French, J. E. (2015). Violent video games and real-world violence: Rhetoric versus data. *Psychology of Popular Media Culture*, *4*(4), 277–295.

Masuchika, G. (2015). Japanese cartoons, virtual child pornography, academic libraries, and the law. *Reference and User Services Quarterly*, *54*(4), 54–60.

McCarthy, J. A. (2010). Internet sexual activity: A comparison between contact and non-contact child pornography offenders. *Journal of Sexual Aggression*, *16*(2), 181–195.

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: A lost opportunity. *Criminal Law Review*, *4*, 245–260.

McLelland, M. (2005). The world of YAOI: The internet, censorship and the global 'boys' love' fandom. *The Australian Feminist Law Journal*, *23*, 61–77.

McLelland, M., & Yoo, S. (2007). The international Yaoi boys' love fandom and the regulation of virtual child pornography: Current legislation and its implications. *Journal of Sexuality Research and Social Policy*, *4*(1), 93–104.

McMahan, J. (2006). On the moral equality of combatants. *Journal of Political Philosophy*, *14*(4), 377–393.

Moehringer, J. R. (1998). Child porn fight focuses on 2 photographers' books. *Los Angeles Times*, 8 March 1998. Retrieved 7 July 2016 from: http://articles.latimes.com/1998/mar/08/news/mn-26778

Mota, S. A. (2002). The U.S. Supreme Court addresses the Child Pornography Prevention Act and Child Online Protection Act in *Ashcroft v. Free Speech Coalition* and. *Ashcroft v. American Civil Liberties Union*. *Federal Communications Law Journal*, *55*(1), 85–98.

Nagel, T. (1979). *Mortal questions*. Cambridge: Cambridge University Press.

Neutze, J., Seto, M.C., Schaefer, G.A., Mundt, I.A., & Beier, K.M. (2011). Predictors of child pornography offenses and child sexual abuse in a community sample of pedophiles and hebphiles. *Sexual Abuse: A journal of Research and Treatment*, *23*(2), 212–242.

Nichols, S. (2008). Sentimentalism naturalized. In W. Sinnott-Armstrong (ed.), *Moral psychology: The evolution of morality, volume 2* (pp. 255–274). Cambridge, MA: MIT Press.

Niveau, G. (2010). Cyber-pedocriminality: Characteristics of a sample of internet child pornography offenders. *Child Abuse & Neglect*, *34*, 570–575.

Norrie, J. (2010). Hardcore comics sidestep porn law. *The Sydney Morning Herald*, 20 February 2010. Retrieved 8 July 2016 from: http://www.smh.com.au/world/hardcore-comics-sidestep-porn-law-20100219-olu8.html

Norris, C. (2009). Manga, anime and visual art culture. In Y. Sugimoto (ed.), *The Cambridge companion to modern Japanese Culture* (pp. 236–260). Cambridge: Cambridge University Press.

Nys, T. (2010). Virtual ethics. *Ethical Perspectives*, *17*(1), 79–93.

Office, H. (2007). *Consultation on possession of non-photographic visual depictions of child abuse*. http://www.gov.scot/Resource/Doc/1099/0048474.pdf. (Accessed 9/08/2016)

Orland, K. (2013). Obama looks to fund study into link between gun violence, video games and 'media'. *Gaming*, 17 January 2013. Retrieved 25 February 2016 from: http://www.wired.co.uk/news/archive/2013-01/17/obama-on-video-games

Ortega-Brena, M. (2009). Peek-a-boo, I see you: Watching Japanese hard-core Animation. *Sexuality & Culture*, *13*, 17–31.

Ost, S. (2010). Criminalising fabricated images of child pornography: A matter of harm or morality? *Legal Studies*, *30*(2), 230–256.

Oswell, D. (2006). When images matter: Internet child pornography, forms of observation and an ethics of the virtual. *Information Communication & Society*, *9*(2), 244–265.

Palmer, T. (2005). Behind the screen: Children who are the subjects of abusive images. In E. Quayle & M. Taylor (eds.), *Viewing child pornography on the internet: Understanding the offence, managing the offender, helping the victims* (pp. 61–74). London: Russell House Publishing.

Patridge, S. (2011). The incorrigible social meaning of video game imagery. *Ethics and Information Technology*, *13*(4), 303–312.

Patridge, S. (2013a). Exclusivism and evaluation: Art, erotica, and pornography. In H. Maes (Ed.), *Pornographic art and the aesthetics of pornography* (pp. 43–57). Basingstoke: Palgrave Macmillan.

Patridge, S. L. (2013b). Pornography, ethics, and video games. *Ethics and Information Technology*, *15*(1), 25–34.

Pohl, K. (2008). Ethical reflection and involvement in computer games. In S. Günzel, M. Liebe, & D. Mersch (eds.), *Conference proceedings of the philosophy of computer games, 2008* (pp. 92–107). Berlin: Potsdam University Press.

