



Between consenting avatars

Clemens Uhing^{1,2} · Bert Heinrichs^{1,2}

Published online: 9 April 2026
© The Author(s) 2026

Abstract

As virtual reality and virtual worlds in general take up more and more space in people's everyday lives, the question of how virtual actions should be morally evaluated becomes increasingly pressing. One of the main challenges in the ethics of virtual worlds, in our view, is to find a normative framework, that is not overly restrictive but still accounts for common intuitions, such as the immorality of discrimination or harassment in virtual realms. In this paper we ask which criterion could be suitable to allow for a more nuanced evaluation of various potentially immoral virtual actions. First, we introduce a case to distinguish possible harms and immoral behaviors in virtual worlds. We then present a concept of consent inspired by principlism in research ethics and discuss criteria for valid consent and its limits. In applying this to our case, we find that consent can render many allegedly immoral actions in virtual worlds – such as virtual killing or theft – morally permissible, while it cannot be used to justify discrimination or violations of personal identity. A consent-based approach is thus shown to successfully reconstruct common moral intuitions, like those on the gamers' dilemma. Finally, we address some objections to our argument, namely objections from deontological ethics and virtue ethics. The result of the paper is a framework for the moral evaluation of virtual actions, that allows for a more balanced approach and takes the particular purposes of virtual worlds and the interests of participating individuals into account.

Keywords Virtual Reality · Virtual actions · Consent · Ethics · Game ethics · Avatars

Introduction

As virtual reality (VR) is used by ever more people and over longer time periods, ethics of technology should take a closer look on potential moral challenges within virtual environments, especially in the consumer market for entertainment and more specifically games.¹ Two characteristics of virtual environments are essential for their ethical

sensitivity: modern technology renders the audiovisual simulation of nearly every “real” event and action possible, allowing for virtual representations of a broad spectrum of activities which would be considered immoral outside virtual environments. In particular since virtual environments often are spaces occupied by several virtually represented “real” persons at the same time, they afford the possibility to subject other “real” persons (or “real” persons' avatars, to be precise) to harmful or in other ways immoral actions.

However, it is controversial how virtual actions are to be morally evaluated, first and foremost because they do not cause any physical harm to people nor do they damage property. Apparently, many of our basic moral intuitions as well as sophisticated normative frameworks are not directly applicable to virtual environments if not accordingly modified. One modification could be to drop the notion of physical harm and focus exclusively on psychological harm instead (Ramirez & LaBarge, 2018). Another rather restrictive approach suggests to conceive of virtual actions as expressions of our character. According to this view, committing immoral actions in a virtual realm constitutes a harm to oneself, as it expresses a non-virtuous character (McGormick,

¹ While the use of VR in science or in professional sectors would be subject to standards of ethical scientific practice or work safety regulations, in European and northern American countries there is next to no regulation on the contents and presentation of videogames.

✉ Clemens Uhing
uhing@iwe.uni-bonn.de

Bert Heinrichs
heinrichs@iwe.uni-bonn.de

¹ Institute for Science and Ethics, University of Bonn, Bonn, Germany

² Institute of Neuroscience and Medicine: Brain and Behaviour (INM-7), Forschungszentrum Jülich, Jülich, Germany

2001). A slightly different position was recently proposed by van den Woldenberg (2024), who holds that immoral virtual actions are only problematic if the agent desires them independent of their virtuality, i.e., desires to do them, no matter if in the real or the virtual. It is also possible to endorse virtual realism, i.e. to claim that virtual environments including virtual actions and avatars are *just as real* as physical environments, actions and people. If one adopts this viewpoint, then there would be no objection to applying established normative frameworks. Virtual harm, then, is *just as real* as physical harm (Chalmers, 2023).

In this paper we want to propose a different solution. We will argue that entering virtual environments usually implies consent to the rules of this environment which, in turn, neutralizes the seemingly immoral character of many virtual actions. We start out by giving a definition of VR and sketching related ontological debates in Sect. 2. Section 3 introduces a case and gives a short overview of four types of potentially immoral virtual actions. In Sect. 4, we briefly present the concept of consent and discuss the limits of valid consent. Section 5 then applies the framework of consent to the case from Sect. 3 and shows how it may neutralize the immoral character of virtual actions, effectively rendering many cases of seemingly problematic behavior unproblematic. In Sect. 6, we subsequently identify virtual actions which cannot be morally “neutralized” by consent from the limits of consent discussed on the basis of relevant literature. Section 7 deals with some objections to our line of argument and introduces some difficult cases. Finally, Sect. 8 summarizes the key findings of our paper and points out some limitations. It should be noted that our argument deals primarily with the ethics of virtual environments in which several real people interact via virtual representations. Its applicability to single-user virtual environments is therefore limited. In Sect. 8 we will, however, try to sketch some tentative ideas on how a transfer of our framework might work.

Definition and ontology of VR

VR is constituted by “immersive, interactive, computer-generated environment[s]” (Chalmers, 2017, p. 312). Immersion means that environments are inclusive, extensive, surrounding and vivid (Slater & Wilbur, 1997). Interactivity refers to the user’s ability to navigate the environment (by changing his position or field of view) as well as to select and manipulate virtual objects (Muhanna, 2015). Interactivity facilitates virtual action. Computer-generation indicates that the virtual objects and events users perceive are physically manifested as data-structures on data storage devices such as hard drives and then rendered on computing units

like Graphics Processing Units (GPU) by rules set down in software.

Although a clear cut distinction between VR and virtual environments realized on two-dimensional screens is hard to make, the use of a Head Mounted Display (HMD) to present a virtual environment is often taken as a sufficient condition to categorize it as a virtual reality (Roettl & Terlutter, 2018). From a psychological perspective, VR is assumed to have the property of inducing a “sense of presence” (Doerner et al., 2022, p. 17) and a “sense of embodiment” (Kiltner et al., 2012, p. 375) in users, making them feel as if they were at another place than they physically are (Witmer & Singer, 1998). Consequently, actions in VR might seem to the user, as if he or she was *really* performing them or *really* subjected to them. Whether actions in VR are evaluated as being real actions or nothing but virtual and thereby somewhat fictional actions depends on the ontological stance one takes towards VR and virtual environments more generally. The most prominent positions are virtual realism and virtual fictionalism.² Virtual realists hold that virtual objects *really* exist and virtual events (and among them actions) *really* take place (Chalmers, 2017, p. 310), namely in a genuinely real virtual reality. Virtual fictionalists, on the other hand, contest the claim of a genuine realness of virtual objects, events and environments and claim that “[...] *VR has fictional aspects all the way down*”³ (Juul, 2019, p. 336).

