



Is it distinctively wrong to simulate doing wrong?

John Tillson¹

Published online: 26 June 2018
© Springer Nature B.V. 2018

Abstract

This paper is concerned with whether there is a moral difference between simulating wrongdoing and consuming non-simulatory representations of wrongdoing. I argue that simulating wrongdoing is (as such) a *pro tanto* wrong whose wrongness does not tarnish other cases of consuming representations of wrongdoing. While simulating wrongdoing (as such) constitutes a disrespectful act, consuming representations of wrongdoing (as such) does not. I aim to motivate this view in part by bringing a number of intuitive moral judgements into reflective equilibrium, and in part by describing the case of a character that I call The Devious Super Geek who simulates wrong to particular people that he knows personally. I build bridging cases from the case of the Devious Super Geek to capture games in which one simulates wrong to imaginary members of extant, morally salient categories. The surprising conclusions that we are led to include not just that simulated wrongdoing is *pro tanto* wrong, but that simulated Just killing is *pro tanto* wrong, and also that the simulated killing of zombies and aliens is also *pro tanto* wrong. Finally, I describe how I propose to handle some potential objections and attempt to weigh the *pro tanto* wrong identified in the paper against some countervailing considerations in some all things considered judgements.

Keywords The gamer's dilemma · Simulation · Wrongdoing · Reflective equilibrium · Videogames

Introduction

Videogames are a multibillion dollar industry. SuperData report that in 2017, the Global Interactive Media Market was worth 105 billion dollars, seeing a market growth of 12% since 2016. It further predicts that by 2020, the Virtual Reality Market (including hardware and software) will be worth 28.3 billion dollars.¹ As money is poured into research and development for new technologies, the experiences that games companies are able to offer their customers are becoming more immersive, textured and realistic. Virtual reality headsets, complete with stereo sound, stereoscopic vision, and head and eye tracking sensors, offer gamers the opportunity to navigate fictional environments from a first personal perspective, using their regular patterns of movement to so do. At the same time, many videogames afford gamers the opportunity to indulge in simulating serious wrongdoing, including, to mention just a few: harvesting children to gain power-ups (*Bioshock*), committing murder

for a series of gruesome snuff films (*Manhunt*), and crushing enemies' skulls with found objects (*Condemned: Criminal Origins*). As such technology and narratives develop apace, games companies will reach a point where they can offer gamers immersion in alternate realities within which they can act out detailed fantasies of serious wrongdoing, arguably without violating the constraints of morality such as the prohibition on doing harm. It is with this prospect in mind that I ask whether it is (as such) *pro tanto* wrong for a player to direct his character to perpetrate immoral action types within a videogame, and, more generally, whether it is (as such) *pro tanto* wrong for an individual to simulate perpetrating immoral action types.² I will argue that the answer to these questions is yes. While I will be concerned with simulating wrongdoing generally, I will focus the following discussion around simulated wrongdoing within videogames

✉ John Tillson
tillsoj@hope.ac.uk

¹ Department of Education Studies, Liverpool Hope University, Liverpool, UK

¹ <https://www.superdataresearch.com/market-data/market-brief-year-in-review/> [last accessed 10 Aug 2017].

² It may help some readers to consider how actions can be wrong as such or wrong given contextual features. An armed robber's squeezing their index finger is not wrong as such, or not wrong as the squeezing of their index finger. It is instead wrong given contextual features, namely as being the causal precipitate of murder by gunshot. On the other hand, committing murder is wrong as such, it is wrong as murder.

as a particular instance of this (partly because that is what much relevant and interesting discussion has focused on).

We should distinguish at the outset between ethical issues regarding the consumption, production and regulation of representations of wrongdoing. This paper is concerned only with ethical issues regarding the *consumption* of representations of wrongdoing. In particular it is concerned with whether there is a moral difference between simulating wrongdoing and consuming non-simulatory representations of wrongdoing.³ I argue that simulating wrongdoing is (as such) a *pro tanto* wrong whose wrongness does not tarnish other cases of consuming representations of wrongdoing. Other examples of consuming representations of fictional wrongdoing include watching depictions and reading accounts of, among other things, fictional murder and rape in George R.R. Martin's *A Song of Ice and Fire* and the *Game of Thrones* television series based thereon. Examples of consuming representations of non-fictional wrongdoing include reading historical accounts of atrocities, and watching films like *Hotel Rwanda*⁴ and *Schindler's List*⁵ which depict them. Other examples of consuming representations of wrongdoing, include experiencing representations of wrongdoing unfolding within the narrative of a videogame that one is playing, without these being wrongs that one is simulating oneself. For example, consider playing a character whose village is suddenly subjected to a violent attack, whose only in-game options are non-violent escape or death. Consuming representations of wrongdoing, I contend, does not (as such) constitute a *pro tanto* wrong, although one may reasonably contend that there are wrongful ways of consuming representations of wrongdoing.

Indeed, one might contend that if anything matters morally about the consumption of representations of wrongdoing (whether simulatory or not), then it is the way in which one consumes them. For instance, perhaps what matters is whether one derives some sadistic or otherwise base or reprehensible pleasure from it. If one derives pride, or inspiration from watching the Nazi propaganda film, *Triumph des Willens*⁶ or the Ku Klux Klan puff piece, *The Birth of a Nation*,⁷ then one is at fault. Perhaps one does wrong to indulge or cultivate such sentiments, when one should aim to eliminate them as flaws about oneself. All the same, while production of these films seems morally objectionable, there

seem to be perfectly respectable and even desirable ways of consuming them. Alternatively, one might contend that nothing about how we consume representations of wrongdoing (whether simulatory or non-simulatory) matters morally at all, for nobody, the argument would continue, has any interest in whether and how others are to consume representations of wrongdoing (at least not as such).

I will begin my argument by distinguishing simulatory and non-simulatory representations before briefly discussing the state of empirical literature on the matter of whether playing videogames has negative effects on behaviour. The discussion is brief since the argument which I develop does not turn on whether or not simulating wrongdoing damages one's character or has negative effects on behaviour. Next I introduce Morgan Luck's 'Gamer's Dilemma' which contrasts the common and seemingly incompatible pair of intuitions that, on the one hand, it is *not* morally permissible to commit virtual paedophilia, but, on the other hand, it *is* morally permissible to commit virtual murder. I seek to dissolve this dilemma by arguing that it is in fact wrong to commit virtual murder, because it is wrong to simulate wrongdoing more generally. Following a procedure of reflective equilibrium to reach this conclusion, I bring to light a range of other relevant judgements, seek a general principle which explains the most and revises the fewest of these with maximal economy, and further seek an error theory for the erroneous judgement in Luck's dilemma. On the way, I judge an Aristotelian attempt to identify the wrong in simulated wrongdoing as inadequate. Ultimately, I argue that simulating wrongdoing is (as such) a *pro tanto* wrong by virtue of constituting a disrespectful act, a view which I motivate through discussion of a character that I call The Devious Super Geek. Simulating wrongdoing, I argue, is a wrong which does not turn on any effects. I suggest tentatively that the occasions on which we are apt to not share this judgement intuitively are plausibly due to the fact that our culture has inoculated us against those judgements. Next I extend the lines of argument to motivate the views that simulating killing more generally, even when the targets are fictional aliens or zombies, or even when the violence one simulates would be justified in real life, is, as such, *pro tanto* wrong by virtue of constituting a disrespectful act. Before summarising the argument I then consider a range of objections and discuss whether the *pro tanto* wrong of simulating wrongdoing can be easily trumped in some *all things considered* cases.

