"May Kantians Commit Virtual Killings that Affect No Other Persons?"

Tobias Flattery
Wake Forest University

Abstract

Are acts of violence performed in virtual environments ever morally wrong, even when no other persons are affected? While some such acts surely reflect deficient moral character, I focus on the moral rightness or wrongness of acts. Typically it's thought that, on Kant's moral theory, an act of virtual violence is morally wrong (i.e., violate the Categorical Imperative) *only* if the act mistreats another person. But I argue that, on Kant's moral theory, some acts of virtual violence *can* be morally wrong, even when no other persons or their avatars are affected. First, I explain why many have thought that, in general on Kant's moral theory, virtual acts affecting no other persons or their avatars can't violate the Categorical Imperative. For there are real world acts that clearly do, but it seems that when we consider the same sorts of acts done alone in a virtual environment, they *don't* violate the Categorical Imperative, because no others persons were involved. But then, how could *any* virtual acts like these, that affect no other persons or their avatars, violate the Categorical Imperative? I then argue that there indeed can be such cases of morally wrong virtual acts—some due to an actor's having erroneous beliefs about morally relevant facts, and others due not to error, but to the actor's intention leaving out morally relevant facts while immersed in a virtual environment. I conclude by considering some implications of my arguments for both our present technological context as well as the future.

Introduction

Is it ever morally wrong to beat, stab, or shoot someone, so long as it's only a character in a video game, and so long as no real persons are affected? Many think the answer is 'no', precisely because no real persons are harmed or misused. But others still worry that something is morally amiss in at least some cases of virtual violence, especially cases of acts that, were they real instead of virtual, would surely be morally wrong, for instance virtual murders.[1] But what might be morally amiss in such cases of virtual violence? I'll call an act of *virtual violence* any act affecting a character in a

---

[1] For a summary of some of the broader moral and policy concerns with violent video games, as well as some of the relevant empirical work on the effects of playing these games, see American Psychological Association 2020.

virtual environment (I discuss virtual environments in the next section) in such a way that, if the environment and character were instead *real*, the character would be harmed.[2]

Perhaps some acts of virtual violence reflect deficient moral character in the actor. From a broadly virtue ethical perspective, for instance, acts of wanton virtual killing might evince vicious underlying moral character. Most philosophers writing on the ethics of virtual violence from this perspective, however, are less centrally concerned with evaluating the moral rightness or wrongness of virtual *acts*, and more concerned about how engaging regularly in acts of virtual violence might degrade one's moral character, hindering the development of virtue or even inculcating vice.[3]

It's trickier to argue that acts of virtual violence should be evaluated as morally wrong. From a consequentialist point of view, for instance, it's difficult to show that such acts bring about more bad results than good results, especially when the virtual characters one harms or kills aren't persons or even player-controlled avatars. Perhaps further empirical work will eventually support a consequentialist argument either for or against some forms of virtual violence.[4]

On Kant's moral theory, however, most authors writing on virtual violence argue or suspect that there isn't a good case to be made that acts of virtual violence could be morally wrong, if no persons or their avatars are affected.[5] Such acts might be morally *inadvisable* for various reasons (see §4), but not strictly speaking *morally wrong*—i.e., not such that they must violate Kant's fundamental moral principle, the Categorical Imperative. But I argue, to the contrary, that Kant's moral theory *can* explain how at least some acts of virtual violence are morally wrong, even ones that affect no other persons or their avatars.[6]

In the next section, I briefly discuss virtual environments and our interactions with them. In §3, I sketch the parts of Kantian moral theory that I'll assume. In §4, I explain why many have thought that on Kant's moral theory, virtual acts that affect no other persons or their avatars *can't* be morally wrong, strictly speaking. For there are real world acts that clearly violate the Categorical Imperative, but when we consider the same sorts of acts done alone in a virtual environment, they *don't* violate the Categorical Imperative, because no other persons were involved. But then, how could *any* virtual acts that don't affect other persons be morally wrong? In §§5-6, however, I argue that there indeed can be such cases of morally wrong virtual acts—some due to an actor's having erroneous beliefs about the morally relevant facts, and others due not to error, but to the actor's

---

[2] My characterization of virtual violence mirrors Luck's 2009 characterization of virtual murder, except that virtual violence is a broader category, since there are acts of virtual violence that aren't virtual murder.
[3] McCormick 2001, Sicart 2005, Coeckelbergh 2011, and Wonderly 2017: 36-7. Note that empirical work is required to determine virtual acts' actual effects on character, and vice versa.
[4] See Schulzke 2020 and Wonderly 2017 for discussion.
[5] See, e.g., McCormick 2001: 282-4; Schulzke 2010: 128-9; Wonderly 2017: 31-32.
[6] It's important to note that because Kant's Categorical Imperative is a fundamental and therefore *fully general* moral principle—see §3—*any* act that violates the Categorical Imperative is morally wrong from a Kantian perspective, whether or not the act is a kind of killing or even violence. In this paper I focus my discussion mostly on acts of virtual violence, especially virtual killing or murder, primarily for convenience and clarity, and also because such acts have been the central focus of much of the literature on the ethics of video gaming and virtual acts. But my overall argument in this paper might be applied to other kinds of virtual acts as well.

intention leaving out morally relevant facts while immersed in a virtual environment. I conclude by considering some possible implications of my arguments for both our present technological context as well as the future.

2. Virtual environments

Since I'm concerned in this paper with acts performed in virtual environments, before moving forward I'll say a bit about what I mean by "virtual reality", "virtual environment", and "immersion" in these environments.

Waving my hand, calling your name, elbow-bumping with you—these are all *acts* that I perform. Similarly, when swinging a sword or firing a rifle in games like *World of Warcraft* or *Call of Duty*; or when adjusting my virtual monitor in *vSpatial*'s Oculus headset-powered virtual workplace, I'm performing acts. I'll call these *virtual acts*, since they're acts done in a *virtual environment*. By 'virtual environment' I mean an environment with which one can engage via *virtual reality* (VR) technology. Brey defines 'virtual reality' broadly as "a three-dimensional interactive computer-generated environment that incorporates a first-person perspective" (1999: 5; 2008: 362), which, as Brey notes (ibid.), makes room for more or less immersive ways of accessing these environments—e.g., via traditional output devices such as monitors and speakers, or the latest VR headsets manufactured by Oculus. Sherman and Craig provide a similarly flexible characterization of 'virtual reality', which includes five key elements: a virtual world, immersion, interactivity, and the persons who participate in and create the virtual world (2019: 6ff). (Since the latter two are extrinsic to a virtual reality, I'll set them aside.) A *virtual world*—a term that, for my purposes, is synonymous with 'virtual environment'—is "an imaginary space often manifested through a medium. … a description of a collection of objects in a space and the rules and relationships governing those objects." (ibid., 8) For instance, the space in *World of Warcraft*'s virtual world—a world manifested through through visual and audio media—spans continents and oceans, and its objects include players' avatars and weapons, all of which are governed by familiar rules largely simulating the real world's physical laws. One *interacts* with objects in a virtual environment by way of a control interface (e.g., a keyboard and mouse), either directly as oneself or indirectly via an *avatar*, that is, a virtual representation of oneself or of a character with which one identifies at least to some degree.[7]

