

The Ethics of Virtual Sexual Assault

By John Danaher

(draft copy of a forthcoming chapter in *Oxford Handbook of Digital Ethics*,
edited by Carissa Véliz)

Abstract: This chapter addresses the growing problem of unwanted sexual interactions in virtual environments. It reviews the available evidence regarding the prevalence and severity of this problem. It then argues that due to the potential harms of such interactions, as well as their nonconsensual nature, there is a good *prima facie* argument for viewing them as serious moral wrongs. Does this *prima facie* argument hold up to scrutiny? After considering three major objections – the ‘it’s not real’ objection; the ‘it’s just a game’ objection; and the ‘unrestricted consent’ objection – this chapter argues that it does. The chapter closes by reviewing some of the policy options available to us in addressing the problem of virtual sexual assault.

Keywords: sexual assault; virtual reality; harm; consent; sexual autonomy; criminal law.

1. Introduction

Roblox is an online gaming platform — or ‘imagination platform’ according to its developers — that allows users to create their own games. Players using the platform interact via onscreen avatars with a cartoonish, distinctively ‘blocky’ appearance. Roblox has more than 100 million actively monthly users (Alexander 2019) and its stated mission is to ‘bring the world together through play’.¹

¹ ‘The Story of Roblox’, video available on YouTube <https://youtu.be/VL6rYNmfrjM>

In June of 2018, Amber Petersen was reading to her seven year-old daughter, while the latter played Roblox on her iPad. At one point, Amber's daughter interrupted her to show her what was happening in her game. Amber was shocked by what she saw. In a Facebook post describing the incident (with some accompanying screenshots), Amber claimed that her daughter's Roblox-avatar was 'violently gang-raped' by two male avatars.² They grabbed her daughter's avatar and forced it to engage in simulated sex acts in a virtual playground. Amber had taken precautions to limit the number of features her daughter could enable on Roblox, but the platform does feature some 'adult' content, and her daughter was unfortunately the victim of a malicious user. The makers of Roblox reacted quickly and banned the user permanently from the platform.

Shocking though it may, the Roblox 'gang rape' is not unprecedented. Unwanted sexual interactions are now widespread in digital and virtual environments. Some of these interactions take place in open-ended user-generated environments such as Roblox or Facebook Spaces; some take place within specific games, such as World of Warcraft, Drunkn Bar Fight or QuiVR (Cortese 2019; Cortese and Zeller 2020). The frequency of such incidents raises a number of important questions. Are these interactions unethical or do the ordinary rules of ethics not apply in virtual environments? If they are unethical then how unethical are they? Are they equivalent to real-world incidents of sexual assault or are they somehow less serious? If they are less serious, how exactly should they be classified and understood?

² The original Facebook post recounting the incident is no longer available but reports of Amber's Facebook post are available at: <https://globalnews.ca/news/4316449/roblox-gang-rape-7-year-old-girl/>; <https://www.bbc.com/news/technology-44697788>; and <https://knowyourmeme.com/memes/events/roblox-gang-rape>. It's not clear what the "gang-rape" consisted of, but from the photos accompanying Amber's post, it appears to have involved one of the male avatars pushing her daughter's avatar's head down toward his crotch/groin and holding it there in order to simulate oral sex. The avatars appear to have been clothed throughout the simulated gang rape.

In this chapter, I take up these questions and try to offer some answers. I do so in four main stages. First, I review some examples of unwanted sexual interactions in virtual environments, consider how widespread the phenomenon might be, and address how best to define and classify them. Second, I offer an initial argument for viewing such incidents as serious moral wrongs. This argument focuses on both the unwanted and harmful nature of such incidents. More specifically, it focuses on how such interactions might violate someone's right to sexual autonomy and/or cause them serious harm. Third, I defend this argument from three major objections: the 'it's not real' objection; the 'it's just a game' objection; and the 'unrestricted consent' objection. Fourth, and finally, I consider how, if this argument is correct, we should classify such incidents. More precisely, I review two possible approaches to classifying them: (i) treating them as equivalent to real world sexual assault or (ii) adopting a graduated scale of seriousness in which some might be equivalent to real world sexual assault and some are less serious, depending on the precise nature of the virtual interaction. I wrap up the chapter by considering some potential solutions to the problem of unwanted virtual sexual interactions.

2. The Prevalence of Virtual Sexual Assault

How prevalent are unwanted sexual interactions in virtual environments? There are two main sources of evidence that can help us answer this question: (i) documented incidents and (ii) surveys.

Let's consider the documented incidents first. These are cases in which the victims of such interactions have reported their experiences, usually in some online media platform. The Roblox example is one. There are several other famous ones that are worth mentioning. Reviewing them helps to give a sense of how widespread and diverse the phenomenon is. Here are some of the better known examples:

The Lambda Moo 'Rape': This incident took place in the early 1990s and is possibly the first documented case of virtual 'rape' (I put this in scare quotes for the time being; we will return to the issue of the appropriate name for such incidents later in this chapter). It was discussed in Julian Dibbell's book *My Tiny Life* (1998) and occurred in the text-based virtual world LambdaMoo. As described by Dibbell, the incident involved one character (Mr Bungle) creating a 'Voodoo Doll' program that allowed him to control the avatars of two other female characters and force them to engage in text-based sexual interactions. Dibbell followed up with the women involved and explained that one of them was particularly traumatised by the incident (Dibbell 1998; Strikwerda 2015; Danaher 2018).

The Second Life Rape(s): Second Life is an open-ended virtual world in which people can create 3D-rendered avatars that interact with one another in complex social environments. It was quite popular in the early 2000s and still exists today. Sexual activity is very common on the platform and sometimes includes unwanted sexual activity or virtual 'rape'. In 2007, the Belgian police decided to investigate an alleged 'rape' that took place on the platform. No charges were brought on foot of that incident but similar incidents have been documented by other users of Second Life (Duranske 2007; Lynn 2007).

The QuiVR Virtual Groping: In 2016, the journalist Jordan Belamire (a pseudonym) wrote a widely-reported piece in which she described being virtually groped while playing the VR game QuiVR. This is an archery game played using a head mounted display (i.e. a set of goggles that projects 3D images directly onto the user's retinas) and so, unlike two previous incidents, involves an 'immersive' virtual environment. In her case, Belamire was approached by another player in the game environment. This player proceeded

to reach out and grab at where her breasts would have been in the real world (Belamire 2016).

The Altspace Unwanted Kissing: Altspace is a social VR app that can be accessed using an immersive head mounted display. It's not a game platform; it's simply a place in which people can 'hang out' virtually — like a sophisticated chat room. In 2016, the journalist Taylor Lorenz and her colleague Alexis Kleinman both reported receiving unwanted 'kisses' when they tried to use the app. The unwanted sexual behaviour was subsequently shut down by a moderator. Similar incidents have been reported by other users (Lorenz 2016).