Considerations on the morality or immorality of virtual actions depend on the ontological stance taken, too. To give only one example: Stealing a virtual object in VR would be more morally problematic, if we held that the respective object was a genuinely real object or at least linked to real values, like money or labor. By contrast, if the virtual object was only a fictional object, stealing it would be equally (un)problematic as stealing someone’s monopoly money.

In this paper, we take an agnostic stance towards the ontological status of virtual reality. We hope to show, with the normative framework proposed here, that it is well possible to evaluate some of the ethical challenges posed by virtual reality without being committed to either virtual realism, fictionalism, or some third position. Moreover, it is subject to further discussion whether the use of VR-technologies warrants an ontological distinction between virtual environments presented on two-dimensional screens and virtual worlds entered via HMDs. For the purpose of this paper, we only assume that VR-environments have no ontological status distinct from screen-based virtual environments. In other words, differences are taken to be purely phenomenal. Screen-based video games are therefore included in our discussion as well as, of course, VR applications. Hereby we

² For a more exhaustive overview of ontological views of virtual worlds see Coeckelbergh (2024).

³ Emphasis from original.

follow Grabarczyk (2024, p. 83), who argues, in our view convincingly, that the use of a HMD instead of a screen doesn't change anything about the virtual environment itself but only about the way we experience it, since it is in both cases the same software.

An overview of potential virtual harm(s)

We will now outline a short scenario which will help to identify potential immoral actions in multi-user VR, which in turn will serve as touchstones for its moral evaluation. In terms of methodology, we build on basic moral convictions, which we assume to align well with the intuitions of many people, as is often done in applied ethics. Of course, this approach is not uncontroversial and some may not share our evaluation in detail. Nevertheless, this approach seems to us to be methodologically reasonable and widely accepted.⁴

Two female friends, Ciri and Lara, play an open-world multiplayer VR survival game. They immerse into the game via a VR-headset and embody two female-looking avatars. The game allows players to build structures such as houses and affords communicating with others through voice chat, cooperation, as well as fighting them. Ciri and Lara spend more than 20 h gathering advanced equipment and building their ideal base. Then a group of other players discovers them. Upon first contact, they refuse cooperation because Ciri and Lara are women who, according to some misogynistic attitudes held by the group, are not considered worthy of being taken seriously in gaming. The group takes virtual actions to attack Ciri and Lara, killing Ciri's avatar and injuring Lara's, who manages to escape to the woods. Afterwards, the hostile group destroys Ciri's and Lara's base and collects the equipment Ciri's avatar lost. They also manage to catch up with Lara's avatar, insult her in generic, non-gender specific but also in gender-specific ways and some of the group members move their avatars' arms so it looks like they were groping Lara's virtual body at intimate spots. Then they go back to the place where the base once stood and leave hate symbols corresponding to Ciri's and Lara's specific religion which the two display in their username to present their faith as part of their (online-) identity.

We can now identify potentially immoral actions from this case of human interaction in virtual environments:

- a) **Violation of avatar bodily integrity and autonomy:** by the act of virtually killing Ciri's or injuring Lara's avatar, the perpetrators violate the bodily integrity of the virtual representations of other human beings. In our case, the consequences in the virtual environment are limited to Ciri's avatar being placed in another part of the world map ("respawn") and Lara's avatar being handicapped, for instance by losing movement speed or by having less health left to handle other dangers. By the virtual groping of Lara's avatar, the offenders disrespect the boundaries Lara would wish to see respected to protect her avatar's bodily autonomy and potentially act in the virtual environment to simulate harassing Lara in the physical world. While the violation of physical integrity and autonomy are two separate issues in the real world, they converge in virtual worlds, as injuries to the virtual body do not cause the user any real pain, but merely impair their options for action in the virtual world and thus their autonomy.
- b) **Violations to virtual property:** by the act of taking the belongings left behind by Ciri's dead avatar and destroying both women's base, the attackers appropriate or eradicate something which has value and meaning to other human beings, which might be an achievement of merit or plainly a resource to sustain an avatar and a comfortable gaming experience.
- c) **Violation of personal identity:** Lara and Ciri both identify themselves as women and religious and also represent themselves this way in the virtual environment. By insulting Lara in misogynist ways and putting up hate symbols corresponding to their religion, the offenders humiliate other human beings and demean the value of their personal identity as composed of e.g. gender, religion or ethnicity.
- d) **Wrongful discrimination:** by not taking women seriously for misogynist reasons, the game's community creates an atmosphere in which women or avatars depicting women are not treated as equals. Therefore, they do not have as many options for cooperation as others and in effect a more difficult time getting virtual resources and security.

It seems that there are widely shared moral intuitions about the moral evaluation of the different actions that are involved in this case. The violation of the avatar's bodily integrity, by killing or injuring them, is generally considered morally permissible. So are the destruction or theft of virtual property and generic insults. One reason for this might be that computer games are often designed around combat, competition, and winning, which can include actions such as killing and stealing. In contrast, actions that seem morally impermissible include virtual groping, violations

⁴ T. Beauchamp and J. Childress (2019) in particular have popularized this approach, building on Rawls' concept of "reflective equilibrium". They speak of "considered judgments" and use these as starting points for ethical considerations. They assume that such considered judgments are most likely to be displayed without distortion. Examples they cite include the wrongness of racial discrimination, religious intolerance, and political repression.

of personal identity and the discriminatory behavior of the attacking group. Although this intuition seems to be widely shared (see for instance in the discussion about the gamer's dilemma) and set down in criminal law of western countries (which do generally not prohibit games with virtual violence, but penalizes hate speech), most authors in video games ethics contradict it. In a Kantian vein of reasoning, some argue that virtual murder necessarily, or at least in some cases, means to treat others as mere means (Flatery, 2021; Kucuk & Yildirim, 2023; Ryland, 2019; Tillson, 2018). According to virtue ethicists, doing wrong in virtual environments is harmful to one's character and implies vicious agency (McCormick, 2001; Sicart, 2005). Van den Woldenberg (2024) follows the virtue ethicists' reasoning and considers the action-motivating desires to be central to their ethical evaluation. Still, he allows for the permissibility of otherwise immoral virtual actions if the agent is motivated by a desire that is limited to doing the action virtually, for instance, to try out a new role. Ostritsch (2017) sees the primary moral responsibility on the developers' side, arguing games should not endorse a morally problematic worldview. Their users would still be morally responsible to correctly judge immoral games as such and disapprove of them. From a Kantian perspective, Ulbricht (2023) argues that it is not the virtual actions as such we should evaluate, but rather the agents' justification, that is, their maxim for the respective action.