Finally, one can be more or less *immersed* in a virtual environment. By 'immersion' I mean an overall state of experiencing oneself *as being in* an environment, even if one isn't consciously thinking things such as "I'm really in this environment".[8] For instance, one is likely to feel more immersed while playing contemporary games like *Elder Scrolls V* or *Call of Duty: Modern Warfare*, which boast realistic graphics, environments in which one can interact from a first-person perspective, and engaging and intense plotlines in which one is an active participant; while one is likely to feel comparatively less immersed while engaging with older, much less realistic virtual environments or even recent but cartoon-like virtual environments such as in *Minecraft*. Features

---

[7] If one strongly identifies with one's avatar, perhaps one's interactions with virtual objects could count as direct.
[8] See Ramirez & LaBarge 2018: 251ff.

of immersed experiences—all or most of which are degreed features, but not all of which need be present at the same time—include deep mental engagement in the events of an environment; a suspension of disbelief concerning the reality of an environment or its events; not attending to one's immediate non-virtual environment; and the experience as of being physically present in an environment, in part mediated by sensory stimuli afforded by visual and audio technologies (e.g., monitors and speakers, or head-mounted displays), and perhaps soon haptic (sense of touch) and olfactory (sense of smell) technologies.[9] Various factors contribute to these features being present to a higher degree, or at least to their doing so more easily. For instance, in general, the more advanced the sensory outputs provided by graphics, audio, or other processing components, the more immersive one's experience in a virtual environment is likely to be. Similarly, the more one's non-virtual surroundings are blocked out—as, e.g., large, high resolution monitors and surround sound speakers do to some extent; or as VR headsets might do to a greater extent; or even as light-blocking curtains do to a lesser extent—the more immersive one's experience is likely to be. But being equipped with the most advanced sensory output devices isn't necessary to have experiences that are immersive to a significant degree.[10] Concerning input devices, one might suspect that interacting with a virtual environment by means of a physical control interface (e.g., a mouse and keyboard) would prevent immersive experiences. But, as any moderately experienced player (or simulation pilot, for that matter) knows, the more one becomes proficient with a control interface, the more that interface fades from one's awareness. At a certain point, one simply acts, no longer attending to the interface, in much the same way that a proficient pianist no longer attends to which keys she strikes when playing chords and arpeggios.[11]

3. Kantian moral theory

For the purposes of this paper, I'll assume what I take to be, for the most part, a mainstream understanding of Kant's moral theory. In this section, I'll outline the parts of this moral theory that are most important for my argument in the sections that follow. Since in this paper I'll be focused primarily on the moral rightness or wrongness of virtual acts, I'll accordingly narrow my focus in this section to parts of the Kant's theory of right and wrong conduct. But I won't attempt to defend what follows as the best interpretation of Kant.

From a Kantian perspective, the question whether an agent's act is morally permissible is a matter of whether the agent's intention, in performing the act, satisfies Kant's fundamental moral principle, the Categorical Imperative, which is binding on all rational agents. One of Kant's most well-known formulations of the Categorical Imperative is the Humanity as an End formulation (HEF). Kant says, "So act that you use humanity, whether in your own person or in the person of

---

[9] Multi-sensory experience in VR is a priority in near-future technology development (Gartner 2019). For discussion of haptic technologies for VR, see Biswas and Visual 2021. For discussion of multi-sensory stimuli in VR, including olfaction, see Melo et al. 2020.
[10] Sanchez-Vives & Slater 2005, Ramirez & LaBarge 2018: 251ff.
[11] Example borrowed from Annas 2011: 13.

any other, always at the same time as an end, never merely as a means." (G 4:429)[12],[13] According to the HEF, each of us is rationally bound to respect the humanity of others (and ourselves), which is to say that we are bound to treat any such being, without condition, as an end of any act, never as something merely to be used as a means for achieving another end. To treat something as an *end* is to treat it as something that has value. To treat something as a *means* is to treat it as something to be pursued not (just) for itself, but as a way to pursue something else that has value. But to treat something *merely* as a means is to treat it as if its *only* value is as a means to some other end. When I ask the bartender for a drink, I treat him as a means, but not *merely* as a means. But if I were to rob the bartender at gunpoint, arguably I would treat him *merely* as a means to some other end (viz., getting money), since I would fail to respect his humanity, i.e., his being a rationally autonomous being (or "person"). So, an act is morally wrong just in case in so acting the agent fails to treat others as ends, treating them instead as mere means; that failure is what makes the act morally wrong.

I'll focus primarily on another of Kant's well-known formulations of the Categorical Imperative, namely, the Universal Law formulation (ULF), according to which we ought to "act only in accordance with that maxim through which you can at the same time will that it become a universal law." (G 4:421). As I read Kant, he intends the ULF, not the HEF, to be the primary method for determining whether or not a course of action is morally permissible.[14] For instance, after having discussed the HEF and ULF—and one or perhaps two other formulations—Kant remarks, "But one does better always to proceed in moral *appraisal* by the strict method and put at its basis the universal formula of the categorical imperative" (G 4:436-7, Kant's italics). In any case, I'll assume that if an act violates any of Kant's formulations of the Categorical Imperative—and thus, if it violates the ULF—the act is morally wrong.

At the heart of the ULF is the notion of acting in accordance with *maxims*. We can put the basic schema of a maxim like this: *in circumstances C, I shall perform act A, in order to attain goal G*. This is a "subjective principle of volition" (G 4:402; cf. G 4:422): It is how the agent *represents to herself* the circumstances, the act, and the goal. So, for instance, suppose Aiden raises his hand in class. We might put Aiden's maxim like this: *when in class, and the instructor says something I don't understand, and raising my hand would get the instructor's attention* [C], *I'll raise my hand* [A] *in order to get my question answered* [G]. No doubt Aiden's maxim could be analyzed further, but this simple version suffices for my purposes.

---

[12] I cite Kant's works by title abbreviation, and using volume:page number which which can be found in the margins of each of the following: G = *Groundwork of the Metaphysics of Morals* (1997a); LE = *Lectures on Ethics* (1997b); MM = *Metaphysics of Morals* (1996).

[13] While Kant here formulates the Categorical Imperative in terms of "humanity", what's important for Kant is not so much being a member of the species *Homo sapiens*, but rather being a rationally autonomous being (or 'person'), whether of the human variety or, if there are (or could be) such things, extraterrestrials or artificial intelligences. See G 4:425. Famously, though, Kant denied that non-human animals have the requisite rational autonomy, and thus that we have no moral duties toward them. See LE 27:413. Some contemporary philosophers—notably, Korsgaard 2018—argue that a Kantian ethics can account for duties toward animals. For the sake of simplicity, in this paper I'll set aside the important question about the moral status of non-human animals. But nothing essential to my overall argument depends on this question.