Sexual Harassment in Rec Room: Somewhat similar to Roblox, Rec Room is an immersive VR platform that allows users to create and play various games together. In 2017, the Facebook VR researcher Michelle Cortese was on a VR field trip with her colleagues. They were due to play paintball on Rec Room. While waiting for others to arrive, an unknown male user 'dumped' virtual water bottles on her and screamed 'WET T-SHIRT CONTEST!' while chasing her around a room. Subsequent to this, Cortese did a public art exhibition in which she exposed herself to the risk of repeated sexual overtures and advances in the social app VR Chat and displayed the experiences to members of the public (Cortese 2019). The intention was to give a vivid depiction of the potential frequency of sexual harassment in virtual spaces.

The list could continue.³ Hopefully these cases are enough to get a sense of what can happen on these platforms.

³ Other incidents are documented in detail here: <https://www.adl.org/resources/reports/hate-in-social-virtual-reality#sexism-racism-and-anti-semitism-in-social-vr-today>

What about survey data the prevalence of such incidents? Jessica Outlaw has performed two studies on this topic. In one, conducted with Beth Duckles, she carried out detailed interviews with 13 ‘tech-savvy’ women (Outlaw and Duckles 2017). She found that most of these women reported fear of sexual harassment or unwanted sexual attention as a major barrier to their use of VR platforms. Furthermore, several had actually experienced unwanted sexual attention when using these platform and some of them went out of their way to adopt non-gendered avatars in order to avoid attracting attention. Subsequently, Outlaw conducted a survey of over 600 users of VR platforms and found that 49% of women and 36% of men had experienced at least one incident of sexual harassment (Outlaw 2018). Users also reported other forms of harassment such as racist or homophobic comments. In a similar vein, Sparrow et al (2019) found that players of online multiplayer games had come to expect some level of abuse and harassment when playing these games. Indeed, although they thought this behaviour was wrong, they had adopted an attitude of almost learned-helplessness in response to it (i.e. they believed that they were unable to do anything to control the problem). This is consistent with PEW research findings suggesting that 40% of Americans have been victims of some form of online harassment.⁴

None of this is surprising. Technology has long been used to facilitate sexual interactions, both wanted and unwanted. The prevalence of pornography and cybersex on the internet is common knowledge. Furthermore, sexual harassment and unwanted sexual contact are common in the real world. Consequently, it’s not surprising to find that this bad behaviour should transfer online.

⁴ Details of this research are available here: http://assets.pewresearch.org/wp-content/uploads/sites/14/2017/07/10151519/PI_2017.07.11_Online-Harassment_FINAL.pdf

But how exactly should we understand these incidents of unwanted sexual interactions in digital and virtual environments? How are they best defined? Some are clearly and unproblematically forms of sexual harassment: menacing, unwanted sexual communications or attention directed at specific persons. Although sexual harassment often must occur in a work environment in order to be legally recognised, colloquially it can be understood to occur in other environments and this can now include virtual environments. Are any of the incidents more than sexual harassment? In past work, I have favoured classifying at least some of them — specifically those that involve enacting or representing sexual contact — as incidents of ‘virtual sexual assault’ (Danaher 2018a; Strikwerda 2015). But this might be thought a problematic designation since it implies an equivalency between these incidents and real world sexual assault (a topic discussed towards the end of this chapter). Avoiding this equivalency, Lucy Sparrow (2019) has suggested that we refer to them using the more descriptive and neutral term ‘UDESIS’ (unwanted, digitally-enacted sexual interactions). This proposal has some merit, particularly if our goal is to scientifically study these incidents as a novel social phenomenon. Nevertheless, as I will argue below, it may fail to do justice to the seriousness of these incidents by failing to grant them some level of equivalency with ordinary sexual assault. For now, I will stick to describing them as unwanted virtual sexual interactions, occasionally adopting the terms virtual sexual assault or harassment as seems appropriate.

More important than terminology, however, is understanding how these incidents arise. If you look at the list of examples given you can immediately see that there are some differences between them. Some involve simulated or virtual touching and some do not; some involve users wearing immersive head mounted displays whereas others do not; some involve social platforms and some involve games. Are there any shared properties across all these incidents?

I have a suggestion. The one thing they all seem to share is that instead of involving direct physical contact between people they each involve a sexual interaction performed via some virtual representation (or avatar). It is these avatars that are used to facilitate sexual communication or engage in representations of sexual acts. In a previous article, building upon work done by Jonny Søraker (2009) and Litska Strikwerda (2015), I have argued this *action-via-a-representation* is perhaps the key defining feature of all these incidents.

Beyond that, it may be worth classifying and distinguishing between these incidents along a number of different dimensions. For example, we could classify them depending on the degree of immersion they involve. This might be based on the intuition that there is something *prima facie* different about using a controller to manipulate an on-screen avatar versus actually embodying that avatar through an immersive headset. In the former case there is more separation between what happens to you and what happens to your virtual representation; thus, the potential harm that results from the incident might be reduced. Whether that intuition holds up to closer scrutiny is something that we will need to consider in more detail in what follows.

We could also classify them depending on whether they involve game environments or not. The intuition behind this strategy is that different ethical rules apply to games versus other kinds of interactions. As the Dutch historian Johan Huizinga noted in his famous study of play in human culture, games seem to take place inside a ‘magic circle’ in which the ordinary moral rules are suspended (Huizinga 1949). There is some *prima facie* support for this idea. After all, we tolerate conduct as part of games and sports that we ordinarily would never tolerate outside of those games. For example, conduct that would ordinarily meet the legal definition of physical assault or battery is a normal feature of sports such as boxing and rugby. Whether this makes any difference

to virtual sexual interactions is another topic that we will need to consider in more detail in what follows, but it is worth noting that the idea that video games exist inside a magic circle has been promoted by game designers and theorists (Zimmerman and Salen 2003).

Finally, we could classify them depending on who exactly is the perpetrator and who is the victim of the unwanted sexual interaction. To this point, and throughout all the examples discussed, the assumption has been that both the victims and perpetrators are real human beings that happen to be acting via virtual representations. But, of course, this may not be the case. Some virtual environments feature artificial characters, who operate according to the rules of a computer program. These characters could, potentially, be both victims and perpetrators of virtual sexual assault. Indeed, I will discuss some examples of this below. Considering these different possibilities is interesting because it raises significant questions about the nature of wrongdoing and responsibility. Can an artificial character be held morally responsible for carrying out a virtual sexual assault? Or does the blame rest with the programmer of that character? Should we care about what people do to wholly artificial characters? If they cannot be moral victims of our actions — if they feel nothing and cannot be harmed — then is it permissible to act out all manner of immoral fantasies on them?

In what follows, I will largely (though not completely) ignore the complexities that arise from cases in which artificial characters are the victims or perpetrators of virtual sexual assault. I do so not because I think these issues are unimportant but because they have been the major focus of my past writings on this topic. To briefly summarise my views: I think that at least some forms of sexual assault that are carried out on artificial characters are morally impermissible, particularly if the artificial character represents a child (Danaher

2017; Danaher 2018a; Danaher 2019); I also think that it may be permissible to hold the designers of artificial characters responsible for their virtual actions, if certain conditions are met (Danaher 2018a). That said, I also believe that there could be cases in which there is no suitably culpable moral agent to hold responsible for the actions of artificial characters (Danaher 2016). And I am not the only one to defend these views. For example, Robert Sparrow (2017) and Stephanie Patridge (2010; 2013) have both argued that representations of rape and sexual assault can be morally problematic. There is also a rich and varied debate about the appropriate take on responsibility for artificial, autonomous systems (Matthias 2004; Danaher 2016; Nyholm 2018; Santoni de Sio and van den Hoven 2018).