In our view, the Kantian theories and the proposals from virtue ethics are too restrictive and not at all compatible with established cultural practices. We do not see a plausible reason why actions in virtual environments should be evaluated stricter than for instance martial arts, pen-and-paper games, particular sexual practices or fictitious stories in books, films or theatre plays. The positions sketched above would therefore entail evaluating many such non-virtual practices as equally impermissible from a moral point of view. Van den Woldenbergs', Ostritschs', and Ulbricht's positions, on the other hand, are nuanced enough as to not fall into double standards, but still have the problem of taking desires, attitudes, and judgments as their main criteria. Referring to facts internal to users of virtual environments – be it desires, judgments, or justifications – makes external ethical evaluation of virtual actions near impossible, since these internal facts are usually inaccessible. Hence, these theories have little import on applied ethics of video games, which should help to evaluate what should be allowed or prohibited in virtual environments.

Against this background, we suggest putting the concept of consent in the center of the ethical evaluation of potentially immoral virtual actions. It allows for a normative view more in line with our intuition on the types of potentially immoral actions described above and established

non-virtual cultural practices. By discussing the limits of consent, we shall demonstrate how discriminatory actions in virtual environments can still be evaluated as being morally impermissible.

The consensual basis of virtual environments

In this section, our aim is to show how focusing on the concept of consent might rid us of some ethical worries which are hard to address by the means of established normative frameworks. We also hope to show how putting consent in the center might allow us to morally evaluate the various actions identified in the case above in a more differentiated way.

For a start, in our society, consent is a very important means to make otherwise immoral actions permissible. Here we understand consent as a morally transformative permission (Hurd, 2018). This is well-known from medicine: a highly aggressive chemotherapy counts as morally permissible or even required, if its predicted medical benefit outweighs its destructive side effects and if the patient consents. Otherwise, it would be a grave violation of bodily integrity and autonomy. In fact, in almost all legal systems treatment without consent counts as assault. But there are also more mundane examples: in the context of sex for instance, consent – at least in the view of many people – makes actions such as restraining, hitting or insulting the partner morally permissible. Tattooing someone is painful for them and technically injures the skin, but if the tattooed person consented, most people would not consider it immoral.

According to Feinberg (1990) harmful actions are usually not wrongful if the harmed person consented to them without coercion. What is disputed though, is how far consent can stretch the boundaries of what is considered to be moral. It is probably uncontroversial that a soft slap in the context of a theatre play is morally unproblematic, while consensual cannibalism seems more dubious.

We won't deep-dive into the ethics of consent, but rather only give a rough sketch of why consent makes a difference in various normative frameworks. According to harm-based accounts, voluntarily giving informed consent to a potentially harmful action might be unproblematic, because by allowing a particular action through our consent, we communicate and prioritize interests which are more important to us than, for instance, not being in any physical pain (cf. Feinberg, 1990). Preference utilitarianism allows us to choose preferences which need not be connected to the maximization of hedonic pleasure. The utilitarian calculus can therefore integrate preferences to experience pain or temporary physical restraints in order to satisfy more

sublime needs and wishes (cf. Hare, 1981). In Kantianism, consent only allows otherwise impermissible actions, if the actions undertaken under consent are embedded in a personal relationship of treating the other as an end and not as a mere means (O'Neill, 1985). Consenting to tattoo someone, for instance, could mean to take their ends – of achieving an aesthetic value in this case – seriously and thereby respect their autonomy to set ends by themselves. Consequently, any action which permanently reduces a human being's autonomy, or is not carried out by the agent because the agent values the subjected person's ends, but for self-interested ulterior motives, is considered ethically problematic even if consented to. The classic example of an act that cannot be validly consented to according to this approach is enslavement, as it permanently eliminates a person's autonomy.

Before we go back to our case, we want to outline some factors which are essential for valid consent. We refer here to principlism, one prominent approach in biomedical ethics, in which the concept of consent has been intensively developed for many years. According to this view consent has to be given by a person who (a) is fully conscious and cognitively capable, (b) not coerced or manipulated and (c) informed of the purpose and consequences of their consent (Beauchamp, 2005). These are the necessary conditions for informed consent. Moreover, it is necessary to give the subject the option to withdraw consent, if it is feasible in the face of the action taken. For example, most surgeries happen under anesthesia and therefore withdrawal of consent mid-action is not possible. This is different for regular MRI-scans where the subject can choose to stop the procedure at any time.

Furthermore, the limits of valid consent might depend on the potentially harmful effects of an action. Some actions have consequences, which are so disadvantageous to a person's well-being or violating their dignity, that the interests of the person to determine their own interests and preferences might be overridden (Owens, 2012). For example, amputating someone's limb, seems not sufficiently justified only by referring to consent, a medical reason seems to be necessary. The boundaries here are certainly difficult to draw and individual cases may be controversial. However, the fact that there are limits to valid consent seems to be widely accepted and is also firmly anchored in many legal systems. Finally, the justificatory strength of consent apparently depends on its breadth. This means that if a person gives consent to being subjected to one single, well specified action, it is easier for the person to assess consequences and be fully informed. On the other hand, if the person gives broad consent telling another person to do whatever they want within an unlimited set of options, it might leave the person who consented in a situation where he or she is

subjected to an action, he or she did not anticipate as being an option.

Consent as a 'game changer' to the ethics of the virtual

With this conception of valid consent, we can elaborate how it should change our ethical evaluation of potentially immoral virtual action types as classified in Sect. 3. To be sure, drawing from consent as a justification for otherwise immoral virtual actions is not a new idea (Beisbart, 2024; Schulzke, 2010). However, existing approaches are rather broad and do not specifically analyze different types of virtual actions. What we aim at is refining the way consent is integrated into the ethics of virtual environments and identifying problematic cases which cannot become moral through consent.

Going back to our story of Ciri and Lara, we now start to re-evaluate violations of avatar bodily integrity and autonomy and violations to virtual property or belongings. Assume both parties knew the virtual environment they joined was a platform for survival and Player-vs-Player (PvP) gameplay.⁵ A disclaimer on the games Steam page reads:

Users have the option to injure or kill each other's avatars. Voice chat allows users to communicate. There is no technical safeguard to in-game property and belongings and can therefore be destroyed or taken by other users.