Powers, T. M. (2003). Real wrongs in virtual communities. *Ethics in Information Technology*, *5*, 191–198.

Prigg, M. 2009. "Violent video game breaks records with 4.7m sales in a day". *London Evening Standard*, 13 November 2009. Retrieved 24 February 2016 from: http://www.standard.co.uk/news/violent-video-game-breaks-records-with-47m-sales-in-a-day-6705082.html

Prinz, J. J. (2007). *The emotional construction of morals*. Oxford: Oxford University Press.

Quayle, E., & Taylor, M. (2002). Child pornography and the internet: Perpetuating a cycle of abuse. *Deviant Behaviour*, *23*(4), 331–362.

Ray, V. R., Kimonis, E. R., & Donoghue, C. (2010). Legal, ethical, and methodological considerations in the internet-based study of child pornography offenders. *Behavioral Sciences and the Law*, *28*, 84–105.

Rea, M. (2001). What is pornography? *Nous*, *35*(1), 118–145.

Ridge, M. (2006). Ecumenical expressivism: Finessing Frege. *Ethics*, *116*(2), 302–336.

Riegel, D. L. (2004). Effects on boy-attracted pedosexual males of viewing boy erotica. *Archives of Sexual Behavior*, *33*(4), 321–323.

Rogers, A. (2009). Protecting children on the internet: Mission impossible? *Baylor Law Review*, *61*(2), 323–356.

Rubenking, B., & Lang, A. (2014). Captivated and grossed out: An examination of processing core and sociomoral disgusts in entertainment media. *Journal of Communication*, *64*(3), 543–565.

Russell, G. (2008). Pedophilies in Wonderland: Censoring the sinful in cyber-space. *Journal of Criminal Law and Criminology*, *98*(4), 1467–1500.

Ryder, B. (2003). The harms of child pornography law. *UBC Law Review*, *36*(1), 101–135.

Sabin, R. (1993). *Adult comics*. London: Routledge.

Samenow, C. P. (2012). Child pornography and the law: A clinician's guide. Sexual addiction & compulsivity. *The Journal of Treatment & Prevention*, *19*(1–2), 16–29.

Sandin, P. (2004). Virtual child pornography and utilitarianism. *Journal of Information, Communication & Ethics in Society*, *2*(1), 217–223.

Scheeres, J. (2002). When kid porn isn't kid porn, *Wired*, 8 May 2002. Retrieved 15 July from: http://www.wired.com/2002/05/when-kid-porn-isnt-kid-porn/?currentPage=2

Schreler, J. (2015). Why most video game 'aggression' studies are nonsense. *Kotaku*, 14 August 2015. Retrieved 25 February 2016 from: http://kotaku.com/why-most-video-game-aggression-studies-are-nonsense-1724116744

Schulzke, M. (2011). Reflective play and morality: Video games as thought experiments. In K. Poels & S. Malliet (eds.), *Vice city virtue: Moral issues in digital game play* (pp. 51–68). Leuven: Acco Academic.

Seto, M. C., & Eke, A. W. (2005). The criminal histories and later offending of child pornography offenders. *Sexual Abuse*, *17*(2), 201–210.

Seto, M. C., Reeves, L., & Jung, S. (2010). Explanations given by child pornography offenders for their crimes. *Journal of Sexual Aggression*, *16*(2), 169–180.

Seto, M. C., Wood, J. M., Babchishin, K. M., & Flynn, S. (2012). Online solicitation offenders are different from child pornography offenders and lower risk contact sexual offenders. *Law and Human Behavior*, *36*(4), 320–330.

Shigematsu, S. (1999). Dimensions of desire: Sex, fantasy, and fetish in Japanese comics. In J. A. Lent (ed.), *Themes in Asian cartooning: Cute, cheap, mad, and sexy* (pp. 127–163). Bowling Green: Bowling Green State University Popular Press.

Sicart, M. (2009). *The ethics of computer games*. Cambridge, MA: MIT Press.

Simpson, B. (2009). Controlling fantasy in cyberspace: Cartoons, imagination and child pornography. *Information & Communications Technology Law*, *18*(3), 255–271.

Singer, P. (2007). Video crime peril vs. virtual pedophilia. *The Japanese Times*, 22 July 2007. Retrieved 4 March 2016, from: http://search.japantimes.co.jp/cgi-bin/eo20070722a1.html

Singh, A. (2009). Brooke Shields nude photograph causes controversy at Tate exhibition. *The Telegraph*, 29 September 2009. Retrieved 6 February 2016 from: http://www.telegraph.co.uk/news/newstopics/celebritynews/6244330/Brooke-Shields-nude-photograph-causes-controversy-at-Tate-exhibition.html

Sullivan, J., & Beech, A. R. (2003). Are collectors of child abuse images a risk to children? In A. MacVean & P. Spindler (eds.), *Policing paedophiles on the internet* (pp. 11–20). London: The New Police Bookshop.