3. The Moral Wrongness of Virtual Sexual Assault

What is it that makes an action wrong? There are many answers to that question but for present purposes we can consider two common ones corresponding, roughly, with the deontological and consequentialist schools of normative ethics: (i) the action violates someone else's rights and/or (ii) the action causes, or at least has the potential to cause, non-trivial harm to another. Accepting these are the two fundamental criteria of wrongness, a simple argument can be made in favour of the view that unwanted virtual sexual interactions are a serious wrong. As follows:

- (1) If an action violates someone else's rights *and/or* causes (or has the potential to cause) them non-trivial harm, then that action is a serious moral wrong.

(2) Unwanted virtual sexual interactions violate someone's rights (specifically their right to sexual autonomy) and/or have the potential to cause them significant harm.

(3) Therefore, unwanted virtual sexual interactions are a serious moral wrong.

Is this argument any good? I presume the first premise is relatively uncontroversial even though it is complicated. The first premise appeals to both deontological and consequentialist theories of moral wrongdoing. It holds that an action is wrong if it violates someone's rights or if it causes non-trivial harm or has the potential to cause non-trivial harm. In what follows, I will consider how both examples of wrongdoing might arise in cases of virtual sexual assault. There are, of course, questions to be asked about what counts as a 'right' and what counts as 'harm'. Answers to those questions have provoked centuries of philosophical debate. That said, there are paradigmatic cases of both that are uncontroversial. As I hope to make clear in what follows, the application of both concepts to this particular argument would seem to fall within the paradigmatic range of cases, namely: the right to sexual autonomy and the physical and psychological trauma that can result from unwanted sexual interactions.

The one aspect of the first premise that may raise a few eyebrows is the bit stating that if an action 'has the potential to cause non-trivial harm' it counts as a serious moral wrong. Why is this included? The answer lies in how we categorise wrongful conduct in advance of that conduct actually taking place. People differ in how much harm they suffer at the hands of others. Some people are physically more resilient than others; some people are psychologically more resilient. Sometimes we can reliably predict who is likely to suffer more harm as a result of a particular action. For example, punching a child in the face is

likely to result in more harm than punching Mike Tyson (in his prime) in the face. We could factor such predictions into our classification of wrongful actions. But we don't always have the capacity to predict the likely outcome and it would be unethical to run an experimental test of harmfulness for every possible case. This is particularly true in the case of actions that are known to cause psychological harm. There is a wide degree of variability in how psychologically traumatising people find different experiences. Some people brush them off with relative ease; others suffer from years of PTSD and depression. Since we cannot know for sure how victims will react in each and every case, and since we can never fully control the consequences of our actions, we have to work with reasonable estimates of the likely harmfulness of our actions. Consequently it seems legitimate to say that if an action has the potential to result in non-trivial harm it is a serious moral wrong.

What about premise (2)? This might be a little bit more controversial than premise (1) but a strong *prima facie* case can be made in its favour. First, all the examples of unwanted sexual interactions discussed in the previous section would seem to involve a violation of the victim's right to autonomy, in general, and sexual autonomy in particular. None of the people in question appear to have invited or desired sexual attention or interactions. They were either playing non-sexual games or participating in general, non-specific social interactions. Other users of the virtual platforms imposed sexual communications and contact on them, against their will. In modern liberal societies we tend to see liberty or autonomy as a foundational right (Gaus 2010). On some classic accounts of this basic liberty right, this includes a set of claim-rights against others to forbear from interfering with one's actions or choices, providing these actions and choices do not infringe someone else's rights, and unless there is consent to those actions (Wenar 2020; Mack 2018). That said, some violations of autonomy are tolerated more than others. This is inevitable given that it is

impossible to live in a society and have one's autonomy respected at all times. We are, almost inevitably, constantly bumping into and interacting with each other. We don't always consent to these interactions. Nevertheless, most societies place particular emphasis on sexual autonomy as an area of individual choice that deserves the utmost respect. This has not always been true – violations of sexual autonomy for women, in particular, have been historically normalised in many cultures – but it certainly appears to be true nowadays, with increasing criticism and scorn poured on those who violate sexual autonomy. Perhaps this is because the potential harm caused by a violation can be significant, and perhaps it is because it is relatively easy to respect sexual autonomy: unlike casually bumping into someone on an crowded train, you have to go out of your way to violate someone's sexual autonomy. Violations of sexual autonomy are thus treated with particular seriousness and a failure to procure consent prior to a sexual interaction is a serious wrong. In a moment, I will consider whether there is some reason to think that consent can be implied from the participation in certain virtual activities or whether there is reason to be less protective of sexual autonomy in virtual spaces. For the time being, however, I conclude that there is a strong *prima facie* case for thinking that unwanted sexual interactions in virtual spaces violate the right to sexual autonomy because they violate the consent requirement.

Second, unwanted sexual interactions in virtual spaces would seem to have the potential to do serious harm. In the absence of haptics — a type of technology that enables the transmission of touch via, at a distance, via virtual representations — they do not involve physical harm but they do seem to involve psychological harm. How do we know this? Well, we have the reports from the victims themselves. As mentioned previously, Julian Dibbell interviewed one of the victims of the LambdaMoo virtual 'rape' for his article on the topic (Dibbell 1998). She was crying and claimed to have been

traumatised by the event. Consider also the testimony of Michelle Cortese, whose experiences of sexual harassment in VR were also described above. A past victim of sexual violence, Cortese (2019) says that her experiences were triggering of her past trauma:

...I panicked. I innately sensed a familiar knee-jerk reaction to an incoming assault. My heart sank. The air evaporated from my lungs. I froze. I braced.

..... I braced, for nothing.

My body responded to virtual stimuli with real-life survivalism [reference omitted]. Ashamed of my emotional lapse and VR n00b reaction, I kept on playing like everything was fine.⁵

Again, none of this should be particularly surprising. There are now ample case reports and studies done on the effects of cyberbullying and cyberharassment that suggest that they can be every bit as psychologically harmful as their real-world equivalents (e.g. Hamm et al 2015). Since unwanted sexual interactions in virtual environments are similar in nature to both of these actions it is reasonable to suppose that they have the potential to do non-trivial harm.