We think Ciri's and Lara's joining the virtual environment can be taken as consenting, even if only implicitly. Ciri and Lara were informed about the possibility of their avatars being injured or killed and their virtual property destroyed or stolen. From the reference to the voice chat-function they were also able to infer that insults can be exchanged in the virtual environment. Because they know the game is a PvP-survival game, they were also informed about the purpose of virtual representations of immoral actions: enabling competition and creating a feeling of danger. In fact, Ciri and Lara probably take part in this particular virtual environment, precisely because they want their gameplay to be embedded in a dangerous atmosphere. Informed consent is secured, provided Ciri and Lara were both sober and there

⁵ Let us further assume that both parties are of legal age and have full moral and legal capacity. A reviewer correctly pointed out that the ethical evaluation becomes more difficult when considering the involvement of minors or adults which are not fully competent. An adequate discussion of these cases would go beyond the scope of this paper.

was no coercion or manipulation involved. In other words, Ciri and Lara consented to virtually being subjected to and experience certain virtual violations to their avatars and their virtual property. Note that they could also withdraw consent anytime, within reasonable limits. By quitting the game, shutting down their device or taking off their HMDs, they can protect themselves from witnessing violent virtual events. In most cases, such games determine a minimum logout-time during which other users can still interact with the avatar of the player logging out.⁶ However, this is only to serve the purpose Ciri and Lara knew about, i.e. creating a sense of competition and danger. If all players were able to instantaneously withdraw their avatar, the purpose of the game would simply be defeated. Such technical limitations do not force Ciri and Lara to witness the respective virtual events, hence allowing them to ‘look away’ to prevent psychological harm. Potentially harmful effects will, very likely, not be enduring long after exiting the virtual environment. Hence, informed consent also holds with regard to the longevity of effects, because probable harmful effects are foreseeable.

Ciri and Lara give a rather broad consent though. It is only specified that it is possible to be injured and killed, not how this might happen in detail. Ciri and Lara might think of gunfire and knife fights. However, there are games, that allow for more “creative” and potentially more disturbing ways of killing and injuring others. In the non-VR-game *DayZ* for instance (Bohemia Interactive, 2018), players can virtually feed others rotten food (including rotten human flesh) or have them die from blood loss by pumping too much of their virtual blood in blood bags. It is unclear whether the game developers had these uses of game mechanics in mind or whether players just found unintended hidden affordances. If there are such options of immoral and potentially disturbing actions, which are hard for consenting players to foresee, this might render the fulfillment of informed consent questionable. In Ciri’s and Lara’s story, this means the virtual groping can be considered immoral. Sexual harassment was neither a possible virtual action indicated in the disclaimer, nor is it an action necessary to sustain the purpose of the virtual environment. In effect, this type of action is not covered by the two women’s consent.

The ethical evaluation of violations to virtual property or belongings goes similar. Because Ciri and Lara know the purpose of the virtual environment and the pertinent game mechanics, which they have been informed about by the disclaimer, they consented to their virtual property and belongings potentially being destroyed or taken from them.

We can conclude that informed consent renders violations of avatar bodily integrity and autonomy, as well as

violations to virtual property or belongings, morally permissible, at least in certain gaming environments.

There are three restrictions on the first two types of potentially immoral virtual actions. First, as pointed out above, consent should not be too broad. Ideally, all possible virtual representations of immoral and potentially disturbing actions should be possible to anticipate by the users by inferring from the purpose of the virtual environment and pertinent gameplay elements and the disclaimer. Developers should therefore include ethical aspects in their beta-testing and make revisions if immoral options are reported, which were initially not anticipated. Using hidden affordances to come up with ways to virtually act in unforeseeable potentially harmful ways, is therefore morally problematic. Second, breaking the rules and technical limitations of the virtual environment, for instance by exploiting glitches, cheating or hacking, is not covered by consent and therefore morally impermissible. Such actions are neither inferable from the purpose of the virtual environment and gameplay, nor from disclaimers, since disclaimers refer to the rules of a particular technologically limited virtual environment. Developers should continue deploying countermeasures against rule-breaking, e.g., temporary or permanent account bans. Third, acting immoral in the physical world or in another virtual environment, in order to coerce persons to do something in a particular virtual environment, is obviously impermissible, because players only consent to being subjected to a set of virtual immoral actions within this particular environment and not within others⁷.

The limits of consent

After having examined the first two types of potentially immoral virtual actions, it might seem to be straightforward to apply the same reasoning to violations of personal identity and unjust inequality. Assume the disclaimer read:

Users have the option to injure or kill each other’s avatars. Voice chat allows users to communicate, also in a derogatory or discriminatory manner. There is no technical safeguard to in-game property and belongings and can therefore be destroyed or taken by other users. In-game social relations and structures are subject to the communities’ preferences and may be unequal, unjust, segregationist or discriminatory.

Should we think that Ciri and Lara consented to potentially being subjected to discriminatory and humiliating

⁶ See for example the wiki entry of the very popular MMORPG *EVE Online* (2025).

⁷ See for example Wildman and McDonnell (2020) for a discussion of virtual theft done by means of physical threats in the physical world.

communication and unjust treatment within the virtual environment by joining after having read the disclaimer? After all, they were informed about purpose and consequences of joining the game. Moreover, they could withdraw consent at any time by quitting the application. What is more, persistent harmful effects of joining are uncertain. Still, we think that there are some important differences between “plain” virtual murder, injury, or violations of property and violations to personal identity and unjust treatment. We now want to show why the latter actions stay immoral, regardless of individual users giving their consent. Drawing on literature on the limits of consent, we shall demonstrate how the group-specific derogatory and discriminatory cases morally differ from virtual murder, injury, destruction and theft. Although we identified violations of personal identity and wrongful discrimination as being two distinct types of potentially immoral actions, we think the reasons why they are impermissible widely overlap. Therefore, we will discuss both of them together subsequently.

Usually only the person subjected to an action can validly consent to it. By consenting, Lara cannot morally transform harmful actions which affect Ciri. Ciri has to do it herself (Müller, 2018). A standard example in the ethics of consent is adultery (Schaber, 2018). If two people secretly consent to cheat on their respective partners, the partners are still harmed, because they did not consent to waive their right to sexual exclusiveness. Consent to actions, which affect multiple persons, might therefore only render an action non-wrongful in respect to some of them. With this in mind, it is possible to see how discriminatory and unjust behavior in virtual environments is morally impermissible, even if an individual consented to it.