³ I acknowledge a slight awkwardness in drawing the distinction, since consuming videogames and simulating wrongdoing more generally seem to be slightly more authorial on the player's part than consuming non-simulatory representations of wrongdoing.

⁴ George (2005).

⁵ Spielberg (1993).

⁶ Riefenstahl (1935).

⁷ Griffith (1915).

Simulatory and non-simulatory representations

Videogames seem to form a sub-section of the larger category of simulations, and simulations form a sub-section of the larger category of representations. However, it is trickier to say what distinguishes simulatory representations from non-simulatory ones. The account of simulation that I will provide does not and need not capture every case of what are commonly called simulations. Instead I will delineate a specific sense of ‘simulation’. A representation is a simulation in the sense I mean, iff it is agential in the sense that a player selects actions to be undertaken by the player’s character(s) (i.e. it comprises agential affordances). Put another way, simulations comprise agential affordances for at least one player (usually via an avatar). To simulate wrongdoing is to direct one’s character to perpetrate some wrong (something that would be wrong if it were done outside of a simulation). Ordinarily there are a range of action options which the player can select from, but even if there is a choice between selecting just one action and not playing, this could still count as simulating wrongdoing in that one initiates the character’s action. Part of simulating an action on my understanding is an interactive interface, so that it is not merely imagined. Instead, there is some external information system that affords actions and is updated with information of one’s having acted.

I do not build into this distinction between simulatory and non-simulatory representations any requirements for or claims about the way in which they are consumed. In particular, it does not matter for the distinction whether one identifies with their avatar, or identifies with them any more than they might identify with the lead character of a film. The distinction turns entirely on whether or not one is able to direct a character to undertake actions. It does not matter whether the experience is immersive, affords one a sense of ‘presence’, or whether one has a sense of embodiment within the game.⁸ Furthermore, I will not make or rely on any judgements about differences that consuming one or the other has on our character or behaviour.⁹

⁸ For further reading, see Sanchez-Vives and Slater (2005) on the concept of ‘presence’, namely “the phenomenon of behaving and feeling as if we are in the virtual world created by computer displays” (p. 332). On embodiment, see Won et al. (2015). On involuntary biological reactions resulting from immersive identification with an avatar, see Fox et al. (2012). On “the relationship effect of immersive system technology on user experiences of presence”, see Cummings and Bailenson (2016).

⁹ An anonymous reviewer has kindly pointed out research indicating there to be little difference between the effects that consuming simulatory and non-simulatory representations have on consumers’ character and behaviour (Anderson et al. 2003).

The claim I argue for, is that for a player to direct their character to perpetrate an immoral action type within a videogame is, *as such*, a *pro tanto* wrong quite apart from the player’s motives and purposes, and quite apart from any effects doing so might have on their character or subsequent behaviours. To simulate wrongdoing more generally is also *as such*, a *pro tanto* wrong quite apart from the player’s motives and purposes and quite apart from any effects doing so might have on their character or subsequent behaviours. I will argue that the same is not true of consuming representations of wrongdoing *tout court*. Due to the “quite apart from any effects” clause of my claim, empirical research into the question of whether simulating wrong doing has detrimental effects on players’ character and behaviour is immaterial to my argument. Information about the effects of simulating wrongdoing on character and behaviour might well make a difference to whether particular instances of simulating wrongdoing are *all things considered* wrong, and we shall return to one consideration motivating this thought at the end. All the same, some readers may find a brief discussion of the state of research on the question of whether simulating wrong doing negatively impacts character and behaviour helpful before I set it to one side.

An aside: does simulating wrongdoing increase non-simulated wrongdoing?

Much empirical research has been undertaken into the effect that violent videogame exposure has on character and behaviour (specifically on aggression and violence). A recent and comprehensive meta-analysis of violent videogames research undertaken by the American Psychological Association found existing data to “demonstrate clear and consistent adverse effects of violent videogame exposure on the social and cognitive behaviors of children, youth, and adults.” At the same time, they found there to be a lack of “studies that examine the link between violent videogame exposure and extreme forms of aggression, that is, violence” (Calvert et al. 2017, p. 138). However, violent videogame exposure is, at any rate, only an imperfect analogue for simulated wrongdoing since not all violent behaviour is morally wrong, not all wrongdoing is violent, and not all violent videogame content is violence that one does.

There are highly significant normative questions about (a) what we ought to do in the absence of sufficient studies to examine potential links between videogame use and violence, and (b) how severe or common negative effects of violent videogame play must be to make their production or consumption morally objectionable, or to warrant more stringent regulation than exists. However these questions are outside the scope of the present paper. Indeed, as indicated already, this paper contends that for a player to direct their

character to perpetrate an immoral action type within a videogame is, *as such*, a *pro tanto* wrong quite apart from any effects doing so might have on their character or subsequent behaviours. The claim would hold true even if one were in no position to hurt anyone, perhaps because one were the last person on earth, or were on a desert island with no hope of being discovered, or simply had such excellent self-control that one would never translate one's simulations to real world actions. Even in such situations, there would still be moral reason not to simulate wrongdoing, and this reason would exist even putting aside concerns about what manner of character we might thereby cultivate through simulating wrongdoing. Similarly, Patridge has also argued that "that there are non-consequentialist resources for morally evaluating our single player game play."¹⁰ For her too, we have "legitimate intrinsic moral reasons to avoid playing [some] videogames".¹¹ On her view, "some videogames contain details that anyone who has a proper understanding of and is properly sensitive to features of a shared moral reality will see as having an incorrigible social meaning that targets groups of individuals, e.g., women and minorities."¹² Her argument focused on social meanings of videogames that we have reason to reject. For example, racist and sexist meanings. The argument of this paper will be wider in scope, and also rather more controversial.

Is it distinctively wrong to simulate doing wrong?

I will argue that for a player to direct their character to perpetrate an immoral action type within a videogame is, *as such*, a *pro tanto* wrong, quite apart from the player's motives and purposes, and quite apart from any effects doing so might have on their character or subsequent behaviours. The same is true of simulating wrongdoing more generally. However, the same is not true of consuming representations of wrongdoing more generally. I will frame my discussion around Morgan Luck's 'Gamer's Dilemma'. 'Is it immoral', asks Luck 'for a player to direct his character to murder another within a videogame?'