There is, however, an important question to be asked as to whether harm is a relevant consideration when it comes to understanding sexual wrongs. It is worth noting that when we legally define crimes such as sexual assault and rape rarely we rarely make any reference to the harmfulness of those actions. This is different from our approach to physical assault, which is often legally

⁵ The text comes from this Medium post <https://medium.com/@ellecortese/virtual-healing-bf2b5f0cbf51>

categorised based on the degrees of harm it causes (i.e. assault, assault causing harm, assault causing grievous harm and so on). If we take these legal classifications serious, then it would appear that we care about sexual assault, legally speaking, because it is a violation of sexual autonomy and not necessarily because it is harmful, even though it may be.⁶ As such, this discussion of the harmfulness of virtual sexual assault or harassment might seem to be surplus to requirements. Nevertheless, I include it here for two reasons. First, if it is the case that these interactions have the potential to do non-trivial harm then this provides an additional reason to treat them as serious moral wrongs. In other words, their potential harmfulness bolsters the argument I wish to make. Second, including some analysis of harm helps to head off a potential criticism of the argument, namely that they are less important because they are not ‘real’.

It is to this criticism and others that I now turn.

4. Three Objections to the Moral Wrongness Argument

In this section I shall consider three objections to the argument given in the previous section: (i) the ‘it’s not real’ objection; (ii) the ‘it’s just a game’ objection and (iii) the ‘unrestricted consent’ objection.

4.1 - The ‘It’s Not Real’ Objection

⁶ Definitions and approaches can vary cross jurisdictionally. Some legal definitions of sexual offences could include a reference to harmfulness, and some judges or juries can take harmfulness into consideration when assessing guilt and, more commonly, appropriate punishments. The point here is simply that it is rare for the core sexual offences (rape and sexual assault) to require harm to the victim in order to be legally recognised. They sometimes require force on the part of the perpetrator, but that’s not the same thing as harm to the victim and, again, very different to how physical assault offences are classified.

The first objection is simply that unwanted virtual sexual interactions should not be viewed as serious moral wrongs because they are not real, or at least they are *less real* than other kinds of unwanted sexual interactions, and so don't carry the same moral weight. This objection could be run in two different ways. It could be run as a general objection to the claim that virtual sexual interactions are morally significant. It could also be run as a more nuanced objection to the claim that virtual sexual interactions carry as much weight as their real world equivalents. I take up the latter possibility in the final section of this chapter. For now, I focus on the former, more general, version of the objection.

From the outset, this objection does not seem very plausible. The claim that virtual interactions are not 'real' and so don't carry the same weight as real physical interactions will seem odd to anyone living in the digital age. The mere fact that an interaction takes place via a virtual representation of some kind (text-based or visual) does not make it less real. We live an increasing portion of our lives via virtual avatars. We communicate with colleagues and friends via our profiles on email and social media all the time. But nobody would claim that these communications and interactions are, consequently, unreal or lacking moral weight. The interactions are certainly missing some of the properties we associate with interactions in the physical world. We cannot (yet, anyway) smell one another in virtual fora, and, in the absence of haptics, we cannot touch each other. Nevertheless, we can hear each other and see each other and have real effects on one another's lives. What someone says in a virtual chatroom can really hurt my feelings; what my boss says to me via email can have real effects on my employment duties; and so on. To claim that this whole swathe of our lives is not real or carries no moral weight would be absurd.

A more plausible version of the objection might claim that by adding an extra representational layer to our virtual interactions we make them *less real*.

Consider two people chatting via a videoconferencing app. Presumably most people would accept that their communications are real in some morally important sense. But now suppose that instead of using unfiltered video to communicate they use cartoon avatars. In other words, they don a virtual skin or costume for the duration of the interaction. Someone might argue that donning the virtual skin turns what would be an ordinary and morally weighty interaction into one that is more akin to make believe or fantasy. This may tip over into the ‘it’s just a game’ objection that I will consider in a moment. That said, at an initial glance, and absent some other rules or norms to suggest that ordinary moral rules should be suspended when we communicate via cartoon avatars, it’s not clear why donning the virtual skin makes the interaction any less morally weighty.

It is worth underscoring this point by considering the empirical work that has been done on the blurred boundaries between what happens to digitally encoded, virtual avatars and what happens to real, physical people. In recent years, a spate of research studies have been carried out on how people relate to their virtual avatars and how what happens to those avatars can have serious and lasting effects on those people. For example, research has revealed that people can have significant ‘self-presence’ with their virtual avatars. In other words, they can come to associate themselves with their avatars in a strong way, viewing their avatars as a core part of their identity. This can occur with non-immersive VR avatars, such as those found in Second Life, and as well as with more immersive forms using head mounted displays. For example, research performed by Elizabeth Behm-Morawitz (2013) on 279 Second Life users found that the degree of presence felt toward the virtual self had an effect on offline health, appearance, and well-being. Similarly, research by Nick Yee and Jeremy Bailenson, on what they call the *Proteus Effect*, suggests that people conform to the expectations of their virtual avatars (Yee and Bailenson 2007). For example,

people embodied in taller virtual avatars tended to be more aggressive than those embodied in shorter avatars. This basic research paradigm has been followed by other researchers who have found that the appearance and behaviour of virtual avatars can have lasting effects on real-world, offline behaviour. Those who embody elderly avatars have been found to save more for retirement (Hershfield et al 2011) and those who embody superhero avatars tend to be more altruistic (Rosenberg et al 2013). Furthermore, in a wide ranging review of the empirical literature, Slater and Sanchez-Vives (2016) document how VR is now being used to train offline behaviour —e.g. using VR as a surgical training tool in medicine has become commonplace — and to treat psychological disorders — e.g. people now use VR exposure therapy for victims of PTSD. All this evidence points to an intimate connection between the virtual world and reality, and shows how actions in the virtual world can have real ‘extravirtual’ effects (Søraker 2009).

In addition, and perhaps most crucially from the present perspective, there is considerable evidence to suggest that people in virtual reality environments can experience what is happening to them as ‘real’ and not simply as some form of make believe. This feeling that a virtual experience is real is possible because the human brain already constructs a virtual model of the real world based on a handful of perceptual cues (Slater and Sanchez-Vives 2016). Virtual reality environments, artificial and unrealistic though they may sometimes seem to be, can exploit this feature of human cognition and trick us into thinking that unusual things are really happening to us in virtual environments. Some of the experimental results are quite dramatic. For instance, research by Blanke and Metzinger (2009 - see also Metzinger 2009) has shown that it is possible to trick people into thinking they are occupying a different body (and, more generally, to have a virtual ‘out of body experience’). Similarly, a series of so-called ‘Pit Studies’, in which subjects stand on a platform raised just over an inch off the

floor while wearing a head mounted display that projects the image of standing on the edge of a pit or canyon into their eyes, has shown that people can be convinced that they are in danger of falling a long way in a VR environment. The subjects know that they are in no danger but experience increased physical stress nonetheless (Meehan et al 2002; Madary and Metzinger 2016). Again, these studies suggest that experiences in virtual reality are not unreal and have the potential of causing non-trivial harm to their subjects.