[14] On this interpretive point, I follow Timmons 2017, ch. 2.

Is Aiden's act—his hand-raising—morally permissible, from a Kantian point of view? If it is, it must be that his maxim satisfies the ULF, which involves satisfying a two-stage requirement: first, we must be able to *conceive* without contradiction that everyone adopt the maxim; second, we must be able to *will* without contradiction that everyone adopt it. A failure in the first stage constitutes a *contradiction in conception*, which is, roughly speaking, to conceive a state of affairs that is in some sense incoherent, perhaps straightforwardly logically contradictory. A failure in the second stage constitutes a *contradiction in the will*, which is, roughly speaking, to will that a state of affairs obtains, but one which any agent, qua rational, already necessarily wills *not* to obtain.[15] Clearly Aiden's hand-raising is morally permissible. We can consistently conceive of everyone adopting his maxim, and we can will it without generating a conflict in the will.

Kantian moral theory is concerned centrally with an agent's *intentions* in acting. It isn't a consequentialist moral theory: The goodness or badness of an act's consequences aren't relevant, strictly speaking, to the rightness or wrongness of the act. So, from a Kantian point of view—and this is a point sometimes missed—even if an act *in fact* harms or misuses no persons, this doesn't show that the act is morally *permissible*. It might still be morally wrong. Consider an example:

Beth the bakery thief

Beth left her apartment in a hurry. She had no time for breakfast and, worse, forgot her wallet in the rush. Speed-walking past the bakery, she spots a platter of pastries on display. No bakery employees are in sight. "It's only a pastry", she reasons. "They won't notice the loss. And I'm starving!" She grabs a pastry without paying, and dashes off toward her workplace. But in her haste, Beth didn't notice the sign above the pastries: "Free pastries. Please take!"

Even though Beth's act didn't in fact deprive the baker of any goods, her act was still morally wrong, from a Kantian point of view. Beth *intended* to steal the pastry. We can suppose Beth's maxim was something like this: *when I'm hungry but can't pay, and when there's food available for the taking without risk to myself, take the food to satisfy my hunger*. One way to explain why Beth's maxim violates the ULF is this: If Beth were to will that everyone adopt her maxim, she would be willing a world in which she might well come across someone who, acting on that maxim, would rob her of her food in a situation in which she needs that food to survive. However, Beth *can't* rationally will this maxim to become a universal law, since, qua rational agent, she necessarily already wills that conditions obtain under which she can continue as a rational agent, which includes having the basic necessities for bodily survival.[16] So, Beth's willing this maxim brings about a contradiction in her will, failing to satisfy the ULF.[17] But the important point here is that, whatever explanation a

---

[15] O'Neill (1975) introduced the terms 'contradiction in conception' and 'contradiction in the will'. The precise nature of these contradictions, however, is a matter of scholarly debate (see Kleingeld 2017). In this paper I've assumed O'Neill's interpretation (1989: 132), but competing interpretations likely wouldn't change my argument in this paper.

[16] Here I follow Herman's general view of the contradiction in the will in such cases (1993: 121). cf. Kant MM 6:419-421.

[17] Some might worry about the 'latitude problem', viz., that contradictions in the will (i.e., second-stage failures of the UFL) violate only imperfect duties, and thus allow the agent some latitude, even for bullying or killing for fun. But Kleingeld (2019) argues persuasively that while Kant's ethics allows some latitude for, e.g., acting on the duty of beneficence, there is no latitude for acting on maxims of non-beneficence, which would

Kantian gives for why Beth's maxim (or any maxim) fails to satisfy the ULF, the explanation won't appeal to any of the actual consequences brought about by Beth's act. Her intention—her maxim—is what matters. So, from a Kantian point of view, an act can be morally wrong *even if* no persons are harmed or misused by the act. But then, at least prima facie, we ought to think the same is true of virtual acts: The fact that a virtual act *doesn't* harm any real persons isn't by itself sufficient to conclude that, from a Kantian point of view, the act *isn't* morally wrong. Thus we can ask: Could someone's act in a virtual environment *ever* be morally wrong, that is, ever violate the Categorical Imperative?[18]

4. The challenge: a baseline pair of cases

Contrary to the standard view, I'll argue that it's indeed possible to act morally wrongly in a virtual environment, even granting a Kantian moral theory, and even when no other persons or avatars are present. According to what I'm calling 'the standard view', one's virtual act can violate Kant's Categorical Imperative only if one is interacting with other persons or their avatars.[19] Of course, an advocate of the standard view might think that, even in the absence of other persons or their avatars, one can perform virtual acts that are in *some sense* morally problematic from a Kantian perspective. For instance, a habit of simulating violence toward non-player characters (NPCs) in virtual environments might well rub against one's *imperfect* duties—e.g., to adopt as a goal the improvement of one's own moral character.[20] Or, since some NPCs might be, somewhat like animals, similar to humans in morally relevant ways, perhaps a habit of mistreating (or simulating the mistreatment of) them might make us more inclined toward mistreating real persons and thus violating the Categorical Imperative.[21] Further, perhaps one can do wrong while knowingly mistreating another person's avatar, for instance, by killing and/or looting it, since this mistreatment can fail to satisfy the HEF formulation of the Categorical Imperative.[22] But, again, on the standard view, one can't violate the Categorical Imperative if one isn't interacting with any other persons or their avatars.[23]

---

include bullying or killing for fun.

[18] One might also ask whether or not, in performing a given virtual act, one is acting *from a good will*. For Kant, one acts from a good will when one's sole motive in acting is for the sake of duty itself. This is a question of the *moral worth* of the act, but not strictly speaking a question of the *moral rightness* of the act. My concern in this paper is the latter, not the former, and so my focus is on whether a virtual act fails to satisfy the Categorial Imperative. A further question is whether an act satisfies the Categorical Imperative if and only if the act is done from the sole motive of duty. This is a matter of interpretive dispute beyond the scope of this paper, but see Timmons 2017, ch. 5; Herman 1993, ch. 1.

[19] The authors cited in note 5, as I read them, hold or at least lean toward this position on the Categorical Imperative and virtual acts.

[20] Waddington 2007: 122-125, Wonderly 2017: 32.

[21] Brey 1999: 9, McCormick 2001: 283, Waddington 2007: 124-125, Schulzke 2010: 128, Wonderly 2017: 32. These authors draw on Kant's discussion of our treatment of animals in LE 27:459 (cf. MM 6:443). For extended discussion of Kant's views on ethics and animals, see Kain 2010.

[22] See Ryland 2019. Cf. McMillan & King 2017. For some relevant empirical work on the effects of persons' connection to their avatars, see Chappell et. al. 2006, Simon et. al. 2009.