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.¹³

Whether or not most people do in fact feel that virtual paedophilia is not morally permissible is an empirical question immaterial to my own line of argument. The moral methodology that I employ is that of Reflective Equilibrium, which I will explain in detail in the next section. Briefly, it involves bringing to light *one's own* judgements, and seeking a general principle which explains the most and revises the fewest of these with maximal economy; this is (ideally) supplemented with an error theory for any judgements revised in the process. It does not involve bringing anyone else's judgements into equilibrium: for then one would have no allegiance to the set of judgements that emerged. To return to the Gamer's Dilemma, I will follow Lucks' more recent preference (in conference presentations and draft manuscripts) for posing the Gamer's Dilemma as an inconsistent triad of propositions, all of which seem correct to (him and) me:

1. Virtual murder is permissible.
2. There is no relevant difference between virtual murder and virtual child molestation.
3. Virtual child molestation is impermissible.¹⁴

Most of the responses to the gamer's dilemma have focused on denying the second claim. For instance, Bartel has argued that proposition two is false because, whereas virtual murder does not constitute murder, virtual paedophilia does constitute child pornography, and child pornography is morally impermissible.¹⁵ However, proposition two seems to me to be correct. While virtual paedophilia does constitute child pornography, it constitutes a virtual sort of child pornography and the question arises as to whether *that* sort of pornography is morally objectionable if, *ex hypothesi*, no children are harmed in its making in just the same way that nobody is harmed in the making of virtual murder. The

¹³ Luck (2009) p. 31.

¹⁴ Luck (unpublished) 'The grave resolution to the gamer's dilemma'.

¹⁵ Bartel (2012), Young (2016) has responded by arguing that child pornography is synonymous with child abuse, that virtual paedophilia need not involve child abuse, and that therefore virtual paedophilia is not an example of child pornography. However, there seems to be little to motivate the claim that child pornography is synonymous with child abuse. Firstly, there is much child abuse which is not child pornography (e.g. neglect). Secondly, child pornography arguably consists in sexually explicit visual representations of children, and such representations could be generated from imagination without requiring any abusive treatment of children at all. So this response seems insufficient.

¹⁰ Patridge (2011), p. 303.

¹¹ Patridge (2011), p. 304.

¹² Patridge (2011), p. 303.

question of what harmless wrong might be perpetrated in the case of virtual paedophilia but not in the case of virtual murder then simply re-emerges for Bartel's response.

In this paper, I evaluate the rival possibilities of claiming that (a) virtual murder and virtual child molestation are both permissible,¹⁶ and of claiming that (b) virtual murder and virtual child molestation are both impermissible. The question of which of these pairings I have most reason to accept will turn on which of them survives the process of reflective equilibrium, a process that I will describe in the next section. It will already come as no surprise that the pair of claims which I contend best survives the process is that which regards both virtual murder and virtual child molestation as impermissible. Or, more accurately, the pair which claims each to be, as such, *pro tanto* morally wrong, rather than always *all things considered* wrong.

Moral methodology: reflective equilibrium

Following Rawls and others, I contend that we should reason ethically by first collecting those of our ethical 'judgements in which our moral capacities are most likely to be displayed without distortion'.¹⁷ These judgements may be of any level of specificity or generality.¹⁸ At some level this first step in reflective equilibrium may seem question begging, since it requires an account of the conditions under which our moral capacities will generate the correct judgements. In doing so it presupposes a grasp of the very things that the procedure of reflective equilibrium is supposed to generate, namely, reliable ethical judgements. However, some conditions seem very likely to allow our moral capacities to function without distortion:

Since the rightness or wrongness of an action, or the justice or injustice of an institution, depends on facts about it, judgments made in ignorance of these facts are unreliable. Factors such as self-interest and emotional distress also make a judgment unreliable, because these factors can interfere with a person's assessment of the morally relevant considerations.¹⁹

Next we should further restrict the set of judgements we consider to those to which we are most firmly committed. This is because by bringing into equilibrium views to which

we are uncommitted, we would then not be committed to the picture which emerged.

With this subset of our ethical judgements in hand, we ought then to attempt to see which principles can be formulated to explain the most and revise fewest of them. Preference would be given to the simplest, least extensive, most explanatorily powerful and coherent set of principles and judgements. 'A moral theory has *power* when it yields judgments not included in the original data base.'²⁰ 'A moral theory has simplicity if it yields a body of judgments out of a relatively sparse amount of theory, deriving the numerous complex variations of the phenomena from a smaller number of basic principles.'²¹ There is more to coherence than mutual consistency: 'We want the principles to hang together, to be mutually supportive, to be jointly illuminated by the moral concepts to which we appeal.'²²

Where one has judgements which are left by the wayside, one would ideally have a theory of error of how one came to have those mistaken judgements. Such a theory of error would point out some flaw in how those mistaken judgements were generated which differs from how our correct judgements were generated, and in virtue of which they are not to be trusted, while our correct judgements are. After all, it was not so long ago that they seemed to be on the same footing. In this way, a theory is generated from judgements.

This, however, is only 'narrow' reflective equilibrium. As Spiecker and Steutel argue, following Rawls, our judgements are to be brought into 'wide' reflective equilibrium, where they are rendered coherent with our other beliefs and theories which 'have some inferential bearing on our considered judgements and ethical principles'²³; 'justifying an ethical belief does not exclusively consist [sic] of showing it to cohere well with other ethical beliefs, but also essentially appealing to non-ethical views, especially to psychological, political and economic theories'.²⁴ Take the following example from Stephen Law: 'The dispute over whether or not women should be allowed to vote was, in part, a dispute over whether women have the necessary intellectual skills to exercise that right properly'.²⁵ Spiecker and Steutel add the sensible requirement of a feasibility test: this is 'to make sure that living according to the ethical views we are defending is possible for creatures like us'.²⁶

What we have in reflective equilibrium is 'a coherence account of justification (as contrasted with an account of

¹⁶ To motivate the first thought, one might urge that neither simulating murder nor simulating paedophilia is immoral, but that both are simply in bad taste, for instance.

¹⁷ Rawls (1972), p. 47.

¹⁸ Rawls (1999), p. 289.

¹⁹ Scanlon (2002), p. 145.

²⁰ Kagan (1989), p. 12.

²¹ Ibid.

²² Ibid.

²³ Spiecker and Steutel (2001), p. 35.

²⁴ Ibid.

²⁵ Law (2006), p. 118.