It's important that this point is not misunderstood. I am not claiming that there are no differences between virtual interactions and interactions in the physical world. There are. The aforementioned absence of physical contact is one of them and could have important implications for how we classify or understand the magnitude of the wrong inherent in unwanted virtual sexual interactions. This is one of the issues I take up in more detail in section 4 of this chapter. Similarly, virtual interactions could involve interactions with artificial characters who might be thought to lack the capacity to be moral victims. This could also make a difference to the moral character of any virtual interactions in which they might be involved. Nevertheless, none of these differences should be taken to imply that virtual interactions, as a whole, carry no moral weight. They do and they can, particularly when they involve humans interacting via virtual avatars. To dismiss their moral salience because they are 'not real' or 'less real' is not persuasive.

4.2 - The 'It's Just a Game' Objection

Related to the previous objection is the 'it's just a game' objection. This objection builds on the aforementioned idea that virtual interactions are a species of fantasy role play or game and adds the claim that the ordinary rules of morality do not apply as a result. We have already encountered some reasons to reject this objection but it's worth giving it a fair hearing. There is, after all,

something to it. Many games do involve the suspension of ordinary moral rules (Huizinga 1949; Zimmerman and Salen 2003). I earlier gave the example of sports such as boxing and rugby. These sports sanction what would ordinarily be called physical assault. There are many other examples. For instance, some board games and games of chance, like Poker and Diplomacy, encourage players to engage in forms of deception and bluffing that would ordinarily be deemed immoral. Doing so is part of the fun of the game.

The notion that games represent an amoral space is a view that is particularly compelling when we turn to the world of videogames. There are now all manner of games that encourage players to engage in simulations of actions that, if they occurred in the real world, would be deemed serious moral wrongs. For instance, games like *Hitman* or *Assassin's Creed* or *Grand Theft Auto* all routinely encourage acts of extreme violence, criminality and cruelty as part of their game play (Luck 2009; Luck and Ellerby 2013; Gooskens 2010; Patridge 2013; Ostritsch 2017; Tillson 2018).

Furthermore, the suspension of ordinary moral rules is not just something that applies to games but that also applies to all manner of fictional activities as well. Many plays and movies depict acts of physical and sexual violence. Are they immoral for doing so? There is a debate about this but the reality is that we often treat fiction that depicts immoral acts with great respect. For example, we regularly reward actors who take on the parts of serial killers, psychopaths, and dictators. Think about Joaquin Phoenix in the *Joker*, Daniel Day-Lewis in *There Will be Blood*, Charlize Theron in *Monster* and Anthony Hopkins in *Silence of the Lambs*. Each of them won an Oscar for depicting a psychopath or serial killer. We find their depiction of these characters psychologically and emotionally compelling and we don't think they do anything morally wrong by pretending to be serial killers or sexual sadists (Gooskens 2010).

Could something similar not be true of actions in virtual environments? Could they all be deemed a kind of game play or make believe? There are three points to made in response to this objection. First, not all unwanted virtual sexual interactions take place in game-like environments or occur as part of some fantasy/make believe. Some of the examples given earlier did involve games, for example the archery game in QuiVR and the paintballing game in Rec Room, but in neither of these cases was the unsolicited sexual interaction part of the underlying structure or purpose of the game. It was something added-on by malicious users of the game platform. So even if we did accept that game-worlds have their own internal ethics, unsolicited sexual interactions would not always be included among the things that ought be tolerated by that internal ethics. Most of the other examples involved open-ended socialising in VR. While there may be an element of fantasy and make-believe to these forms of socialising, and while some online platforms have their own community norms and standards, those norms typically do not tolerate unsolicited sexual interactions, nor are they always transparent to users. Certainly, in the examples cited above, none of the users of the relevant platforms thought that unsolicited sexual interactions were acceptable, even if they did anticipate them to some extent.

Second, even if these interactions did occur within environments in which unsolicited sexual interactions were encouraged, tolerated, or part of the underlying game structure, it's not clear that this would save them from moral reproach. Although we do sometimes tolerate the suspension of ordinary moral rules within games, there are limits to how much we are willing to do this. Extreme or risky violence is rarely tolerated in games that permit rough physical contact. You cannot deliberately snap someone's neck in a game of rugby. And knowledge about dangerous, post-game effects of such physical contact can

alter our moral attitude towards games. Fears of permanent brain damage arising from repeated concussion, and the social costs this has both to the affected individuals and their families, have, for example, changed some people's attitudes toward sports such as boxing and American football in recent years. This has, in turn, affected some of the internal rules of such sports (heightened safety protocols and equipment). That virtual interactions can have harmful and lasting real-world effects (as documented above) may give us reason to think that we shouldn't tolerate any suspension of the ordinary moral rules when it comes to unwanted sexual interactions.

It is important to bear in mind that this point is not the same as the typical 'effects' argument that is made in relation to exposure to violent or sexually explicit media. I am not claiming here simply that acting out a sexually violent fantasy can have negative downstream effects for our interactions with other people. As I have suggested elsewhere, I am sceptical about our capacity to adequately research such effects (Danaher 2017b; Danaher 2019). What I am arguing instead is that acting out these fantasies can be *immediately* and *intrinsically* harmful, in the sense that they can corrupt or undermine one's moral psychology and character, that these harms are not obviously compensated for by other gains, and that they *may* have other negative repercussions. Others have made this point too, suggesting that video games can be a training ground for the moral virtues and game designers should pay attention to these possibilities when creating the game structure (Sicart 2009).

Third and finally, immoral sexual interactions within games or other fantasy realms are not as easily tolerated as other kinds of immorality within games and fantasy realms. This topic has been the subject of much debate in the literature on the ethics of videogames. Stephanie Patridge (2010; 2013), for example, has argued that a game player who enjoys and tolerates the immoral sexual actions

depicted in games such as *Custer's Revenge* and *Rapelay* — both games that encourage users to engage in represented acts of sexual assault and violence — is displaying a troubling insensitivity to the moral meaning of such representations. Her view is that we are, consequently, within our rights to question their moral character. This is an argument that I and others such as Robert Sparrow have applied to the adjacent debate concerning acts of sexual assault and violence toward sex robots (Danaher 2017a; Robert Sparrow 2017 — and see also Sterri and Earp, chapter X of this volume for more).

The philosopher Morgan Luck (2009) has also noted the odd inconsistency in our attitudes toward sexual immorality in video games and crafted it into something he calls the 'Gamers Dilemma'. According to this dilemma, most people seem to think it is immoral to depict and play a game involving acts of virtual child sexual abuse, but most people are tolerant of acts of virtual murder and physical torture. Luck claims that it is very difficult to reconcile these two attitudes. You should either tolerate both or accept both. Luck's formulation of the Gamer's Dilemma has generated a small cottage industry of replies since it was first published (Luck and Ellerby 2013; Patridge 2013). The details of those replies need not detain us here. What is noteworthy about them, however, is that most contributors to the debate accept the starting position that representing and engaging in acts of virtual child sexual abuse is immoral. This stance is backed up by studies of users of virtual social platforms such as Second Life, who generally see sexual 'ageplay' (when adult users of the platform play at being children engaging in sexual acts) as something that transgresses or lies at the extreme edge of what is acceptable on the platform (Reeves 2018). This supports the notion that there is something different about virtual sexual immorality in most people's minds. Furthermore, it is worth noting that this debate about virtual sexual immorality usually concerns actions carried out by game players against wholly artificial characters, not human-controlled avatars.