Tate, T. (1992). The child pornography industry: International trade in child sexual abuse. In C. Itzin (ed.), *Pornography: Women, violence and civil liberties* (pp. 203–216). Oxford: Oxford University Press.

Tavinor, G. (2009). *The art of videogames*. New York: Wiley.

Taylor, J. (2009). Graphic artists condemn plans to ban erotic comics. *The Independent*. Monday 23 March 2009. Retrieved 9 July, 2016 from: http://www.independent.co.uk/news/uk/politics/graphic-artists-condemn-plans-to-ban-erotic-comics-1652270.html

Vanacker, B., & Heider, D. (2012). Ethical harm in virtual communities. *Convergence: The International Journal of Research into New Media Technologies*, *18*(1), 71–84.

Veber, M. (2004). Virtual child pornography. *Public Affairs Quarterly*, *18*(1), 75–90.

Walzer, M. (1977). *Just and unjust wars: A moral argument with historical illustrations*. New York: Basic Books.

Webb, L., Craissati, J., & Keen, S. (2007). Characteristics of internet child pornography offenders: A comparison with child molesters. *Sexual Abuse*, *19*, 449–465.

White, A. (2006). *Virtually obscene: The case for an uncensored internet*. London: McFarland & Company, Inc.

Whitty, M. T., Young, G., & Goodings, L. (2011). What I won't do in pixels: Examining the limits of taboo violation in MMORPGs. *Computers in Human Behavior*, *27*(1), 268–275.

Wilkinson, E. (2011). 'Extreme pornography' and the contested spaces of virtual citizenship. *Social & Cultural Geography*, *12*(5), 493–508.

Williams, K. S. (2004). Child pornography law: Does it protect children? *Journal of Social Welfare and Family Law*, *26*(3), 245–261.

Williams, N. (1991). *False images: Telling the truth about pornography*. London: Kingsway Publications.

Williams, Z. (2003). Panic on the screens. *The Guardian*, 14 January, 2003. Retrieved 6 February 2016 from: http://www.guardian.co.uk/media/2003/jan/14/newmedia.childrensservices

Wilson, B. (1999). Becoming Japanese: *Manga*, children's drawings, and the construction of national character. *Visual Arts Research*, *25*(2, issue 50), 48–60.

Wood, R. T. A., Griffiths, M. D., Chappell, D., & Davies, M. N. O. (2004). The structural characteristics of video games: A psycho-structural analysis. *CyberPsychology & Behavior*, *7*(1), 1–10.

Wright, P. J., Tokunaga, R. S., & Kraus, A. (2016). A Meta-Analysis of pornography consumption and actual acts of sexual aggression in general population studies. *Journal of Communication*, *66*(1), 183–205.

Young, G. (2013a). Enacting taboos as a means to an end; but what end? On the morality of motivations for child murder and paedophilia within gamespace. *Ethics and Information Technology*, *15*(1), 13–23.

Young, G. (2013b). *Ethics in the virtual world: The morality and psychology of gaming*. London: Routledge.

Young, G. (2014). A meta-ethical approach to single-player gamespace: Introducing constructive ecumenical expressivism as a means of explaining why moral consensus is not forthcoming. *Ethics and Information Technology*, *16*(2), 91–102.

Young, G. (2015a). Are there some things it is morally wrong to make-believe? An examination of imaginative resistance as a measure of the morality of pretence. *Cogent Arts and Humanities*, *2*(1), 1–14.

Young, G. (2015b). Violent video games and morality: A meta-ethical approach. *Ethics and Information Technology*, *17*(4), 311–321.

Young, G., & Whitty, M. T. (2011). Should gamespace be a taboo-free zone? Moral and psychological implications for single-player video games. *Theory and Psychology*, *21*(6), 802–820.

Young, G., & Whitty, M. T. (2012). *Transcending taboos: A moral and psychological examination of cyberspace*. London: Routledge.

Young, K. S. (2001). *Tangled in the web: Understanding cybersex from fantasy to addiction*. Bloomington: Authorhouse.

Zagal, J. P. (2009). Ethically notable videogames: Moral dilemmas and gameplay. In B. Atkins, T. Krzywinska, & H. Kennedy (eds.) *Proceedings of the Digital Games Research Association international conference (DiGRA)*, 2009. London: DiGRA.

Žižek, S. (1997). *The plague of fantasies*. London: Verso.

Zumbach, J., Seitz, C., & Bluemke, M. (2015). Impact of violent video game realism on the self-concept of aggressiveness assessed with explicit and implicit measures. *Computers in Human Behavior*, *53*, 278–288.

INDEX