²⁶ Spiecker and Steutel (2001), p. 36.

truth)'.²⁷ For Spiecker and Steutel: 'The more comprehensive and coherent an ethical conception is, and the better it takes into account our initially tenable beliefs, the stronger is the justification of its constituent parts'.²⁸ Coherence among claims 'is not just a function of logical consistency but also a matter of connectedness by inferential relations' so that they constitute 'a network of mutually supporting claims,' rather than a motley assortment of 'mutually indifferent judgements',²⁹ what Kagan calls a 'laundry list'³⁰; 'Other things being equal, a more comprehensive coherent set of beliefs is more credible than a less comprehensive one'.³¹ 'The credibility of our ethical view on human interaction would be increased if it could be integrated with those other [coherent sets of] views [on different topics] into a more comprehensive coherent belief set'.³²

Initial judgements

As per the process of reflective equilibrium, we should reason ethically by first collecting those of our judgements (of all levels of generality and specificity), in which our moral capacities are most likely to be displayed without distortion and to which we are most firmly committed. The judgements that I begin with are these:

1. Simulating murder is not *pro tanto* morally wrong.
2. Simulating paedophilia is *pro tanto* morally wrong.
3. There is no morally salient difference between simulating murder and simulating paedophilia.

At least one of these must be mistaken. In line with the methodology spelt out in the previous section, to determine which of these judgements is mistaken we need more relevant judgements, a principle that would explain the most and revise the fewest of these judgements (and which would include the revision of proposition one or two), and (ideally) a plausible error theory to account for the mistaken judgement.

First then, let us supplement the above list with some further, intuitive judgements:

4. Simulating torture (and otherwise cruel and unusual treatment) is *pro tanto* morally wrong.

5. We have duties to refrain from actions which others have a significant interest in us not performing, from which refraining comes at no significant cost to ourselves.
6. Actions can significantly impact us negatively without constituting an unpleasant or otherwise unwelcome experience. To deny this means, implausibly, that:

even if a man is betrayed by his friends, ridiculed behind his back, and despised by people who treat him politely to his face, none of it can be counted as a misfortune for him so long as he does not suffer as a result. It means that a man is not injured if his wishes are ignored by the executor of his will, or if, after his death, the belief becomes current that all the literary works on which his fame rests were really written by his brother.³³

Before exploring unifying, vindictory principles for the greatest number of these judgements, let us consider error theories that explain why it is that our judgements about whether virtual paedophilia or virtual murder is permissible, might be erroneous.

Error theories

Three theories that explain why it is that our judgements about the permissibility of virtual paedophilia and virtual murder might be erroneous suggest themselves; call these *The Taboo Thesis*, *The Inoculation Thesis* and *The Squeamishness Thesis*.

The Taboo Thesis: Notwithstanding the views of the Ancient Greeks, we in the West have been conditioned to react negatively to content related to paedophilia even beyond content which would be morally impermissible. While it is possible to talk about paedophilia in the rather distanced way that I currently am, it is decidedly more unpleasant for both speaker and listener in general than it is to discuss murder.

The Inoculation Thesis: We have been surrounded by simulated violence since a young age, and so it seems perfectly ordinary to us, and it is hard to recognize what is perfectly ordinary as morally objectionable. We have not been surrounded by simulated paedophilia since a young age, and so it does not seem perfectly ordinary to us.

The Squeamishness Thesis: We frequently find it hard to keep simulation and reality separate since we are evolutionarily hardwired to confound them. By way of comparison, consider the 'Rubber Hand Illusion'

²⁷ Daniels (2016).

²⁸ Spiecker and Steutel (2001), p. 33.

²⁹ Ibid, p. 32.

³⁰ Kagan (1989), p. 11.

³¹ Spiecker and Steutel op. cit, p. 33.

³² Ibid.

³³ Nagel (1970), p. 76. To be sure, Nagel argues that these are indeed unacceptable results.

in which a subject knowingly places their hand behind their back, out of harm's way, and has a prosthetic hand placed in position consistent with where their own hand would naturally sit. The rubber hand (on some variations) is then threatened, and the subject feels a phantom attachment to this prosthetic limb, and reacts as if it were their real hand.

These error theories are empirical conjectures that imply certain predictions that would confirm or disconfirm them. The Inoculation Thesis (by itself) predicts that had we not been inoculated against simulated murder, we would think that both virtual paedophilia and virtual murder were impermissible. It also predicts that had we also been inoculated against virtual paedophilia we would think that both virtual paedophilia and virtual murder were permissible. The Squeamishness Thesis together with The Inoculation Thesis predicts that in the absence of inoculation against either, we would squeamishly object to both. The Inoculation Thesis derives some plausibility from the judgement that, *prima facie*, simulating torture (and otherwise cruel and unusual treatment) is *pro tanto* morally impermissible. The simulation of torture (and otherwise cruel and unusual treatment) seems perhaps even more obviously continuous with the simulation of murder than the simulation of paedophilia does, and is equally something that most of us have not been inoculated against.

Drawing on these resources, one line of thought contends that the judgement that virtual paedophilia is impermissible is an error and that The Squeamishness Thesis together with The Inoculation Thesis, or else The Taboo Thesis all by itself explains this error. After all, its proponent will continue, nobody is harmed or has claims over us that do not involve the maintenance of their interests, and it is hard to see what interest anyone else has in whether or not we should indulge in these activities (I will return to this point later, arguing that we do have such interests). Another line of thought contends that the judgement that virtual murder is permissible is an error and that The Inoculation Thesis explains this error. The proponent of this line of thought will contend that simulating either action type is morally impermissible (or at least as such, *pro tanto* wrong).

It should be acknowledged that even if we accept an error theory for one of our intuitive judgements, we may still continue to feel its pull. In such cases we should recognize their pull as 'a sort of ineliminable moral illusion, similar to certain optical illusions which do not lose their intuitive hold on us even when our theory tells us better.'³⁴ This may seem to put some strain on one of the error theories. If we want to explain away our judgement that virtual paedophilia is

wrong by appealing to The Squeamishness Thesis, we then need to explain why we do not feel similarly about virtual murder via the same mechanism. In that case one needs to contend, entirely plausibly, that inoculation overwrites that intuition.

It seems at this point that we have reached an impasse, since plausible error theories are available to all sides: we cannot work from these error theories alone to resolution to the gamer's dilemma. Our next objective is to consider what principles might best explain our considered, most reliably formed moral judgements (bearing in mind that where we left the discussion in this section, either of our intuitions could plausibly be explained away). What would the wrong making feature be, and what other actions would it also make wrong?