Whatever objections people might have to those actions would seem to apply *a fortiori* to cases involving human-controlled victims.

The typical counterexample to the suggestion that role-playing sexual violence or assault is morally problematic is, of course, the world of BDSM. To outsiders, practitioners of BDSM appear to tolerate and encourage sexual violence and to endorse fantasies of domination and submission that run contrary to our norms of consent. But this is not a true counterexample since the BDSM community typically adopts a complex ethical code grounded in explicit consent and a clear statement of boundaries (Nielsen 2010). People who step outside those boundaries or ignore consent are doing something immoral even within the community norms of BDSM. This is very different from tolerating unsolicited sexual interactions in a virtual environment.

In sum, as Ostritsch notes in his discussion of the amorality of games and fiction, ‘amoralist reasoning...implausibly severs all ties between the fictional and the non-fictional’ (Ostritsch 2017, 122). It does not make sense to sever such ties with the virtual world, particularly given that more and more of us are living more and more of our lives through virtual avatars.

4.3 - The ‘Unrestricted Consent’ Objection

One final objection to the claim that unwanted virtual sexual interactions are serious moral wrongs concerns the possibility that people consent to those interactions by participating in certain virtual games or VR social platforms. In other words, these are not ‘unwanted’ sexual interactions at all. People are consenting to them by entering VR platforms, thereby removing any hint of immorality. Presumably, the idea here would be that either (a) people who use these VR platforms should know that they are liable to be exposed to this kind of experience, hence by using them they signal consent to an unrestricted set of

sexual interactions or (b) because it is relatively easy to withdraw from a virtual environment (take off your headset or switch off your computer) anyone who stays inside a virtual environment knowing that they are at risk of experiencing an unwanted sexual interaction is signalling consent by failing to withdraw.

This objection should not detain us for too long. As I have noted in past writings on this topic, this objection comes up frequently in conversations and talks that I have given about the topic of virtual sexual assault, but it seems highly implausible (Danaher 2018). To put it bluntly, it is a virtual equivalent of the 'asking for it' defence one sometimes hears in real-world cases of sexual assault and rape. It places the onus on the victim of a sexual assault to communicate non-consent and not on the perpetrator to elicit affirmative consent. Although there was a time when consent was readily inferred by courts and tribunals from non-verbal behaviour, the ethical and legal consensus seems to be shifting away from this norm (Dougherty 2013; 2015). It is unfair to force a victim to communicate their resistance to sexual assault or to 'run away' from a sexual aggressor. They may feel afraid or shocked or otherwise psychologically incapable of doing so. This is true even in virtual environments. Both Jordan Belamire and Michelle Cortese, for example, reported finding themselves initially 'frozen' and in a state of shock when they were subjected to virtual sexual assault. Requiring victims to remove themselves from virtual environments if they don't want to be subjected to sexual interactions would be unfair for an additional reason: it would mean that they are effectively excluded from these environments. Since many of the victims of virtual sexual assault are women, and since these virtual environments are becoming more common and integrated into our daily lives, this exclusion would just perpetuate a social injustice.

5. What Type of Moral Wrong Is It?

If the argument defended above is successful, then an unwanted virtual sexual interaction should be viewed as a serious moral wrong, which raises a follow-up question: what kind of moral wrong? We usually organise our moral wrongs into graduated hierarchies. This is most commonly, but not only, reflected in the criminal law. In other words, some wrongs are taken to be more serious than others. Giving someone a dead leg might be a form of physical assault, for example, but it is not as serious as beating a person to death with an iron bar. These graduated hierarchies are typically incorporated into the legal punishments we attach to different wrongs. For example, minor assault is defined differently from, and attracts a lower penalty than, intentional murder.

This is also true when it comes to sexual wrongs. Most legal jurisdictions around the world distinguish between rape and other kinds of sexual assault or non-consensual sexual activity. Rape is usually deemed the most serious kind of sexual wrong, with some additional gradations introduced for more serious forms of rape such as child rape and statutory rape. Sexual assault is a lesser kind of sexual wrong, and can come in various grades too. If we are to fit unwanted virtual sexual interactions into these graduated hierarchies, then where should we fit it? Is it less serious than rape? Is it a form of common sexual assault? Or is it something entirely *sui generis*?

In some ways, the simplest answer is to say that it is *sui generis*. It is a new phenomenon, made possible by new developments in technology, and should be treated as its own thing. To some extent, the present chapter has worked from the presumption that it is *sui generis*. After all, if it just another form of sexual assault — no different from the ones with which we are familiar — it wouldn't need a separate chapter dedicated to it in a book like this. Following this approach, we should analyse virtual sexual assault on its own terms and, if we

wish to criminalise it or legally ban it, then we should create a separate and novel category of legal offence to address it.

But, of course, treating it as a *sui generis* phenomenon would be to miss the point of much of the preceding discussion. The argument for moral wrongness and the subsequent analysis of objections to that argument hinges largely on drawing out the analogies between unwanted virtual sexual interactions and other kinds of moral wrongs, including specifically sexual wrongs. So it is at least worth asking whether we can just subsume this new phenomenon within the existing categories of rape and sexual assault. When we attempt to do this two possibilities appear to arise.

First, we could view unwanted virtual sexual interactions as essentially the same as rape and sexual assault. To do this, however, we would probably need to argue for a new understanding of rape and sexual assault. In most countries, both of these offences are defined in such a way that they require physical contact between the parties. Consider, for example, the definitions of rape and sexual assault in the English and Welsh Sexual Offences Act 2003. Under Section 1 of that act, rape is defined as an offence that requires penetration, by a penis, of the vagina, anus or mouth of another person (section 2 builds on this in creating an offence of ‘penetrative sexual assault’ that covers non-penile penetration of the vagina or anus). Under section 3 of that act, sexual assault is defined in a way that it requires physical touching. Given that virtual sexual interactions tend to lack this element of physical touching (though see the qualification below), they could not count as sub-types of rape or sexual assault.

Nevertheless, there is some reason to think that this physicalist paradigm ought to be abandoned. Human sexuality is a complex thing. People can experience sexual arousal and response in the absence of direct physical contact.