[23] Might virtual acts in such cases violate the Categorical Imperative in virtue of violating a duty to *oneself*? (See Waddington 2007: 124 for this sort of suggestion.) Perhaps in some cases, but I suspect that violations of duties to self won't explain what is primarily morally concerning about many virtual acts, including some of

The standard seems persuasive. In order to make clearer the challenge my argument against this view faces, I'll begin by presenting a pair of cases meant to highlight further why, prima facie, the standard view seems correct. The first is a 'real world' case—i.e., one set in a non-virtual environment—in which the agent's act is intuitively morally wrong, an act that quite plausibly fails to satisfy the ULF. The second case is analogous to the first, except that the person acts in a virtual environment. In the second case, however, while we might have moral qualms with the person's act, it's plausible that the act *satisfies* the ULF and thus isn't, strictly speaking, morally wrong. And that's just what we should think, according to the standard view.

Consider the following case of an assault in the real world.

Bob in the Alley

Bored one evening, Bob hides in the shadows of the alley between his apartment building and the adjacent pub. Moments later, Chuck leaves the pub, entering the alley on his way home. Bob leaps out, pushes Chuck against the brick wall, and punches him for several minutes before sending him off bruised and bleeding, Bob laughing all the while.

Was Bob's punching Chuck morally wrong? Intuitively so. It's also plausible that the maxim on which Bob acted fails to satisfy the ULF. We can suppose that Bob's maxim was something like this: *when I'm in a place where I can catch someone unaware and punch them, and where there's little risk, punch them in order to have some fun*. Bob can't rationally will this maxim to become a universal law. In doing so, Bob would will a world in which he'd likely come across someone who would harm him. But, qua rational agent, Bob *must* will that conditions obtain under which he would continue as an autonomous agent; and therefore, qua rational agent, he necessarily wills that conditions *don't* obtain that would preclude his rational autonomy. So, Bob's willing his maxim brings about a contradiction in his will, failing to satisfy the ULF. So, Bob's act is morally wrong, which intuitively is the correct verdict.

Now consider an analogous case, but one in a virtual environment, and one which seems to *satisfy* the ULF.

Vernor in the Alley

Vernor is playing an immersive, first-person perspective video game—one with an open plot line, allowing him to pursue whatever career he likes. Bored one evening, Vernor hides in the shadows of the alley between his virtual apartment building and the adjacent virtual pub. Moments later, Nolan, an NPC, leaves the pub, entering the alley on his way home. Vernor leaps out, pushes Nolan against the brick wall, and punches him for several minutes before sending him off bruised and bleeding, Vernor laughing all the while.

Was Vernor's virtual attack morally wrong? Some readers' intuitions will pull toward the judgment that Vernor's act is morally wrong, or at least that there's *something* morally problematic about the case. For instance, we might be inclined to evaluate negatively Vernor's *character*, since Vernor

---

those done in the absence of other persons or their avatars. Since I'll be arguing against the standard view without appeal to duties to self, I'll set this question aside.

seems to take pleasure in simulating harming innocent people. But set that aside. The question on which I'm focused is whether Vernor's act of virtual punching is morally right or wrong. Plausibly, Vernor's act *satisfies* the ULF, in which case his act is *not* morally wrong, strictly speaking. We can suppose that Vernor's maxim was something like this: *when I'm acting in a virtual environment where I can catch an NPC unaware and punch them, and where there's little risk, punch them in order to have some fun*. In this case, Vernor *can* rationally will this maxim to become a universal law. One can successfully conceive of a world in which everyone adopts this maxim; and, plausibly, one also can will without contradiction that it become a universal law. Since Vernor isn't himself an NPC, a world in which everyone punches NPCs in virtual environments needn't conflict with anything he necessarily wills as a rational agent. Even if virtual punching of NPCs should become widespread, this wouldn't seem to conflict with his continuing to persist as a rational agent. But then Vernor's virtual punching *satisfies* the ULF, and thus *isn't* morally wrong, strictly speaking. And it's fairly obvious why: Vernor's act has no *persons* (or even avatars) as its objects; and so Vernor's act, in a person- and avatar-free virtual environment, involves no intention to harm or misuse any persons.

Here, then, is the challenge: given that the Kantian moral verdicts diverge in the analogous Bob and Vernor cases, how could *any* case of acting in a person- and avatar-free virtual environment be morally wrong, since in all such cases, ex hypothesi, there are no other persons or avatars present? Notice, too, that Bob's and Vernor's respective maxims *represent accurately* the morally relevant facts of their cases, including facts concerning the objects of their acts.

In the following two sections, I'll consider a couple of broad kinds of cases. In §5, I'll consider cases of maxims that, due to errors in the agent's beliefs, *misrepresent* morally relevant facts. In §6, I'll consider cases in which, despite the agents' having no erroneous beliefs, their maxims still leave out morally relevant facts while immersed in the virtual environment. I'll argue that, for both kinds of cases, there *can* be virtual acts that fail to satisfy the ULF and thus are morally wrong from a Kantian perspective.

5. Cases involving erroneous beliefs

From a Kantian perspective, what matters for determining the rightness or wrongness of one's act aren't the act's *consequences*, but rather one's *intention*—in Kantian terms, one's *maxim*—in acting. So, even if one's act turns out to have bad results—and even if those results are worse than the results of something else one might have done—it nevertheless doesn't follow immediately that the act was morally wrong. But then it's true, too, that even if one's act turns out to have *good* results, it doesn't follow that the act *wasn't* morally wrong. Sometimes we do wrong, but things still turn out well, despite our bad intentions. And sometimes when we do wrong in this way, things turn out well precisely because we had erroneous beliefs about some of the morally relevant facts in the circumstances. Clearly this can happen in everyday non-virtual contexts. Consider such a case:

Manny and the Mannequin

Atop a parking structure, across the street from a shopping mall's entrance, Manny makes the final distance adjustments on his rifle as the sun sinks below the horizon. After months of simmering anger, this morning he resolved to punish someone, anyone. Peering through the

scope, he plans to shoot the first employee to exit. Reflections on the glass doors obscure his view inside, so he'll have to be ready. But unbeknownst to Manny, the police were tipped off hours earlier. Mall employees already have been ushered out a rear exit. Hidden SWAT team members stand by to apprehend Manny, but they don't have his location. Aware of the reflections on the glass doors, officers inside the mall's entrance have attached a mannequin to a remote-controlled bomb disposal robot. They begin to drive the mannequin out of the doors. Manny sees the human form and fires.

Intuitively, Manny's act—his shooting at the mannequin—was clearly morally wrong. His erroneous belief that the mannequin was a human doesn't absolve him of moral wrongdoing. Manny *intended* to kill a person. That is, as a result of Manny's belief that the mannequin was human, Manny's maxim represented the circumstances of his act as involving a person as the object of his act, the target of his shooting. Manny's maxim fails to satisfy the ULF, since attempting to universalize his maxim would result in a conflict of the will. By willing that everyone adopt his maxim, he would thereby will a world in which he might find himself in the crosshairs of someone intending to kill him. But Manny cannot rationally will this, since qua rational agent he necessarily wills that conditions obtain under which he continues as an autonomous rational agent. So, Manny's act is morally wrong, even though some of the morally relevant facts (e.g., that the mannequin isn't a person) are misrepresented by his maxim.