Against a virtue ethics frame

Matthew McCormick develops the thesis that simulating wrongdoing amounts to self-harm, and that a prohibition on self-harm is the relevant wrong-making feature for both virtual paedophilia and virtual murder. McCormick takes an Aristotelian lens for discussing the wrong of simulated killing. Raising the example of harmless (to others) virtual paedophilia, McCormick seeks to secure our intuition that there is 'something wrong with the activity without regard to what might happen' beyond the simulation.³⁵ He thinks this can be explained by Aristotle's account of virtue ethics according to which 'a deep, fulfilled happiness or flourishing (*eudaimonia*),' which we each have most reason to pursue, 'can only be achieved by pursuing the development of the capacities that are the unique function of human beings.'³⁶ McCormick worries that 'By participating in simulations of excessive, indulgent, and wrongful acts, we are cultivating the wrong sort of character.'³⁷ He concludes that 'By engaging in such activities, you do harm to yourself in that you erode your virtue, and you distance yourself from your goal of *eudaimonia*'; 'the harm that may occur is best construed as harm to one's character.'³⁸

One may recall a standard objection that it is not clear that one's flourishing is indeed tightly connected with one's virtue, and that the two can and often do come apart. However, if we do allow a tight connection, it is important to ask why it matters *morally*. Certainly we might worry that in cultivating a bad character we will treat others significantly worse, but McCormick argues in the same article that we have no

³⁴ Kagan op. cit, p. 15.

³⁵ McCormick op. cit, p. 284.

³⁶ McCormick op. cit, p. 285.

³⁷ Ibid, p. 284.

³⁸ Ibid, p. 285, 286.

reason to think that this is true. Furthermore, it certainly would not tell us that there is ‘something wrong with the activity without regard to what might happen’ beyond the simulation. While we might admit that we do indeed have self-regarding duties, such duties do not demand our cultivating our own *perfect* wellbeing. If all we do in simulating wrongdoing is forego our *perfection*, then it is hard to see how we have violated self-regarding duties. Instead, the case would need to be made that simulating serious wrongdoing does such serious damage to us that we owe it to ourselves not to participate in such actions. Here we need an account of damage that does not beg the question. For if the damage just consists in an increased disposition to engage in similar activity in the future, no case has been made at all.

Indeed, it is not clear that the virtues (such as kindness, say) should impede virtual wrongs in a virtual reality. Following John McDowell, let us grant that, “a kind person has a reliable sensitivity to a certain sort of requirement which situations impose on behaviour,”³⁹ that the kind person is responsive to the requirements of the situation and that analogous reliable sensitivity and responsiveness constraints typify each of the virtues. A virtuous person, of course, would not molest a child; the requirements of almost every situation are such that child molestation is ruled out, the virtuous person will be reliably sensitive to these requirements and will act on them. What situational features prohibit the virtual molestation of virtual children? Unless this question can be convincingly answered, one could contend that real virtue is not incompatible with virtual paedophilia. I will develop an answer to this in the next section, but my concern will be primarily with what we owe to others rather than to ourselves.

Perhaps most fundamentally for the way McCormick develops the Aristotelean case, we should ask which capacity that is the unique function of human beings is incompatible with simulating serious wrongdoing such as this. In the absence of a plausible answer to this question, the McCormick’s attempt to make the Aristotelian case against virtual paedophilia simply cannot get off the ground. I do not mean to say that an Aristotelian case cannot be made, but merely that we have much further to go in it being made. Without pretending to have refuted it, I want to suggest that even if it is correct, it cannot be the whole of the story of what is wrong about simulating wrongdoing, or the most fundamental part of it. I want to motivate this view with the imaginary case of The Devious Super Geek.

The Devious Super Geek: Imagine that a student in one of your classes from school was a brilliant computer programmer and built a program featuring detailed vir-

tual models of you and your other classmates. Imagine in that program he proceeded to simulate doing serious wrongs to you and your classmates.

It seems to me that in this case you should be angry with The Devious Super Geek for what he has done, and it seems that your ground of anger would not be that he has harmed himself, but that he has wronged you. He would have wronged you even if he were never to be found out. It would be a wrong he had done to you and your classmates without its being positively unpleasant for any of you. The wrong, it seems would consist in demonstrating disrespect for you.⁴⁰

In favour of a deontological, respect frame

At the end of the last section, I invoked the hypothetical case of The Devious Super Geek to motivate the view that we can do wrong to particular people by simulating doing wrong to those particular people. A general principle that would explain that is the following: it is wrong to simulate doing wrong to particular people. An explanation of why it should be wrong to simulate doing wrong to particular people is that it constitutes a disrespectful act towards those particular people to whom wrong is simulated. This prohibition on disrespectful actions goes beyond causing suffering or otherwise unwelcome experiences to explaining why it is wrong to ridicule people behind their back, or to insult and humiliate someone with an intellectual disability even though they cannot know that they are being ridiculed, or to urinate on a grave at night with nobody around to see and no marks left. The idea that I have begun to develop with the case of The Devious Super Geek is that to simulate harm to the occupants of morally salient categories is disrespectful to the occupants of those morally salient categories. Now I suggest that we can move from a judgement that it is wrong to simulate wrongdoing to particular people to this more general conclusion by some bridging cases.

The Devious Super Geek was going to make a simulation of you and your classmates. But we can make the example less intimate in the following ways.

The Paired Devious Super Geeks: Imagine that a student in one of your classes from school was a brilliant computer programmer and built a program featuring detailed virtual models of you and your other classmates. Imagine that he then teamed up with a friend who was similarly talented, who made a program featuring detailed virtual models of people that he knew.

³⁹ McDowell (1979).

⁴⁰ For sophisticated discussion of non-consequential wrongs still within an Virtue Ethic Frame, see Partridge (2011).

Imagine that the two simply swapped the videogames to play in the way they were going to play with the videogame featuring simulants of the people in their own class. Imagine that within these programs each proceeded to simulate doing serious wrongs to the other's classmates.

The Impersonal Devious Super Geek: Imagine that a brilliant computer programmer built a program featuring detailed virtual models of people that he had never met. Imagine in that computer program he proceeded to simulate doing serious wrongs to these people that he had never met.

These cases would also seem morally objectionable, and so the fact of knowing the specific people in question can be ruled out as the moral difference maker. Perhaps the difference maker is the fact that the people in the examples are actual individuals. To test this conjecture, we can adjust the example to make it less specific, supposing instead that the simulants were not simulations of particular people.

The Non-specific Devious Super Geek: Imagine that a brilliant computer programmer built a program featuring detailed virtual models of imaginary members of real groups of people, perhaps the homeless, immigrants, prostitutes, children, women, or Jewish people.⁴¹ Imagine in that program he proceeded to simulate doing serious wrongs to these people.