Instead of focusing on physical contact as being the defining feature of sexual assault we could, perhaps, focus on violations of sexual autonomy or sexual agency as being the defining feature (Danaher 2018a takes up this idea). If sexual activity can extend into the virtual world — and the prevalence of ‘cybersex’ and other kinds of virtual sexual activity would support this idea — then it is possible to have one’s sexual agency undermined and compromised through these sexual acts. A non-consensual virtual sexual act could then count as a form of sexual assault and may perhaps rise to a level of seriousness that justifies calling it ‘rape’. This suggestion is not as outlandish as it might first sound. There are some countries that have moved away from the traditional, physicalist paradigm — with its obsessive focus on what bits go where — when legally defining rape. For example, under Chapter 6, Section 1 of the Swedish Penal Code, rape can occur without physical penetration.⁷ This recently led to a man being convicted of rape for encouraging children to engage in sexual activity via a webcam.⁸

What if you are not persuaded by this proposal and think that we should stick with the physicalist paradigm? Does that mean that unwanted virtual sexual interactions can never count as rape or sexual assault? Not necessarily. There are haptic technologies that allow for touch, including sexual touch, to be transmitted via the internet. The world of ‘teledildonics’, for example, involves the creation of smart sex toys that allow people to engage in physically immersive forms of cybersex. It is possible to combine the use of teledildonics with immersive VR.⁹ If users of teledildonics have their interfaces hacked by

⁷ An English language translation of the Swedish criminal code can be found here: <https://www.government.se/498621/contentassets/7a2dcae0787e465e9a2431554b5eab03/the-swedish-criminal-code.pdf>

⁸ See [https://www.vice.com/en_us/article/pazyn7/in-a-first-a-man-is-charged-for-rape-over-the-internet. This non-physicalist approach to defining sexual assault could also have implications for how we view something like revenge porn \(or, to give it its more euphemistic name, the ‘non-consensual sharing of intimate images’\). Sharing such images non-consensually could also be viewed as a violation of sexual agency and autonomy.](https://www.vice.com/en_us/article/pazyn7/in-a-first-a-man-is-charged-for-rape-over-the-internet. This non-physicalist approach to defining sexual assault could also have implications for how we view something like revenge porn (or, to give it its more euphemistic name, the ‘non-consensual sharing of intimate images’). Sharing such images non-consensually could also be viewed as a violation of sexual agency and autonomy.)

⁹ The Dutch company Kiiroo makes a range of such devices. See <https://www.kiiroo.com/> for more.

malicious actors or are deceived into interacting sexually with someone with whom they did not consent to such interactions, then we may have cases of unwanted virtual sex that count as a genuine form of rape and sexual assault. Other forms of unwanted virtual sex — that do not involve the use of teledildonics — could then be treated as a separate category of moral and legal wrong. That said, it is worth bearing in mind that, even if they are not emphasised in the definitions of the offences, there are risks associated with physical sexual contact that do not arise with virtual sexual contact, e.g. sexually transmitted diseases and unwanted pregnancies. These risks may warrant differential treatment.

Conclusion

In this chapter I have analysed the phenomenon of unwanted virtual sexual interaction, described some of the forms it can take, and looked at some evidence suggesting that it is quite prevalent. I have argued that it counts as a serious moral wrong from both deontological and consequentialist perspectives because it violates rights and can cause harm. I have also argued that the objections to this view are not persuasive. I have also suggested that there is reason to treat these unwanted sexual interactions as a new type of sexual assault and rape, and not simply as a *sui generis* category of sexual wrongdoing.

One final issue that is worth addressing is what can be done about this new form of sexual wrongdoing. Three possibilities suggest themselves. First, and most importantly, designers of VR technologies and platforms have a major responsibility here. They create the possibilities for unwanted sexual actions within their virtual platforms and so they have a duty to work hard to ensure that people are not subjected to unwanted sexual interactions in VR spaces. In this regard, the proposals from Cortese and Zeller (2020) should be taken

seriously. They argue that there are ways to build ‘consent’ into the design of VR apps by creating features that allow users to control who they interact with, under what terms, and enable easy withdrawal from threatening spaces. More drastically, we could demand that designers prevent all forms of sexualised interaction in virtual spaces. Admittedly, this may be practically difficult in some cases, and may have a moral cost insofar as it could prevent desired VR sexual interactions, nevertheless there are presumably some kinds of VR platform (e.g. gameworlds populated primarily by children) where designing out the possibility of sexual contact is appropriate. Second, where sexual interactions are possible or desirable, there should be strong community standards that outlaw or ban unwanted sexual interactions in virtual spaces. These standards should be monitored and enforced by moderators of these virtual spaces. This is already being done, to some extent (as is clear from some of the stories discussed earlier on) but it is an important line of defence and should not be neglected. Third, and as implied above, there is a role for legal regulation of VR spaces. In particular, the criminal law should not treat what happens in VR spaces as somehow different or peripheral to its mission. Legislators and legal reformers should work to include virtual sexual wrongdoing within the scope of the law.

References

Alexander, Julie (2019). Roblox surpasses Minecraft with 100 million monthly players. *The Verge* 2nd August 2019, available at <https://www.theverge.com/2019/8/2/20752225/roblox-100-million-users-minecraft-youtube-twitch-pewdiepie-keemstar>

Behm-Morawitz, E. (2013). Mirrored selves: The influence of self-presence in a virtual world on health, appearance, and well-being, *Computers in Human Behavior*, Volume 29, Issue 1, 2013,

Belamire, J. 'My First Virtual Reality Groping' 5 October 2016, *Medium: Athena Talks*, available at <https://medium.com/athena-talks/my-first-virtual-reality-sexual-assault-2330410b62ee>

Blanke, O., and Metzinger, T. (2009). Full-body illusions and minimal phenomenal selfhood. *Trends Cogn. Sci.* 13, 7–13.
doi:10.1016/j.tics.2008.10.003

Cole, Samantha (2017). In a First, a Man Is Charged for Rape Over the Internet. *Vice* 1st December 2017, available at https://www.vice.com/en_us/article/pazyn7/in-a-first-a-man-is-charged-for-rape-over-the-internet

Cortese, Michelle (2019). Virtual_ Healing. *Medium* 18 November 2019, available at <https://medium.com/@ellecortese/virtual-healing-bf2b5f0cbf51>

Cortese, Michelle and Zeller, Andrea (2020) “Designing Safe Spaces for Virtual Reality: Ethics in Design and Communication” in DeRosa, Andrew, and Laura Scherling, eds (2020), *Ethics in Design and Communication*, London: Bloomsbury Academic Press.

Danaher, J (2016) 'Robots, law and the retribution gap'. *Ethics And Information Technology*, 18 :299-309

Danaher, J. (2017a). Robotic Rape and Robotic Child Sexual Abuse: Should they be criminalized? *Criminal Law and Philosophy* 11(1): 71-95

Danaher, J. (2017b). The Symbolic-Consequences Argument in the Sex Robot Debate. In Danaher, J. and McArthur, N. (eds). *Robot Sex: Social and Ethical Implications*. Cambridge, MA: MIT Press.

Danaher, J. (2018b). The Law and Ethics of Virtual Assault. In Barfield, W. and Blitz, M. (eds) *The Law of Virtual and Augmented Reality*. Cheltenham, UK: Edward Elgar Publishing.

Danaher, J (2019) 'REGULATING CHILD SEX ROBOTS: RESTRICTION OR EXPERIMENTATION?'. *Medical Law Review*, 27 :553-575

Dibbell, I. 'A Rape in Cyberspace: How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database Into a Society' in Ian Dibbell *My Tiny Life* (New York: Holt Publishers, 1998)

Dougherty, T. 'Sex, Lies and Consent' (2013) 123(4) *Ethics* 717–744

Dougherty, T. 'Yes means yes: communication as consent' (2015) 43(3) *Philosophy and Public Affairs* 224-253

Duranske, Benjamin (2007) "Reader Roundtable: 'Virtual Rape Claim Brings Belgian Police to Second Life', 24 April 2007, *Virtually Blind* available at <http://virtuallyblind.com/2007/04/24/open-roundtable-allegations-of-virtual-rape-bring-belgian-police-to-second-life/>

Gaus, G. (2010). *The Order of Public Reason*. Cambridge, UK: Cambridge University Press.