What of an analogous virtual case? Unlike the pair of baseline Bob and Vernor cases in which the moral verdict in the virtual case doesn't mirror the verdict in the real-world case, in the following virtual case the moral verdict *does* mirror the one in the Manny case above. While the following case isn't feasible given today's technology, it is nevertheless *possible*—perhaps even plausible at some point in the future, if VR technology continues to advance. For convenience, I've borrowed from Ernest Cline's (2011) *Ready Player One* the idea of the OASIS, an advanced virtual universe accessible using a full-body "immersion rig", which provides completely immersive haptic (i.e., touch-stimulating), audio, and visual sensory experiences, often indistinguishable from non-virtual reality.

Snyder the Unwitting Virtual Sniper

After months of simmering anger, this morning Snyder resolved to punish someone, anyone. He spent much of the day logged into a simulation of his city in the OASIS—though one populated only by NPCs—preparing and practicing his plan to exact his punishment. Satisfied with his preparation, Snyder switched off his VR helmet, disengaged from his immersion rig, and logged out of the OASIS. Soon after, atop a parking structure across the street from the entrance to a local shopping mall, Snyder makes the final distance adjustments on his rifle as the sun sinks below the horizon. Peering through the scope, he plans to shoot the first employee to exit. Reflections on the glass doors obscure his view inside, so he'll have to be ready. Moments later, an employee exits, car keys twirling on her finger. Snyder squeezes the trigger. The employee falls. But unbeknownst to Snyder, he never *actually* switched off his VR helmet, disengaged from his immersion rig, or left the OASIS, despite its appearing so to him. Police were tipped off hours earlier, and in a plan to protect the public and arrest Snyder for conspiracy to commit murder, police technical staff hacked Snyder's OASIS account and spoofed his logout process. Snyder

isn't aware he's still online—until the police raid his home and arrest him.

Intuitively, Snyder's act—his shooting the NPC employee—was morally wrong. His erroneous belief that the NPC was a real human doesn't absolve him of moral wrongdoing. Snyder *intended* to kill a *person*, despite his maxim's misrepresenting the NPC as a person. So, his act is morally wrong: His maxim fails to satisfy the ULF, and for reasons similar to those in Manny's case.[24]

Of course, cases like Snyder's, in which one mistakenly believes oneself to be in a non-virtual environment, are extremely unlikely to occur given present technology. But nothing about such cases is impossible. Perhaps in the future—how far into the future is hard to gauge—VR and related technologies will become advanced enough to fool our senses, at least temporarily.[25] We ought to think through these sorts of moral situations before we find ourselves in them, so that we'll be more likely to be better moral agents when we get there, if and when the technology ever arrives.

Even though the technology required to make cases like Snyder's a reality is, at best, still a ways off into the future, nevertheless such cases are sufficient to show that it's indeed *possible*, from a Kantian perspective, to do moral wrongs in virtual environments in which one is alone. But are there any plausible cases of acting morally wrongly in virtual environments employing *today's* technology? In the following section, I argue that there are.[26]

6. Cases involving no error, but that leave out morally relevant facts

It's plausible that sometimes, while we don't have any erroneous beliefs about any of the morally relevant facts about our circumstances, we end up doing what's morally wrong because our maxim—the way we represent our circumstances and goal when we're acting—*leaves out* some of what's morally relevant about our circumstances. Of course, in many cases maxims that leave out a few morally relevant details *won't* render our acts wrong—when, for instance, what a maxim *does* represent entails that no persons are possible or at least likely objects of the act. However, plausibly for some (or maybe even many) of our acts, these kinds of maxims *do* render these acts morally wrong. This shouldn't be shocking, since we act morally recklessly to one degree or another when our maxims leave out morally relevant details of our circumstances. Put more colloquially,

---

[24] Snyder's maxim, like Manny's, represented the circumstances of his act as involving a person his target. But then, willing that everyone adopt his maxim would produce a conflict in the will: he'd be willing a world in which he might be the target of someone intending to kill him, conflicting with his necessarily willing the conditions of his own continuance as a rational agent.

[25] For general discussion about how the future of VR and related technologies might unfold, see Greengard 2019. For discussion of increasingly immersive multi-sensory VR technologies, see note 9. Another example of a new, more sensory-immersive technology is gaze-tracking, by which a VR display tracks the user's eye and head movements as inputs. See Pai et al 2019. A technology closed related to VR is augmented reality (AR), in which perception of one's environment is overlaid with virtual, computer-generated imagery and/or other sensory information. Development of AR/VR contact lenses, which might have increased potential to fool our senses at least temporarily, is already on the horizon (Savitz 2021). Finally, a much more speculative frontier of VR involves whole-brain emulation and mind uploading, which, if possible and ever implemented, could provide a virtual environment completely indiscernible from the non-virtual world. See Sandberg and Bostrom 2008.

[26] I won't, however, attempt to answer the question *how often* these kinds of cases happen. That's a question best left to empirical study.

sometimes we don't think enough about what we're doing, and sometimes this amounts to doing wrong. Sometimes we don't form beliefs about some of the morally relevant facts, and so we don't represent those facts when we act, as in Rhonda's case below. But other times, even though we've got beliefs about the morally relevant facts, when we're acting 'in the moment', some of those beliefs fade from our awareness, and so those beliefs aren't reflected in the way we represent our situation when acting, as in Reed's case below. In such cases, *were* one to have adequately and accurately represented the morally relevant facts about one's circumstances, one's act may (or may not) have been morally permissible; it depends on whether the adequate and accurate maxim would satisfy the ULF.

Consider how this might happen in a non-virtual case:

Rhonda the Rural Shooter

Rhonda enjoys sitting on her front porch, rifle in hand, waiting to shoot at whatever crosses the prairie in her field of view. In the past, she's shot at the unfortunate few people who have taken a shortcut across her property. She's also shot at wild deer and the occasional escaped cow from a nearby farm. But, partly due to her failing eyesight and deteriorating temperament, Rhonda no longer gives a thought to what *kinds* of things she aims at. So long as it seems to be moving itself across her prairie, she gets a kick out of shooting at it. This morning, she spots an object crossing her property fifty yards out. She aims and fires.

Intuitively, in my view, Rhonda's act is morally wrong. But why, from a Kantian point of view? Rhonda's intention represented nothing about shooting at a *person*. While her thoughtlessness surely justifies evaluating negatively her moral character, does it give us grounds to judge the moral rightness or wrongness of her act? It might not be immediately clear how her act could violate the ULF. After all, as noted above, Rhonda's moral thoughtlessness also entails her lacking any intention to harm a person. We can suppose that Rhonda's maxim is something like this: *whenever there's something big moving itself across my property, shoot the thing, for the thrill of it.*[27] Her maxim represents no humans as targets; nor do the circumstances represented entail that humans are at risk. But then, *can* Kantian ethics vindicate the intuition that Rhonda's act is morally wrong? I believe so.