Here it still seems morally wrong to simulate wrongdoing to these simulants qua simulations of members of these categories. However, simulating wrongdoing to groups that have no history of oppression, such as white adult males, would not vindicate its moral status, even if it would mitigate it. What seems to be wrong here is that by simulating harm to the occupants of morally salient categories we thereby disrespect the moral value of actual occupants of those categories. Equally, humanity in general is a morally salient category, to which disrespect is evinced (or a failure of respect is in evidence) in enacting simulations of crimes to imaginary members of that category. Furthermore, humanity is not the most general morally salient category. By focusing on just the category of humanity we leave out non-human animals that could be morally salient in respect of their being persons, or in respect of their being sentient. Indeed, some argue that it is in respect of a capacity to suffer that human beings are morally salient, and not in virtue of being human at all.⁴²

By these bridging cases then, it seems that we wrong actual members of morally salient categories when we simulate doing wrong to members, and possible members, of those categories. However, given that both aliens and zombies seem to constitute possible members of morally salient categories (the categories of persons and or of sentient creatures) it would follow that we should refrain from playing videogames in which we simulate doing wrong to aliens or zombies. One might now contend that in such games, it is rare that the object is to simulate doing wrong to such creatures, but I can now extend my case. I have spoken throughout about the ethics of simulating wrongdoing, rather than e.g. the ethics of playing violent videogames (as McCormick, Schulzke and Waddington have debated).⁴³ Now one might think of games such as those in the Medal of Honour game series in which:

- (1) One cannot simulate killing one's allies.
- (2) One cannot simulate killing civilians.
- (3) One is enabled to simulate killing Nazi soldiers in the context of World War II.

I take it to be fairly uncontroversial that the Allies' cause in World War II was just (even if it was not always fought in just ways). Very arguably, playing Medal of Honour constitutes an example of simulated *rightdoing* rather than simulated *wrongdoing*. The arguments that I have developed above would then seem to have no power against such games. However, the following reply can be made. Grant that the Nazis, and German soldiers more generally, had some moral status qua their humanity, rationality or sentience, and that it was a *pro tanto* wrong to shoot and kill them. Grant that the *pro tanto* wrong of killing Nazis, and German soldiers more generally, was outweighed by considerations of saving occupied Europe from the terror and inhumanity that they wrought. Now consider by comparison the act of simulating the killing of Nazis and German soldiers. If I am right, then it is a *pro tanto* moral wrong to simulate killing them. However, nothing constitutes the conditions of justification here in the way that the occupation of Europe and existence of, for instance extermination camps, did. It would seem then, that one ought not to simulate even justified killing.

Why is it that wrongness of wrongdoing 'comes through' to the simulation, but the rightness of right doing does not come through? First, in simulating justified killing, one actually manufactures the conditions of justified killing. For instance, consider a case of one's organizing the indoctrination of a group of people such that they will go about perpetrating violent wrongs, unless one violently prevents them. It seems that, overall, such actions count as wrong. One can

⁴¹ Cf Patridge (2011).

⁴² "[T]he question is not, can they reason? Nor, can they talk? But, can they suffer?" Bentham (1970).

⁴³ McCormick (2001), Waddington (2007), Schulzke (2010).

avoid playing the games altogether without moral cost, but one cannot avoid the justified killing of World War II without moral cost, i.e. the moral cost of allowing the Nazis to expand their oppressive and murderous regime without opposition. Second, we don't have as much of an interest in having rightdoing simulated to us as we do in not having wrongdoing simulated to us. While it is distinctly *against my interests* that you should simulate murdering me, it is not as distinctly *in my interest* to have you simulate giving to me charitably. This is borne out through considering whether it is preferable to have *some* people simulate charitable giving to you, or whether it is preferable to have *nobody* simulate murdering you, while the former might be nice, it is trumped by the nastiness of the latter. The surprising conclusions that we are led to include not just that simulated wrongdoing is, as such, *pro tanto* wrong, but that simulated justified killing is, as such, *pro tanto* wrong, and that the simulated killing of zombies and aliens is also, as such, *pro tanto* wrong (assuming, plausibly, that these will belong to the morally salient categories of either agents, or sentient entities).

Two objections

In this section, I propose to answer two potential objections. The first is that Devious Super Geek argument generalizes to generate a *pro tanto* objection to consuming all representations of wrongdoing, rather than just simulating wrongdoing. The second is that what makes The Devious Super Geek's acts disrespectful is that his victims do not know what he has been up to, but that quite generally we human beings do know that acts of wrongdoing are simulated in games. Let us consider each objection in turn.⁴⁴

The first objection contends that the above arguments apply to consuming representations of wrongdoing more generally, and that this is too counterintuitive a conclusion to be acceptable. Suppose that the Devious Super Geek was an author, and, instead of creating a virtual world in which players could simulate wrongdoing, he simply wrote descriptions of similar wrongdoing. While the production of this fiction seems objectionable, as I said at the outset, the production of representations of wrongdoing is not within the scope of this paper, and must be set aside for another occasion. Suppose now that someone read these descriptions; i.e. consume them—something decidedly within the scope of the present paper. Would the mere reading of these descriptions be similarly disrespectful to playing a game in which one simulates the acts described? If so, *mutatis mutandis*, then the disrespectfulness of reading such descriptions would generalize

to include reading descriptions of wrongdoing *tout court*. However, it seems to me that reading such descriptions, as such, would not be similarly disrespectful. Rather, any disrespectfulness evident in reading such descriptions would have to consist not in the fact of reading, but in the attitude that one has to the wrong represented. To illustrate the distinction, consider the intuitive moral difference between watching Spielberg's, *Schindler's List* and playing a videogame based on that film from the point of view of the Nazis, with the same objectives as the Nazis. While one certainly could watch *Schindler's List* in a disrespectful way, perhaps cheering through murders, or jeering at lucky escapes, no such inappropriate emotions need be had for there to be a *pro tanto* wrong in the act of playing the hypothetical videogame, one which would require significant countervailing considerations in order to be outweighed.

The second objection contends that what makes The Devious Super Geek's acts disrespectful is that his victims do not know what he has been up to. However, it is not that his victims *do not know* that they have been replicated and had wrongs simulated to them that makes those simulated acts disrespectful. The point of invoking Nagel's considerations is rather to show that even if they did not know, they would no less have been wronged. Knowing about the wrong does not mitigate its moral status, any more than knowing that one has been insulted mitigates the moral status of the insult; often knowing about the wrong exacerbates it by constituting an unpleasant experience. More arguably, one might suggest that A Less Devious Super Geek could simulate doing wrong without thereby doing wrong if he sought permission from those he was going to build virtual versions of and sought permission for the sorts of things he planned to do to simulate doing to them. On this understanding, simulating wrongdoing has a defeasible, default quality of disrespectfulness that disappears as soon as consent is given for it. Supposing that such permission was granted, one might then contend that simulating wrongdoing to these parties would not itself be wrong. However, firstly it is not generally the case that people who simulate wrongdoing have sought permission to simulate wrongdoing. Secondly, it is likely impossible for all those potentially wronged in the simulation of wrongdoing to consent to it. This is both because the class of those potentially wronged is likely to be very large (e.g. humanity) and its occupants are likely not all able to consent (e.g. sentient beings). More fundamentally however, it seems that gaining consent for treating someone in a disrespectful way does not remove the *pro tanto* wrong of treating someone in that way, it simply removes a further *pro tanto* ethical barrier, namely that wrong of making decisions about a person without involving that person in the decision. I will take an example from personal experience to illustrate this point. One student at a school where I trained to teach was perfectly happy to kowtow to a more powerful