Gooskens, G. “The Ethical Status of Virtual Actions” (2010) 17(1) *Ethical Perspectives* 59–78

Hamm MP, Newton AS, Chisholm A, et al (2015). Prevalence and Effect of Cyberbullying on Children and Young People: A Scoping Review of Social Media Studies. *JAMA Pediatr.* 169(8):770–777.
doi:10.1001/jamapediatrics.2015.0944

Hershfield, H., Goldstein, D., Sharpe, W., Fox, J., Yeykelis, L., Carstensen, L., et al. (2011). Increasing saving behavior through age-progressed renderings of the future self. *J. Mark. Res.* 48, S23–S37. doi:10.1509/jmkr.48.SPL.S23

Huizinga, J. (1949). *Homo Ludens: A Study of the Play Element in Human Culture*. London: Routledge and Kegan Paul.

Lorenz, Taylor (2016). Virtual Reality Is Full of Assholes Who Sexually Harass Me. Here's Why I Keep Going Back. *Mic* 26 May 2016, available at <https://www.mic.com/articles/144470/sexual-harassment-in-virtual-reality>

Luck, M. ‘The Gamer’s Dilemma: An analysis of the arguments for the distinction between virtual murder and virtual paedophilia’ (2009) 11 *Ethics and Information Technology* 31-36.

Luck, M. and Ellerby, N. ‘Has Bartel Resolved the Gamer’s Dilemma?’ (2013) 15 *Ethics and Information Technology* 229-233

Lynn, R. ‘Virtual Rape is Traumatizing But is it A Crime?’, 4 May 2007, *Wired Magazine*.

Mack, Eric. (2018). Robert Nozick's Political Philosophy. In *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2018/entries/nozick-political/>>

Madary M and Metzinger TK (2016) Real Virtuality: A Code of Ethical Conduct. Recommendations for Good Scientific Practice and the Consumers of VR-Technology. *Front. Robot. AI* 3:3. doi: 10.3389/frobt.2016.00003

Matthias, A (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics Inf Technol* 6, 175–183
<https://doi.org/10.1007/s10676-004-3422-1>

Meehan, M., Insko, B., Whitton, M., and Brooks, F. (2002). Physiological measures of presence in stressful virtual environments. *ACM Trans. Graph.* 21, 645–652.

doi:10.1145/566654.566630

Metzinger, T. (2009a). *The Ego Tunnel. The Science of the Mind and the Myth of the Self*. New York: Basic Books

Morten Ebbe Juul Nielsen (2010) Safe, Sane, and Consensual—Consent and the Ethics of BDSM. *International Journal of Applied Philosophy* 24 (2):265-288 (2010)

Nielsen, Morten Ebbe Juul (2010). Safe, Sane, and Consensual—Consent and the Ethics of BDSM. *International Journal of Applied Philosophy* 24(2): 265-288

Nyholm, S. Attributing Agency to Automated Systems: Reflections on Human–Robot Collaborations and Responsibility-Loci. *Sci Eng Ethics* 24, 1201–1219 (2018).

Outlaw, J. (2018). Virtual Harassment: The Social Experience of 600+ Regular VR Users. The details of this report are available here:
<https://drive.google.com/file/d/1afFQJN6QAwmeZdGcRj9R4ohVr0oZNO4a/view>

Outlaw, J. and Duckles, B (2017). Why Women Don't Like Social Virtual Reality: A Study of Safety, Usability and Self Expression in Social VR. Full text available at: <https://extendedmind.io/social-vr>

Ostritsch, S. (2017). The amoralist challenge to gaming and the gamer's moral obligation. *Ethics and Information Technology* 19(2): 117-128.

Patridge, S. (2010). The Incurable Social Meaning of Video Game Imagery. *Ethics and Information Technology* 13(4): 303–312;

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

Reeves, C. (2018). The virtual simulation of child sexual abuse: online gameworld users' views, understanding and responses to sexual ageplay. *Ethics*

and Information Technology, 20(2), 101–113. <http://doi.org/10.1007/s10676-018-9449-5>

Rosenberg, R., Baughman, S., and Bailenson, J. (2013). Virtual superheroes: using superpowers in virtual reality to encourage prosocial behavior. *PLoS ONE* 8:e55003. doi:10.1371/journal.pone.0055003

Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful Human Control over Autonomous Systems: A Philosophical Account. *Frontiers in Robotics and AI*, 5, 73. <http://doi.org/10.3389/frobt.2018.00015>

Sicart, M. (2009) *The Ethics of Computer Games*. Cambridge: Massachusetts Institute of Technology Press

Slater M and Sanchez-Vives MV (2016) Enhancing Our Lives with Immersive Virtual Reality. *Front. Robot. AI* 3:74. doi: 10.3389/frobt.2016.00074

Søraker, Johnny (2010) *The Value of Virtual Worlds and Entities – A Philosophical Analysis of Virtual Worlds and their Impact on Well-being*. Inskamp: Enschede 2010.

Sparrow, Robert. (2017) Robots, Rape, and Representation. *International Journal of Social Robotics* 9, 465–477. <https://doi.org/10.1007/s12369-017-0413-z>

Sparrow, Lucy (2019). The moral (im)permissibility of groping in virtual reality games. Proceedings of DiGRAA 2019: What's Next? - available at http://digraa.org/wp-content/uploads/2019/01/DIGRAA_2019_paper_9.pdf

Sparrow, Gibbs and Arnold (2019). Apathetic Villagers and the Trolls Who Love Them. <https://doi.org/10.1145/3369457.3369514>

Strikwerda, L. 'Present and Past Instances of Virtual Rape in Light of Three Categories of Legal Philosophical Theories of Rape' (2015) *Philosophy and Technology* 28 (4):491-510

Smith, P. K. (2012). Cyberbullying and cyber aggression. In S. R. Jimerson, A. B. Nickerson, M. J. Mayer, & M. J. Furlong (Eds.), *Handbook of school violence and school safety: International research and practice* (p. 93–103). Routledge/Taylor & Francis Group.

Tillson, J. (2018). Is it distinctively wrong to simulate wrongdoing? *Ethics and Information Technology* DOI: 10.1007/s10676-018-9463-7

Wenar, L. (2020). Rights. In *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), available at <<https://plato.stanford.edu/archives/spr2021/entries/rights/>>.

Yee, N., and Bailenson, J. (2007). The proteus effect. The effect of transformed self-representation on behavior. *Hum. Commun. Res.* 33, 271–290. doi:10.1111/j.1468-2958.2007.00299.x

Zimmerman, E. and Salen, K. (2003). *Rules of Play: Game Design Fundamentals*. Cambridge, MA: MIT Press.