According to the “representational model” of virtual actions proposed by J.-H. Heinrichs, media and virtual actions can represent problematic attitudes just like regular speech acts do, as Heinrichs calls them (Heinrichs, 2021). By saying something sexist to a woman one expresses, for instance, a hostile attitude towards all women (in contrast to using a generic insult). The virtual groping in our scenario can similarly be considered to be an expression of such hostile attitude, because groping as such is often an expression of disrespect for women’s physical boundaries and also a reflection of gender-based power asymmetries. The semantic content of symbols and speech acts can easily be reproduced ontologically in virtual environments. Although they are virtual, such speech acts are qualitatively the same as their “real” counterparts (Brey, 2013). This counts for the misogynist expressions as well as for the hate symbols from our story. Now, such symbols or speech acts are either designed to humiliate people qua their personal identity, which is understood as being flawed from being tied to some “inferior” group, or to demean groups by insulting one

of their members by referring to their group membership. Regardless, the demeaning semantic content of a symbol or a speech act does not only offend the individual it is directed to, but rather the whole group. A hate symbol corresponding to a specific religion, for instance, is not designed to offend one specific member of that religion, but to demean all its members. In other words, the circle of affected people is much wider in such cases. It includes many people, who do not take part in a particular virtual environment. Since they do not partake, they have not given their consent to being subjected to demeaning or humiliating actions mediated by their group membership. Therefore, any virtual action which carries a semantic content that is derogatory towards people outside “the magic circle” (Huizinga, 2016) is morally impermissible, since it does not satisfy the conditions of valid consent. Consenting to violations of personal identity seems too general a consent for hostilities towards the self and others, for being validly consented to. Note that generic insults are different: calling someone an “arse” does not refer to any other person than the person insulted.

One could object to our argument, by denying that persons who do not partake in a particular virtual environment are harmed; they might not notice the discriminatory and humiliating practices in a virtual environment towards a group they belong to. We partly accept this objection. Surely, but unlikely, there might be cases of discrimination and unjust treatment in virtual environments, which are not leaked. To have a more accessible image in mind, consider the pub around the corner. There might be a disclaimer at the door, saying that any form of verbal humiliation and discrimination is accepted and consented to by any person entering. Therefore, everyone who enters and gets discriminated against is not wrongfully harmed. Suppose no one ever tells on what happens in the pub. In effect, no one would be wrongfully harmed by means of their group being discriminated against in a closed and consensual pub-community. The same holds for virtual environments. However, we would like to point out three limits to the scope of this objection.

First, such cases of discrimination, which never extend beyond the group of people who consented, should be very rare. Popular virtual environments would be subject to discussion in gaming communities or reported on in newspapers or games magazines. In 2017 for instance, the online magazine Vice published an article, which described how communities in the popular online survival game Rust erected Hitler counterfeits or reconstructed concentration camps (Schott, 2017). We take from this that the occurrence of wrongful harm by discrimination in virtual environments is never excluded.⁸

⁸ We could include a lengthy discussion on whether the discriminatory behavior itself is the problem or rather the messenger, who

Second, there is a substantial difference in the way the indiscriminate killing of a persons' avatar might be harming, in contrast to how discriminatory and derogatory practices against specific identities do harm. The former is only harmful if the person who controls the avatar identifies with the avatar and values it. Upon entering a virtual environment, which affords killing and injuring avatars, participants could just decide not to care for their avatar. Vulnerability to killing and injury of ones' avatar in a virtual environment is therefore created by the partaking persons themselves. This is different for offenses directed towards the personal identity, e.g. being religious or identifying as female. Obviously, this person could choose not to identify with the respective identity in order to reduce vulnerability. However, this would force the person to give up something which is often quite important and valuable to people. This is to say, the stakes are higher when consenting to being subjected to discrimination as compared to one's avatar being vulnerable. It therefore seems problematic to present people with this choice in the first place.

Third, as discussed by Elise Woodward (2022) in regard to sexual practices, consent can never be conceptualized independent of psychological pressure or social coercion. Explicit consent is a minimum prerequisite for consent to do its morally transformative magic, but it is not sufficient. We can apply this thought to the question of consenting to discrimination in virtual environments. We can imagine that a woman agrees to partake in a virtual environment within which she will be structurally discriminated against. However, since such discriminatory structures might merely be a reproduction of discriminatory structures from the real world, it is plausible to assume, that psychological pressure or social coercion make persons from discriminated groups consent to the same discriminatory practices in virtual environments. This, again, casts doubt on the possibility of valid consent to being discriminated against in a virtual environment.

In our view, these arguments taken together show that typical cases of discrimination, i.e. violations of personal identity, remain morally impermissible in virtual environments even if consent was given by individual participants.⁹

makes it public. This is outside our current scope. However, shooting the messenger in the case of virtual environments, would also mean having to shoot the messenger in adultery cases. Secret adultery would therefore only be wrongful harm, if it became known to the respective cheated partners and the harming person would be whoever told on the infidels. We think most readers would be inclined to agree that such reasoning is at least highly unintuitive.

⁹ Of course, there are more complex cases. What about violating someone's gender identity they chose for themselves, but that is not shared by anyone else? What about offences against groups or identities that are usually not regarded as vulnerable to discriminatory practices (or, to pick up the notion from the literature on discrimination, groups that are not "socially salient"; cf. Lippert-Rasmussen, 2006)?

Inequal treatment of persons in a virtual environment based on them being part of some social group which also exists in the "real" world (as demarked by skin color, gender, sexual orientation, etc.) is problematic, because it would, in most cases, be connected to the use of derogatory speech or symbols as discussed above. It is therefore morally impermissible in a derivative sense: keeping up a segregation by avatar skin color in a virtual environment would in most cases entail using semantic content which conveys violations to personal identities in order to explain virtual discriminatory and unjust practices. Additionally, systematic discriminatory and unjust practices towards virtual representations of members of "real" social groups, exclude members of these groups from embodying their personal identity in the respective virtual environment, except, if they consented to being discriminated against virtually. Therefore, the opportunities of "real" people to choose among a range of virtual environments to partake in, in accordance with their preferred way of behaving in the virtual environment, are limited compared to those of other people, just for the reason of the former being members of a particular group. To go back to our pub image: anyone entering the pub, who is successively discriminated against, is not wrongfully harmed, because they consented to the pub's rules. However, in comparison to their peers who do not belong to a discriminated group, they have one less attractive option when it comes to pub choice, solely based on some morally irrelevant property of theirs.¹⁰ This constitutes unjust treatment. To make our case clearer, consider a more extreme example. Say every pub in the world had a policy of allowing discrimination based on distinctive, but morally irrelevant features of persons, like eye color. Persons who would be affected if they entered a pub, would not be wrongfully harmed directly by the ongoing in the pubs themselves. One wrongful harm, among others, would instead consist in them unjustly not having an attractive option to take part in the cultural practice of drinking in pubs. And to this unjust design of society, they did not consent. Analogously, unjust virtual environments take attractive options to partake in virtual environments from persons who are part of the groups affected by injustice. In the most extreme case, which assumes that all (and not just some) game communities are discriminatory and toxic,

Although interesting, discussing such cases is beyond the limited scope of this paper.