The class of objects represented by Rhonda's maxim doesn't *rule out* humans as possible objects of the act; nor do the circumstances represented rule this out. Rhonda's maxim leaves out a morally relevant fact, namely, *what* the objects of the act *are*. And clearly the class of things represented as possible objects of the act—viz., big things moving across her property—*include* humans as a sub-class. And it isn't as if humans are unlikely to be members of this class. So, it's plausible that Rhonda's maxim fails to satisfy the ULF. Rhonda can't rationally will her maxim to be a universal law, since in doing so she would will a world in which she might well find herself walking into the crosshairs of someone like herself. Willing for this world to obtain would conflict with what she, qua rational agent, necessarily wills, namely, the conditions of her own continuance as a rational agent. So, even though Rhonda didn't *represent* her act as harming or misusing

---

[27] For readers wondering about the 'problem of relevant descriptions', see note 28.

persons, her maxim nevertheless evinces a moral recklessness that is morally wrong both intuitively but also according to the ULF.

Had Rhonda been more thoughtful in acting—had her maxim represented all of the morally relevant facts—would she have acted morally permissibly? It depends. Had she represented her target as, for instance, a wild animal, perhaps her maxim would satisfy the ULF.[28] Or, supposing she lives near Boston Dynamics's laboratory, and represented her target as Spot, their quadruped robot, her maxim might satisfy the ULF. By contrast, had she represented her target as a human being, then, even supposing this further specification would render her maxim complete with respect to the morally relevant facts, clearly this wouldn't improve her maxim morally.

So, in some cases one *can* act morally wrongly even when one doesn't intend to harm any persons, and even when one isn't in error about one's circumstances. One can act morally recklessly—i.e., act on maxims that leave out morally relevant facts—in such a way that one's maxim *would* generate a conflict in the will and thus fail to satisfy the ULF. But can one act morally recklessly in this way while in a *virtual* environment with no other persons or avatars present? Plausibly so. Consider the following case, based on the contemporary, immersive online game, *Red Dead Online*. This case is longer than those above, since a fuller description of the virtual environment helps to bring out the agent's circumstances.

Reed the Virtual Outlaw

Reed spends his free hours of unemployment playing *Red Dead Online*, which is set in the nineteenth century American West, and features an open virtual world in which players are free to explore and interact. The game boasts graphics approaching photo-realism, especially when run on high-powered computers like Reed's. *Red Dead Online*'s NPCs are programmed with robust AI, and each has hundreds of conversation responses recorded by voice actors. Players are outlaws, building their avatars' in-game experience and notoriety through interactions with both NPCs and other players, often with virtual violence.

For highly immersive gaming, Reed always employs the game's more immersive first-person perspective. He surrounds himself with ultra-wide monitors, with the lights off and sun-blocking curtains drawn; he sees nothing beyond the game's virtual environment. Nor does he hear anything but the game's audio, thanks to sound-isolating headphones. Countless hours in *Red Dead Online* and years of experience with similar games have made Reed an expert player; he rarely notices the game's control scheme or his own use of input devices. When acting virtually, he doesn't represent himself as using such and such button combinations and mouse movements; rather, he *simply acts* in the virtual environment—he *runs*, he *fires his pistols*, and so on.

Reed aspires to be the most notorious outlaw in *Red Dead Online.* He's planned and executed attacks against rival players, and, most recently, robbed and killed NPC sheriffs and

---

[28] At least, Kant would've thought so, since he thought that only *rationally* autonomous beings are persons, and thus that we have no moral duties toward animals. See note 20. Although in this paper I've set aside the question whether animals have moral status, the essentials of my discussion needn't change to account for animals' having moral status.

shopkeepers in a string of towns. Tonight, Reed enters yet another town—one without any other players' avatars present—sneaks into the sheriff's office, and dispatches the sleeping deputy with his shotgun. On his way to the general store across the street, Reed is confronted by a sheriff and another deputy—NPCs both—with guns drawn, who yell, "Freeze, mister!" Reed runs, firing his shotgun, downing the deputy. He bursts into the general store, startling the (NPC) shopkeeper. The sheriff pursues, firing shots at Reed, punching holes in the shop's door. Reed returns fire, toppling the sheriff onto his deputy in the street. The trembling shopkeeper quickly raises his hands and begs for mercy. Grinning, Reed takes a moment to aim, and guns down the shopkeeper. And, as he's done several times tonight already, he laughs aloud, watching the blood pooling on the floorboards. He loves the thrill of the kill, every time.

Minimally, Reed's behavior is morally worrisome. Some of one's moral reaction to Reed's case likely homes in on his character. He appears to take pleasure in, and might even be obsessive about, virtual violence. But it's also plausible that some of Reed's virtual *acts* are morally wrong—for instance, his virtual killing of the NPC shopkeeper. If, from a Kantian perspective, Reed's virtual killing is indeed morally wrong, it must be that the maxim on which he acted violates the ULF. But then the question is: *Could* Reed's maxim violate the ULF? It might seem not. After all, unlike Snyder, Reed never held a belief that the target of his act was a person; on reflection, or if asked, Reed would affirm that the shopkeeper was only a virtual person. And unlike in Rhonda's case, Reed's act took place in a virtual environment with no other persons or avatars present, so no persons ever were in fact plausible objects of his act.

And yet I think Reed's virtual act *does* fail to satisfy the ULF. Of course, Reed surely held a background belief that the town's sheriffs, deputies, and shopkeepers aren't real persons; and so, had Reed gunned down the NPC shopkeeper moments after switching on his computer and signing in to *Red Dead Online*, his maxim would likely be something like: *when I see a virtual person*, or, *when I see a character in the game, shoot them for the thrill of it*. Reed would still be attending to the facts that the town is *virtual*, that the character he targets is *virtual*, that he's playing *a game*. But in the case above, Reed had come to be acting in a highly immersed and intense state. The virtual environment's visual and auditory realism, Reed's sensory isolation, his expertise with the game and its control scheme, his hours-long gaming session, and his frantic encounter with, and reactions to, the sheriff and deputy—all of these factors contribute to it being plausible that Reed became sufficiently immersed in the game's virtual environment so that, when he shot the shopkeeper, he *didn't* represent himself as navigating a virtual environment and manipulating a control interface to fire virtual bullets at a virtual character. Rather, while 'in the moment', immersed in the game's environment, events, and overall narrative, as well as his character's narrative—with which Reed identifies to some extent—it's plausible that Reed simply represented his act as *shooting the shopkeeper*. And this is true even though Reed would, if asked or if he took a moment to reflect, affirm that the shopkeeper was only a *virtual* person. Nevertheless, when he acted in a state of immersion, Reed was no longer attending to the morally relevant fact that his target was—and so he no longer represented his target as—a *virtual* shopkeeper, and likewise for the fact that the town was virtual, and so on. Under these conditions, the fact that his environment and its objects were

merely virtual things had dropped out of his representation of them.[29] Put another way, it's plausible that, when shooting the shopkeeper, Reed represented his act and circumstances in a way that is indistinguishable from how he might represent his act and circumstances in an analogous but non-virtual situation. Thus, it's plausible that Reed's maxim, while in his state of immersion, was something like this: *when I'm in the presence of someone who's defenseless and cornered, shoot them for the thrill of it*.[30] Were Reed to will that everyone adopt this maxim, he would thereby will a world in which he might well come into contact with someone intent on shooting him, conflicting with his necessarily willing, qua rationally autonomous being, the conditions of his own continuance as such a being. Such a maxim would fail to satisfy the ULF.