⁴⁴ Both of these objections have been suggested by an anonymous reviewer.

student and be belittled by that person, even in public spaces. She could not explain why she was happy to put up with this treatment, although we could speculate. Her mistreatment included having her hair used to clear up food from a table in a busy school canteen. The girl consented to this treatment. The consent made the case a somewhat awkward for teachers to deal with, but intervene they did on the grounds that she should not have let herself be debased in this way, and should have had more self-respect. One might additionally contend that unless the devious super geek was intending to disrespect the people to whom he simulated wrongdoing, then he had not really disrespected them at all. However, disrespect does not require an intention to disrespect, but a failure to show respect or simply an intention to do the thing which is disrespectful quite apart from whether they regard it as disrespectful: such as pushing in front of someone in a queue. One can show a kind of respect to others by not simulating wrongdoing to them, and by refusing to simulate wrongdoing to them.

All things considered wrongs

Briefly, I want to evaluate whether some considerations might trump the *pro tanto* wrong of simulating wrongdoing. Some have suggested, in line psychoanalytical theory, that *simulating* wrongdoing could have a cathartic effect and thereby reduce the chances of one *actually* doing wrong. Such a cathartic effect might, if strong enough, plausibly give one *all things considered* reasons to simulate wrongdoing. However, Calvert et al. empirical grounds for doubting that this is the case:

Although some areas did not have sufficient studies to test effects, there were no outcomes supporting a reduction in any kind of aggressive outcome, as predicted in psychoanalytic theory through catharsis.⁴⁵

It can be urged that by simulating wrongdoing, we can learn moral lessons and that this can give us reason to so do. This could render it *all things considered* permissible to simulate wrongdoing. One possibly morally educative example of simulated wrongdoing that has been suggested to me, is the case of an imaginary videogame based on the experiences of the Sonderkommando. In such a videogame the player would be tasked with simulating the morally ambiguous or excusable actions undertaken by the Sonderkommando.⁴⁶ The Sonderkommando were work units made up of Nazi death camp prisoners who, on threat of death, were

tasked to aid with the disposal of gas chamber victims during the Holocaust. However, a game such as this, in which one simulates morally ambiguous and excusable acts, does not straightforwardly involve simulating wrongdoing, and so does not straightforwardly constitute an *all things considered* example where simulating wrongdoing is not itself wrong. The lesson one might learn from playing such a game is just how difficult decision making can become, and how unclear it can be as to what is permissible or merely excusable.

The educative value of playing the real videogame, *Spec Ops: The Line* (2012) might, *all things considered*, also be suggested to trump the *pro tanto* wrong of simulating wrongdoing. That videogame begins by tasking the player with simulating morally ambiguous or excusable actions undertaken by US Special Forces in a hypothetical scenario. However, the game begins to task its players with simulating straightforwardly morally impermissible actions. Somewhat like Michael Haneke's *Funny Games*⁴⁷ which is a film he did not want his audience to watch, *Spec Ops: The Line* seems to be a game that the game makers do not want its audience to play. Walt Williams, the writer of *Spec Ops: The Line*, commented on the game in an interview:

[T]he ultimate real choice of any videogame is not the choice that we've given you in the game. It's the choice of, "Do I want to play a game where I do these things, or do I not like to play that?" Turning off the game is a valid player choice ... it's about looking at what you're comfortable with doing and realizing that you're simulating truly terrible acts. Even though they are simulated, even though they are not in the world that we are in, you are still choosing to do them over and over.⁴⁸

Williams seems to allow that it is acceptable to make such a game, but not to play such a game and so again not to constitute an *all things considered* example in which simulating wrongdoing is not itself wrong. Williams' phrasing is of course more permissive, "turning off the game is a valid player choice." However, it is much more plausible to interpret Williams as giving his player every reason to turn off, and little reason to carry on playing, so that turning off is more than merely valid. About half way through, the game presents players with the choice between firing white phosphorous at troops from a mortar in order to continue playing, and not playing at all. The game finishes with a discussion in which the player character, Walker, is asked by Konrad, 'do you feel like a hero yet?' and confronted with his (the player's) orders having killed '47 innocent people.' Walker says that he is 'done playing games', Konrad replies 'I assure

⁴⁵ Calvert et al. (2017), p. 138.

⁴⁶ I am grateful to Adam Slavny for suggesting this example in conversation.

⁴⁷ Haneke (1997).

⁴⁸ Grayson (2013).

you, this is no game'. Walker further argues that 'What happened here was out of my control', to which Konrad answers, 'Was it? None of this would've happened if you just stopped. But on you marched. And for what?' All of this indicates that, rather than being merely a 'valid choice', putting down the controller is the reaction the game maker thinks most appropriate. In this connection, compare Haneke's comment regarding his 1997 film, *Funny Games* 'I turn the viewer into the killer's accomplice and in the end, I chastise the viewer for playing that role.'⁴⁹ Indeed, more generally it is difficult to imagine what one might learn from simulating straightforwardly immoral acts, for one already knows that they are immoral. There does not seem to be any morally educative function, for instance, to simulating wrongdoings of the kind that saturate games in the Grand Theft Auto series.

My thesis has been only that simulating wrongdoing is *pro tanto* wrong, and so there could perhaps be occasions on which that wrong is outweighed by other considerations. However, it is not so obvious what those occasions are in the normal course of game playing. It might be urged that the *pro tanto* wrong can be easily outweighed by other positives such as the pleasure that people derive from simulating wrongdoing. Taking pleasure in such acts seems to exacerbate rather than vindicate their morally problematic status. The pleasure that a bully and his consorts derive from teasing a victim with an intellectual disability, who does not know he is being abused, in no way vindicates or even mitigates the moral status of the act, but only exacerbates it. How can the *pro tanto* wrong be outweighed? I think very clearly the interests of national and international security warrant war games and tactical training which involves the simulation of violent right doing and also the simulation of a wide range of wrongdoing. Police and soldiers sadly often have reason to take the lives, due to lethally aggressive actions, and need to be ready to end such aggressions quickly and with minimal collateral damage, which often means taking the lives of the aggressors. To be ready to do this requires training in circumstances which mirror likely scenarios as closely as possible. Any wrongs of disrespect are easily outweighed by considerations of domestic and national defence. Col. Matthew Caffrey, professor of war gaming and planning at the Air Command and Staff College, observes that 'Combat veterans live longer' adding that 'one reason we use war games is to make virtual vets'.⁵⁰