¹⁰ A reviewer objected that this doesn't really matter, since there would be no justifying reason why *anyone* from a minority group should have *any* interest to enter that pub, except if there was no other pub. In contrast, we believe that there might be some such cases: The discriminatory pub might be in a very good location or it might be particularly cheap or have a unique and spectacular interior design. Then, at least some members of a minority group might have some interest to enter the pub.

some people would not have any option to partake in virtual environments without consenting to having their personal identity violated.

We can now make an additional argument why virtual violations of personal identity are immoral independent of the consent everyone who takes part in one particular virtual environment has given: it makes this particular virtual environment less attractive to groups of persons for morally irrelevant reasons. The entire set of virtual environments provided by society is therefore less attractive to members of some ‘real’ social groups as compared to members of other groups. If Ciri and Lara are hunted down for being women in one virtual environment, they have one attractive option less, and this leaves them worse off than non-female persons.

Objections

We argued that persons give valid consent to a set of harmful virtual actions upon entering a virtual environment, if they are cognitively capable and conscious, not manipulated or coerced, and sufficiently informed about possible harmful actions and their purpose within the virtual environments’ setting and gameplay. Virtual actions such as injuring or killing avatars, stealing or destroying virtual property or generic insults in voice chat, are therefore morally transformed by consent to be permissible. This argument does not hold, however, for virtual actions violating personal identities or establishing an unjust treatment of members of “real” social groups. In these cases, persons outside the “magic circle” would usually be harmed by the semantic content of such actions. Since they didn’t consent, they are wrongfully harmed. We provided some additional arguments and then went on to show how violations of personal identities render particular virtual environments less attractive for some groups of people, therefore leading to injustice, which spills over to the real world, by giving some groups fewer or no attractive options to partake in virtual environments. We now want to address two possible objections to our proposal.

From a deontological point of view, participants consenting to being subjected to harmful actions might seem insufficient to make the respective actions morally permissible. The optimistic assumption about consensual actions would be, that A gives consent to action *h*, because action *h* or its outcome is an end α to A. B then does that action *h*, because she takes A and his ends $\alpha, \dots$ to be her own, since she respects A’s autonomy and wants to promote it. B’s action would therefore be motivated by her respecting and sharing A’s ends (O’Neill, 1985), or by a benevolent attitude towards A (Bergelson, 2010), to phrase it differently.

But it is equally – if not more – plausible, that B does action *h* which A consented to, for ends entirely different from A’s ends α . Think of a surgeon who just enjoys inflicting injury on persons and does only care for their consent, in order not to get into legal trouble. His motivation of performing a medically beneficial procedure would seem to us morally flawed because what he actually wants is to inflict harm on persons. The same goes for the hostile group in our case. Under the optimist assumption of consent, they attack Ciri and Lara to give them a thrill and make their gameplay more interesting. Under a pessimist assumption, they attack them in order to psychologically harm Ciri and Lara, to discourage them from playing or to play out a fantasy of wrongfully harming them. One could now argue that Ciri and Lara did not consent to being instrumentalized in that way but only to the purpose of having an exciting and thrilling experience with a community.

We think this objection from the pessimist assumption to be quite powerful. Still, it does not render our argument invalid. Ciri’s and Lara’s end to feel a thrill from danger within the virtual environment is not necessarily best promoted by others taking Ciri’s and Lara’s ends as their own. The uncertainty about virtual threats, that is the uncertainty about whether other participants are hostile or not, partly stems from their motives and ends being obscured and seemingly arbitrary. But it is this uncertainty which renders the gameplay especially thrilling and exciting. Consider a soccer match: if both teams agreed on the common end to have an exciting match, they would in some moments hold back on putting pressure on the other team, as for the score not to get too unbalanced to provide an exciting match. The match would become foreseeable and hence stop being exciting. In the context of games, it is necessary for the ends of A and B to partly diverge. However, there still must be a level on which the ends of Ciri and Lara and other players converge. In order to play morally, everyone needs to share the end of doing justice to the purpose of the game, and for this reason obey the game’s rules and, in particular, refrain from violations to personal identity. We conclude that it is entirely possible for the participants’ ends to diverge on some level, as long as they converge in the agreement to partake in the virtual environment within the limits of consent described above. Therefore, aiming at being the “bad guy” and acting as such is morally permissible, because there is no violation of personal identity or wrongful discrimination connected to it. What the group in the case of Ciri and Lara does is different, however, because they targeted the two women in order to violate their personal identities and push them out of the game – ends not covered by valid consent for reasons explained above.

The second objection is related to ethical positions which take internal facts about players, like desires, judgments,

justifications, or virtues as a criterion for the morality of virtual action (Ess, 2013; Kucuk & Yildirim, 2023; McCormick, 2001; Ulbricht, 2023; van den Woldenberg, 2024). In their extreme form, these positions condemn almost all instances of virtual violence. McCormick (2001), for example holds that committing virtual violence means opting for vice, in consequence harming one's own character. Kucuk and Yildirim (2023) claim acting immoral towards avatars or NPCs means holding a maxim conflicting with Kant's Humanity as an end-formulation, because there is no reliable phenomenological distinction from real persons. These objections can be seen as an extension of our argument against violations of personal identity in virtual environments. The problematic internal facts, like wrong maxims on the treatment of others or a non-virtuous self-conception, could be offensive to everyone outside the "magic circle" or just harmful to one's own character. We do not think, however, that acting violently in virtual environments necessarily requires the agent to adopt a problematic stance towards human beings in general or towards themselves in particular. In our understanding, this is consistent with the inconclusive or contradicting results of studies on the correlation between violent video games and aggressive behavior in the real world.¹¹ We suggest a rather simple explanation: virtual environments bring forward a rich setting to contextualize the player's experiences, avatar, and actions. Players may, for instance, find themselves in a dangerous fantasy-world inhabited by dragons, embodying an orc and killing their spouse on the altar of Mephala (Bethesda Game Studios, 2011). Proponents of these internal-oriented views might now claim that killing the virtual spouse is evidence of the desire to kill any spouse and endorses a problematic world-view, that ritual killings are morally permissible, because they could have chosen virtuous desires or righteous maxims instead but opted for vice or wrongness. However, to pick out the problematized action as being generally (and not only virtually) desired by the player but not assume that the player at the same time desires to be in a dangerous fantasy world instead of being in their home or to be an orc and not a human being seems premature.

¹¹ A meta-analysis by Prescott et al. (2018), for instance, found overall significant effects confirming a link between violent video games and aggression, while a meta-analysis of longitudinal studies by Drummond et al. (2020) concluded studies do not support "substantive long-term links", but found significant effects only because of methodological weaknesses. A reviewer pointed out that empirical studies in this field can reveal multiple correlations without confirming causal links between violent video games and aggressive behavior of players.