Did Reed slip into *believing* that the shopkeeper was a real person? No, in my view. As noted above, if we were to have asked Reed whether the shopkeeper was a person—either by interrupting his gameplay or by asking him later on—he would have attended more fully to the relevant circumstances, denied that the shopkeeper was a person, and affirmed that the shopkeeper was merely virtual, a character in a game.

But did Reed slip into *representing* the shopkeeper *as being a person*? On the one hand, we might think, no, not quite. Despite Reed's having been immersed in the virtual environment, with person-negating representational contents such as 'virtual' and 'only in the game' having faded from his occurrent psychological state, still, it seems unlikely that 'is a person' or 'person' was *added* to his representation of the shopkeeper. And so perhaps Reed was representing the shopkeeper *as if* it were a person, but not explicitly representing it *as being* a person. But even if not, Reed's maxim would fare no better with the ULF, for the same reasons given above. And his maxim still would be indistinguishable from one he might employ under analogous but non-virtual circumstances—after all, we surely don't always explicitly represent every person as *being a person* or as *a real person* in our everyday interactions with people.

On the other hand, we might think that Reed *did* indeed slip into representing the shopkeeper as being a person. Perhaps his state of immersion, and the fading out of representational contents like 'virtual' or 'only in the game', was indeed enough so that Reed *just was* representing the shopkeeper as being a person, even without any explicit 'is a person' content. If so, then Reed's maxim even more straightforwardly violates the ULF.

I won't try to settle here this question about precisely how we should understand Reed's representation of the shopkeeper while immersed. But on either of the options sketched above, the content of Reed's maxim leaves out the morally relevant facts that the shopkeeper and the

---

[29] Perhaps entering an environment with the background belief that no persons are present might even *lower* one's moral inhibitions or preclude priming one's moral sensitivities. If so, this could make it more likely, at least in circumstances like Reed's, that one does wrong.

[30] One might ask: how do we know how much detail Reed's maxim actually represents? Perhaps the maxim represents 'shopkeeper' rather than 'someone who's defenseless and cornered'. First, it's plausible that often, in our maxims, many of the particular features of the circumstances are abstracted away. It's highly unlikely that we represent our circumstances in full detail. In any case, the problem noted here—often called the *problem of relevant descriptions*—is one that any Kantian moral philosopher must address, so it isn't a problem for my argument in particular. For discussion, see Anscombe 1958, O'Neill 2004, and Timmons 2017, ch. 2.

circumstances were virtual. And so, either way, it's plausible that Reed's maxim would violate the ULF. Thus it *does* seem plausible that one can act morally wrongly, from a Kantian perspective, even in virtual environments with no other persons or avatars present, and even using contemporary technology.

Concluding Remarks

If what I've argued is correct, even Kantians ought to be—and ought to think we all ought to be—morally on guard when immersed in virtual environments. The virtual is not a moral safe haven against the Categorical Imperative.

With an eye towards the future, video games and other sorts of virtual platforms are likely to become more immersive as VR and related technologies continue to develop, and their use will become more widespread. Circumstances in which we end up, like Reed, committing morally wrong virtual acts (from a Kantian perspective) might become more common. There might even at some point come to be a realistic risk that we could sometimes find ourselves mistaking—or being led to mistake—the virtual for the real, in which case the odds of doing virtual wrongs only increase. Further, if mind uploading is indeed metaphysically possible, and if we develop the technology to make it actual, these sorts of mistakes could become frighteningly easy to make.[31]

With an eye to the present, increasing numbers of people spend increasing numbers of hours playing video games,[32] and these games are increasingly realistic and immersive, many of them meeting the broad definitions of 'virtual reality' sketched in §3 above. Many of the acts that players perform in these games' virtual environments are of the sort that would be morally worrisome or even abhorrent in the real world. If what I've argued in this paper is correct, the odds only increase that people will sometimes, in some virtual contexts, perform acts that Kantians ought to evaluate as morally wrong. Moreover, new frontiers for immersive virtual environments are being imagined—e.g., virtual prisons,[33] the use of VR by prisoners,[34] and even the provision of VR simulations as a kind of catharsis or alternative outlet for pedophiles.[35] These novel uses of virtual environments plainly have moral implications, including the potential for users to act in ways that, under immersion, might employ maxims that would fail to satisfy the Categorical Imperative. It behooves us, now rather than later, to think through the implications for ourselves as moral agents in these novel sorts of virtual circumstances—both the broader social implications, but also the narrower implications for the moral rightness or wrongness of our acts.[36]

---

[31] For the philosophical implications of mind uploading, see the essays in Blackford and Broderick 2014. For optimistic appraisals of the possibility of mind uploading, see Chalmers 2014 and Sandberg and Bostrom 2008.

[32] Anderson 2019.

[33] Gash 2020.

[34] Lewis 2018.

[35] Levy 2002, Cisernos 2002. For discussion, cf. Luck 2009, Patridge 2013.

[36] For helpful comments on the ideas treated in this paper, I'm grateful to Don Howard, Ting Cho Lau, Christian Miller, Raphael Mary Salzillo, audience members at the Georgia Philosophical Society, and two anonymous referees.

References

American Psychological Association. 2020. "Resolution on Violent Video Games: February 2020 Revision to the 2015 Resolution."

Anderson, Kevin. 2019. "Research Report Shows How Much Time We Spend Gaming." *Forbes.* https://www.forbes.com/sites/kevinanderton/2019/03/21/research-report-shows-how-much-time-we-spend-gaming-infographic/

Annas, Julia. 2011. *Intelligent Virtue.* Oxford University Press.

Anscombe, G. E. M. 1958. "Modern Moral Philosophy." *Philosophy* 33 (124):1-19.

Biswas, S., Visell, Y. 2021. "Haptic Perception, Mechanics, and Material Technologies for Virtual Reality." *Advanced Functional Materials.* 2008186. https://doi.org/10.1002/adfm.202008186

Blackford, Russell and Damien Broderick (eds). 2014. *Intelligence Unbound: The Future of Uploaded and Machine Minds.* Wiley Blackwell.

Brey, Philip. 1999 "The Ethics of Representation and Action in Virtual Reality," *Ethics and Information Technology*, 1(1):5–14.