Conclusion

I have argued that it is *pro tanto* wrong for a player to direct his character to perpetrate immoral action types within videogames (as such), and that it is more generally *pro tanto* wrong for an individual to simulate perpetrating immoral action types (as such). I have argued further that simulating wrongdoing is a *pro tanto* wrong whose wrongness does not tarnish other cases of consuming representations of wrongdoing. What matters morally about the consumption of representations of wrongdoing (as such) may only be the way in which one consumes them (and what effects they have on the way we are disposed to treat others). However, to simulate wrongdoing to morally salient individuals and to actual and possible occupants of morally salient categories is disrespectful of those individuals and to actual occupants of those morally salient categories. The occasions on which we are apt to not share this judgement intuitively can, at least plausibly, be explained away by our culture having inoculated us against those judgements. I aimed to motivate this view by starting with the case of The Devious Super Geek simulating wrong to particular people that he knew personally, and then building bridging cases to capture games in which one simulates wrong to imaginary members of extant morally salient categories. Finally, I described how I proposed to handle some potential objections, and discussed whether the *pro tanto* wrong of simulating wrongdoing was trumped in some *all things considered* cases. The surprising conclusions that we were led to included not just that simulated wrongdoing is *pro tanto* wrong, but that simulated just killing is *pro tanto* wrong, and also that the simulated killing of zombies and aliens is also *pro tanto* wrong. Furthermore, these *pro tanto* wrongs are not easily outweighed in *all things considered* deliberations.

Acknowledgements Viktor Ivanković, Morgan Luck, Chris Mills, David Rowthorn, Adam Slavny, Nathan Wildman and two anonymous referees are to thank for reading and commenting helpfully on drafts. Craig Bourne, Emily Caddick Bourne, Rebecca Davnall, Michael Hand, Hwa Kim, and Felix Pinkert are to thank for very interesting conversations on its arguments. Christopher J Bartel and Derek Matravers are to thank for organising the Video Games and Virtual Ethics conference in July 2017 at which I delivered the presentation on which this paper is based. Emma Fenn, Philp Gaydon and Sally Latham are also to thank for inviting me to speak to their A Level students about video game ethics.

References

- Anderson, C. A., Carnagey, N. L., & Eubanks, J. (2003). Exposure to violent media: The effects of songs with violent lyrics on aggressive thoughts and feelings. *Journal of Personality and Social Psychology*, *84*, 960–997.
- Bartel, C. (2012). Resolving the gamer's dilemma. *Ethics and Information Technology*, *14*, 11–16.

⁴⁹ Haneke (2005).

⁵⁰ Singer (2009).

- Bentham, J. (1970). *An introduction to the principles of morals and legislation*. London: The Athlone Press.
- Calvert, S. L., Appelbaum, M., Dodge, K. A., Graham, S., Nagayama Hall, G. C., Hamby, S., ... Hedges, L. V. (2017). The american psychological association task force assessment of violent videogames: Science in the service of public interest. *American Psychologist*, 72, 126–143.
- Cummings, J., & Bailenson, J. (2016). How immersive is enough? A meta-analysis of the effect of immersive technology on user presence. *Media Psychology*, 19(2), 272–309.
- Daniels, N. (2016). Reflective equilibrium. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*, Winter Ed., Retrieved from <https://plato.stanford.edu/archives/win2016/entries/reflective-equilibrium> Accessed 21 Jan 2018.
- Fox, J., Bailenson, J. N., & Ricciardi, T. (2012). Physiological responses to virtual selves and virtual others. *Journal of Cyber-Therapy & Rehabilitation*, 5(1), 69–72.
- George, T. (2005). *Hotel Rwanda*.
- Grayson, N. (2013). *Spec ops, fc3 writers on art, treating players intelligently*. Retrieved from <https://www.rockpapershotgun.com/2013/04/11/spec-ops-fc3-writers-on-art-treating-players-intelligently>. Accessed 21 Jan 2018.
- Griffith, D. W. (1915). *The birth of a nation*. David W. Griffith Corp.
- Haneke, M. (1997). *Funny games*. Filmfonds Wien.
- Haneke, M. (2005). “Hidden: Interview with Michael Haneke by Serge Toubiana. DVD Special Features. *Hidden (Caché)*. Dir. Michael Haneke. Madman. Retrieved from <https://www.youtube.com/watch?v=roOI9PvEPjs>. Accessed 21 Jan 2018.
- Kagan, S. (1989). *The limits of morality*. Oxford: Clarendon Press.
- Law, S. (2006). *The war for children's minds*. Abingdon: Routledge.
- 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.
- McCormick, M. (2001). Is it wrong to play violent videogames? *Ethics and Information Technology*, 3, 277–287.
- McDowell, J. (1979). Virtue and reason. *The Monist*, 62, 331–350.
- Nagel, T. (1970). Death. *Noûs*, 4, 73–80.
- Patridge, S. (2011). The incorrigible social meaning of videogame imagery. *Ethics & Information Technology*, 13, 303–312.
- Rawls, J. (1972). *A theory of justice*. Oxford: Clarendon Press.
- Rawls, J., Freeman, S. (1999). *Collected papers*. Cambridge: Harvard University Press.
- Riefenstahl, L. (1935). *Triumph des Willens*. Leni Riefenstahl-Produktion.
- Sanchez-Vives, M. V., & Slater, M. (2005). From presence to consciousness through virtual reality. *Nature Reviews: Neuroscience*, 6, 332–339.
- Scanlon, T. M. (2002). Rawls on justification. In S. Freeman (Ed.), *The Cambridge companion to rawls* (pp. 139–167). Cambridge: Cambridge University Press.
- Schulzke, M. (2010). Defending the morality of violent videogames. *Ethics and Information Technology*, 12, 113–127.
- Singer, P. W. (2009) *Wired for war: The robotics revolution and conflict in the 21st Century*. New York: Penguin Press.
- Spiecker, B., & Steutel, J. (2001). Reflective equilibrium as a method of philosophy of education. In F. Heyting, D. Lenzen & J. White (Eds.), *Methods in philosophy of education*. London: Routledge.
- Spielberg, S. (1993). *Schindler's List*. Universal Pictures.
- SuperData. (2017). Market brief—Global games 2017: The year to date. Retrieved from <https://www.superdataresearch.com/market-data/market-brief-year-in-review>. Accessed 21 Jan 2017.
- Waddington, D. I. (2007). Locating the wrongness in ultra-violent videogames. *Ethics and Information Technology*, 9, 121–128.
- Won, A., Bailenson, J., Lee, J., & Lanier, J. (2015). Homuncular flexibility in virtual reality. *Journal of Computer-Mediated Communication*, 20, 241–259.
- Young, G. (2016). *Resolving the gamer's dilemma: Examining the moral and psychological differences between virtual murder and virtual paedophilia*. London: Palgrave.