Conclusion

In this paper, we examined the role of consent for the ethics of virtual actions in VR and non-VR applications. We started out with the case of Ciri and Lara, who encounter several potentially immoral actions in an open-world multiplayer VR survival game. We distinguished violations of avatar bodily integrity and autonomy, violations of virtual property, violations of personal identity, and wrongful discrimination. The intuition we presented was that the former two are morally permissible, while the latter are not justifiable. We argued for our view by drawing on the concept of consent that is well-established in research ethics, but also from perceived cultural practices such as martial arts or particular sexual practices. Immoral actions in virtual environments, which are necessary to sustain a virtual world's purpose, gameplay, and setting, can generally be rendered morally permissible if all participants consent, at least implicitly, to be affected by such actions.

Violations of personal identity and wrongful discrimination, however, can—according to our reasoning—never be permissible. Discriminatory symbols, speech acts, and practices carry a problematic semantic content, which not only offends the respective participants of a virtual environment, but also persons outside "the magic circle," who didn't consent. Additionally, the potential gravity and scope of violations of personal identity cast doubt on whether anyone can validly consent to them. Moreover, we noted that discriminatory practices in virtual environments might turn into real-world discrimination, since some groups might have fewer or no attractive options to join virtual worlds. Virtual violations of personal identity and wrongful discrimination, therefore, have the potential to aggravate injustice in real societies.

Finally, we discussed and rejected two potential objections to our arguments. According to the first, virtually treating persons immorally means not taking them for ends in themselves, but instrumentalizing them for one's own concept of fun. While this might be true for the relationship we take towards the avatars of other participants, by abiding by the rules of a particular virtual environment and sharing the common aim to have fun with its gameplay and setting, overall, we still take the ends of fellow participants as our own. The second objection comes from positions which take internal facts, such as desires, judgments or justifications about the respective agent, as a criterion for the moral evaluation of virtual actions. We rejected this objection on the grounds of it not being plausible for virtual representations of immoral actions to necessarily be connected to such problematic internal facts. Choosing to virtually act immorally does not imply endorsing such behavior in the real world any more than choosing to play an orc implies

wanting to be an orc in real life. Furthermore, evaluating the morality of virtual actions along internal facts about the agent makes it impossible to give a reliable interpretation in the absence of a truthful testimony.

One limitation of the framework proposed in the present paper is that it is designed to morally evaluate virtual action in multi-user virtual environments. In single-user virtual environments, there is no consent between avatars. However, our arguments for the immorality of violations of personal identity and discriminatory behavior may still hold if the respective world is either designed to promote these actions or if these actions become public (e.g., by a stream). If this is not the case, i.e. if no one except the agent experiences them, we tend to concede that such actions would be morally unproblematic. However, this does not imply that the underlying attitude itself is morally unproblematic too.

Sharing virtual environments with other persons means taking up a responsibility to respect the boundaries of valid consent, but also to create an exciting experience for everyone by taking the role of the “bad guys” from time to time and virtually acting immorally. No matter if good or bad—virtual environments are spaces between consenting avatars.

Acknowledgements The authors thank Dr. Tommaso Bruni for reviewing the first draft of the manuscript and Dr. Tijs Vandemeulebroecke for reviewing the manuscript before submission.

Author contributions B.H. and C.U. conceptualized the paper. B.H. and C.U. wrote the manuscript. B.H. and C.U. reviewed and revised the manuscript before submission. All authors have read and agreed to the submitted version of the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL. This research has been funded by the state of North-Rhine-Westphalia as part of INVIRTUO under grant no. PB22-063A.

Data availability No datasets were generated or analysed during the current study.

Declarations

Competing interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Beauchamp, T. L. (2005). Who deserves autonomy, and whose autonomy deserves respect? In J. S. Taylor (Ed.), *Personal Autonomy: New Essays on Personal Autonomy and its Role in Contemporary Moral Philosophy* (pp. 310–329). Cambridge University Press. <https://doi.org/10.1017/CBO9780511614194.015>
- Beauchamp, T., & Childress, J. (2019). Principles of biomedical ethics: Marking its fortieth anniversary. *The American Journal of Bioethics*, 19(11), 9–12. <https://doi.org/10.1080/15265161.2019.1665402>
- Beisbart, C. (2024). *Was heißt hier noch real? Wirklichkeiten in Zeiten von Computersimulationen und virtueller Realität*. Reclam.
- Bergelson, V. (2010). Consent to Harm. In F. G. Miller & A. Wertheimer (Eds.), *The ethics of consent: Theory and practice* (pp. 163–192). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195335149.003.0007>
- Bethesda Game Studios. (2011). *The Elder Scrolls V: Skyrim [Computer software]*. Bethesda Softworks.
- Bohemia Interactive (2018). *DayZ [Computer software]*.
- Brey, P. (2013). The Physical and Social Reality of Virtual Worlds. In M. Grimshaw-Aagaard (Ed.), *The Oxford handbook of virtuality* (pp. 42–54). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199826162.013.029>
- Chalmers, D. J. (2017). The virtual and the real. *Disputatio*, 9(46), 309–352. <https://doi.org/10.1515/disp-2017-0009>
- Chalmers, D. J. (2023). *Reality+: Virtual worlds and the problems of philosophy*. WW Norton & Co.
- Coeckelbergh, M. (2024). All too real metacapitalism: towards a non-dualist political ontology of metaverse. *Ethics and Information Technology*, 26(2). <https://doi.org/10.1007/s10676-024-09768-4>
- Doerner, R., Broll, W., Jung, B., Grimm, P., Göbel, M., & Kruse, R. (2022). Introduction to virtual and augmented reality. In R. Doerner, W. Broll, P. Grimm, & B. Jung (Eds.), *Virtual and Augmented Reality (VR/AR): Foundations and Methods of Extended Realities (XR)* (pp. 1–37). Springer. https://doi.org/10.1007/978-3-030-79062-2_1
- Drummond, A., Sauer, J. D., & Ferguson, C. J. (2020). Do longitudinal studies support long-term relationships between aggressive game play and youth aggressive behaviour? A meta-analytic examination. *Royal Society Open Science*, 7(7), 200373. <https://doi.org/10.1098/rsos.200373>
- Ess, C. M. (2013). Ethics at the Boundaries of the Virtual. In M. Grimshaw-Aagaard (Ed.), *The Oxford handbook of virtuality* (pp. 683–697). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199826162.013.009>
- EVE Online (2025, June 17). *Weapon and Logoff Timers*. <https://support.eveonline.com/hc/en-us/articles/203208932-Weapon-and-Logoff-Timers>
- Feinberg, J. (1990). *Harm to others* (4. print). *The moral limits of the criminal law / Joel Feinberg: Vol. 1*. Oxford University Press.
- Flattery, T. (2021). May Kantians commit virtual killings that affect no other persons? *Ethics and Information Technology*, 23(4), 751–762. <https://doi.org/10.1007/s10676-021-09612-z>
- Grabarczyk, P. (2024). What is real in virtual reality? *Roczniki Filozoficzne / Annales De Philosophie / Annals of Philosophy*, 72(1), 79–98. <https://doi.org/10.18290/rf24721.5>
- Hare, R. (1981). *Moral thinking: Its levels, method, and point*. Clarendon. <https://academic.oup.com/book/12440> <https://doi.org/10.1093/0198246609.001.0001>
- Heinrichs, J. H. (2021). Virtual action. *Ethics and Information Technology*, 23(3), 317–330. <https://doi.org/10.1007/s10676-020-09574-8>
- Huizinga, J. (2016). *Homo ludens: A study of the play-element in culture*. Angelico.