Brey, Philip. 2008 "Virtual Reality and Computer Simulation," in Himma K.E. & Tavani, H.T. (eds.) *The Handbook of Information and Computer Ethics*, Wiley-Interscience, pp. 361-384.

Chalmers, David J. 2014. "Uploading: A Philosophical Analysis." In Russell Blackford & Damien Broderick (eds.), *Intelligence Unbound: The Future of Uploaded and Machine Minds*. Wiley Blackwell. pp. 102-117.

Chappell, D., Eatough, V., Davies, M. N. O., and Griffiths, M. 2006. "EverQuest—It's Just a Computer Game Right? An Interpretative Phenomenological Analysis of Online Gaming Addiction." *International Journal of Mental Health and Addiction*, 4(3), 205-216.

Cline, Ernest. 2011. *Ready Player One.* Crown.

Cisneros, Dannielle. 2002. "Virtual Child Pornography on the Internet: A 'Virtual' Victim?", *Duke Law & Technology Review* 1-8 (2002)

Coeckelbergh, Mark. 2011. "Virtue, Empathy, and Vulnerability: Evaluating Violence in Digital Games," in K. Poels & S. Malliet, (eds.) *Vice City Virtue: Moral Issues in Digital Game Play*. Leuven/Den Haag: Acco Academic, 89-105.

Gartner. 2019. "Gartner Top 10 Strategic Technology Trends for 2020". https://www.gartner.com/en/newsroom/press-releases/2019-10-21-gartner-identifies-the-top-10-strategic-technology-trends-for-2020.

Gash, Tom. 2020. "Forget prisons, the future of punishment will be virtual." *Wired.* https://www.wired.co.uk/article/virtual-prisons

Greengard, Samuel. 2019. *Virtual Reality.* The MIT Press.

Herman, Barbara. 1993. *The Practice of Moral Judgment.* Harvard University Press.

Kain, Patrick. 2010. "Duties regarding animals." In Lara Denis (ed.), *Kant's Metaphysics of Morals: A Critical Guide*. Cambridge University Press. pp. 210-233.

Kant, Immanuel. 1996. *The Metaphysics of Morals* [MM], trans. and ed. M. Gregor. Cambridge University Press.

Kant, Immanuel. 1997a. *Groundwork of the Metaphysics of Morals* [G], trans. and ed. M. Gregor. Cambridge University Press.

Kant, Immanuel. 1997b. *Lectures on Ethics* [LE], trans. P. Heath, Eds. P. Heath and J.B. Schneewind. Cambridge University Press.

Kleingeld, Pauline. 2017. "Contradiction and Kant's Formula of Universal Law." *Kant-Studien* 108 (1):89-

115.

Kleingeld, Pauline. 2019. "A Contradiction of the Right Kind: Convenience Killing and Kant's Formula of Universal Law." *Philosophical Quarterly* 69 (274):64-81.

Korsgaard, Christine M. 2018. *Fellow Creatures: Our Obligations to the Other Animals*. Oxford University Press.

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

Lewis, Nicole. 2018. "A View of Tomorrow." *The Marshall Project*. https://www.themarshallproject.org/2018/07/17/a-view-of-tomorrow

Luck, Morgan. 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 (1):31-36.

McCormick, Matt. 2001. "Is it wrong to play violent video games?" *Ethics and Information Technology*, 3:277-287.

McMillan, John and Mike King. 2017. "Why Be Moral in a Virtual World?" *Journal of Practical Ethics* 5(2):30-48.

Melo, M., G. Gonçalves, P. Monteiro, H. Coelho, J. Vasconcelos-Raposo and M. Bessa. 2020. "Do Multisensory stimuli benefit the virtual reality experience? A systematic review," in *IEEE Transactions on Visualization and Computer Graphics*, doi: 10.1109/TVCG.2020.3010088.

O'Neill, Onora. 1975. *Acting on Principle: An Essay on Kantian Ethics*. Columbia University Press.

O'Neill, Onora. 1989. *Constructions of Reason: Explorations of Kant's Practical Philosophy*. Cambridge University Press.

O'Neill, Onora. 2004. "Modern Moral Philosophy and the Problem of Relevant Descriptions." *Royal Institute of Philosophy Supplement* 54: 301-316.

Pai, Y.S., Dingler, T. & Kunze, K. "Assessing hands-free interactions for VR using eye gaze and electromyography." *Virtual Reality* 23, 119–131 (2019). https://doi.org/10.1007/s10055-018-0371-2

Patridge, Stephanie L. 2013. "Pornography, ethics, and video games." *Ethics and Information Technology* 15 (1):25-34.

Ramirez, Erick Jose and Scott LaBarge. 2018. "Real moral problems in the use of virtual reality." *Ethics and Information Technology*, 20:1249-263.

Rescorla, Michael. 2020. "The Computational Theory of Mind." *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2020/entries/computational-mind/>.

Ryland, Helen. 2019. "Getting away with murder: why virtual murder in MMORPGs can be wrong on Kantian grounds." *Ethics and Information Technology* (2).

Sanchez-Vives, M. V., & Slater, M. 2005. "From presence to consciousness through virtual reality." *Nature Reviews: Neuroscience*, 6, 332–339.

Sandberg, Anders and Nick Bostrom. 2008. *Whole Brain Emulation: A Roadmap*. Technical Report 2008-3. Future of Humanity Institute, University of Oxford.

Savitz, Eric J. 2021. "Apple Will Lead in AR, Analyst Says. Watch for Its Helmet, Glasses, and Contact Lenses" *Barron's*. https://www.barrons.com/articles/apple-will-lead-in-ar-analyst-says-watch-for-its-helmet-glasses-and-contact-lenses-51615223178

Schulzke, Marcus. 2010. "Defending the morality of violent video games." *Ethics and Information Technology*, 12:127-138.

Sherman, William R. and Alan B. Craig. 2003. *Understanding Virtual Reality. Interface, Application, and*

Design. Morgan Kaufmann Publishers.

Sherman, William R. and Alan B. Craig. 2019. *Understanding Virtual Reality. Interface, Application, and Design*, Second Edition. Morgan Kaufmann Publishers.

Sicart, M. 2005. "Game, Player, Ethics: A Virtue Ethics Approach to Computer Games." *International Review of Information Ethics*, 4(12), 13-18.

Simon, B., Kelly Boudreau, and Mark Silverman. 2009. "Two players: biography and 'played sociality' in EverQuest." *Game Studies*, 9(1).

Timmons, Mark. 2017. *Significance and System: Essays on Kant's Ethics*. Oxford University Press.

Waddington, David I. 2007. "Locating the wrongness of ultra-violent video games." *Ethics and Information Technology*, 9:121-128.

Wonderly, Monique. 2017. "Video Games and Ethics." In Joseph C. Pitt & Ashley Shew (eds.), *Spaces for the Future: A Companion to Philosophy of Technology*. Routledge.