- Hurd, H. M. (2018). The normative force of consent. In A. Müller, & P. Schaber (Eds.), *Routledge handbooks in applied ethics. The Routledge handbook of the ethics of consent* (pp. 44–54). Routledge Taylor & Francis Group.
- Juul, J. (2019). Virtual reality: Fictional all the way down (and that's OK). *Disputatio*, 11(55), 333–343. <https://doi.org/10.2478/disp-2019-0010>
- Kiltner, K., Groten, R., & Slater, M. (2012). The sense of embodiment in virtual reality. *Presence: Teleoperators and Virtual Environments*, 21(4), 373–387. https://doi.org/10.1162/PRES_a_00124
- Kucuk, E. E., & Yildirim, C. (2023). Is categorical imperative meta-versal? A kantian ethical framework for social virtual reality. In A. Schmidt (Ed.), *ACM Digital Library, Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–7). Association for Computing Machinery. <https://doi.org/10.1145/3544549.3585911>
- Lippert-Rasmussen, K. (2006). The badness of discrimination. *Ethical Theory and Moral Practice*, 9(2), 167–185. <https://doi.org/10.1007/s10677-006-9014-x>
- McCormick, M. (2001). Is it wrong to play violent video games? *Ethics and Information Technology*, 3(4), 277–287. <https://doi.org/10.1023/A:1013802119431>
- Muhanna, M. A. (2015). Virtual reality and the CAVE: Taxonomy, interaction challenges and research directions. *Journal of King Saud University - Computer and Information Sciences*, 27(3), 344–361. <https://doi.org/10.1016/j.jksuci.2014.03.023>
- Müller, A. (2018). Moral obligations and consent. In A. Müller, & P. Schaber (Eds.), *Routledge handbooks in applied ethics. The Routledge handbook of the ethics of consent* (pp. 119–130). Routledge Taylor & Francis Group.
- O'Neill, O. (1985). Between consenting adults. *Philosophy & Public Affairs*, 14(3), 252–277.
- Ostritsch, S. (2017). The amoralist challenge to gaming and the gamer's moral obligation. *Ethics and Information Technology*, 19(2), 117–128. <https://doi.org/10.1007/s10676-017-9420-x>
- Owens, D. (2012). *Shaping the normative landscape*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199691500.001.0001>
- Prescott, A. T., Sargent, J. D., & Hull, J. G. (2018). Metaanalysis of the relationship between violent video game play and physical aggression over time. *Proceedings of the National Academy of Sciences of the United States of America*, 115(40), 9882–9888. <https://doi.org/10.1073/pnas.1611617114>
- Ramirez, E. J., & LaBarge, S. (2018). Real moral problems in the use of virtual reality. *Ethics and Information Technology*, 20(4), 249–263. <https://doi.org/10.1007/s10676-018-9473-5>
- Roettl, J., & Terlutter, R. (2018). The same video game in 2D, 3D or virtual reality - How does technology impact game evaluation and brand placements? *PLOS ONE*, 13(7), 1–24. <https://doi.org/10.1371/journal.pone.0200724>
- Ryland, H. (2019). Getting away with murder: Why virtual murder in MMORPGs can be wrong on Kantian grounds. *Ethics and Information Technology*, 21(2), 105–115. <https://doi.org/10.1007/s10676-019-09498-y>
- Schaber, P. (2018). Consent and wronging a person. In A. Müller, & P. Schaber (Eds.), *Routledge handbooks in applied ethics. The Routledge handbook of the ethics of consent* (pp. 55–64). Routledge Taylor & Francis Group.
- Schott, D. (2017, November 16). Menschenjagd und nachgebaute KZs: Unterwegs in der Nazi-Community von Rust. *VICE*. <https://www.vice.com/de/article/menschenjagd-und-nachgebaute-kzs-unterswegs-in-der-nazi-community-von-rust/>
- Schulzke, M. (2010). Defending the morality of violent video games. *Ethics and Information Technology*, 12(2), 127–138. <https://doi.org/10.1007/s10676-010-9222-x>
- Sicart, M. (2005). Game, player, ethics: A virtue ethics approach to computer games. *International Review of Information Ethics*, 4(12), 13–18. <https://doi.org/10.29173/iriel163>
- Slater, M., & Wilbur, S. (1997). A Framework for Immersive Virtual Environments (FIVE): Speculations on the role of presence in virtual environments. *Presence: Teleoperators & Virtual Environments*, (6), 603–616. <https://doi.org/10.1162/pres.1997.6.6.603>
- Tillson, J. (2018). Is it distinctively wrong to simulate doing wrong? *Ethics and Information Technology*, 20(3), 205–217. <https://doi.org/10.1007/s10676-018-9463-7>
- Ulbricht, S. (2023). A Kantian response to the Gamer's Dilemma. *Ethics and Information Technology*, 25(3), 39. <https://doi.org/10.1007/s10676-023-09710-0>
- Van den Woldenberg, D. A. (2024). Vices in Gaming: Virtue ethics and endorsement view. *The Journal of Value Inquiry*, 58(3), 425–441. <https://doi.org/10.1007/s10790-022-09897-3>
- Wildman, N., & McDonnell, N. (2020). The puzzle of virtual theft. *Analysis*, 80(3), 493–499. <https://doi.org/10.1093/analys/anaa005>
- Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 7(3), 225–240. <https://doi.org/10.1162/105474698565686>
- Woodard, E. (2022). Bad sex and consent. In D. Boonin (Ed.), *The Palgrave Handbook of Sexual Ethics* (1st ed. 2022, pp. 301–324). Springer International Publishing; Imprint Palgrave Macmillan. https://doi.org/10.1007/978-3-030-87786-6_18

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