

# The Law and Ethics of Virtual Sexual Assault

Forthcoming in Barflied, W. and Blitz, M. *The Law of Virtual and Augmented Reality* (Cheltenham, UK: Edward Elgar Publishers, 2018).

By John Danaher

In 1993, Julian Dibbell wrote an article in *The Village Voice* describing the world's first virtual rape.<sup>1</sup> The incident he described took place in a virtual world called LambdaMOO, which is a text-based virtual environment that still exists to this day. People in LambdaMOO create avatars (onscreen 'nicknames') and interact with one another through textual descriptions. Dibbell's article described an incident in which one character (Mr. Bungle) used a "voodoo doll" program<sup>2</sup> to take control of two other users' avatars and force them to engage in sexual acts with both his avatar and one another. At least one of the victims found the incident extremely traumatising, with Dibbell describing her later as having 'posttraumatic tears' streaming down her face when she spoke about it.<sup>3</sup>

In 2003, a similar incident took place in *Second Life*. Second Life is a well-known virtual world, created by the engineer and entrepreneur Philip Rosedale. It is visual rather than textual in nature. People create virtual avatars that can interact with other user's avatars in a reasonably detailed visual virtual environment. In 2007, the Belgian Federal Police announced that they would be investigating a 'virtual rape' incident that took place in Second Life back in 2003.<sup>4</sup> Little is known about what actually happened, and the investigation does not appear to have progressed to a charge, but taking control of another character's avatar and forcing it to engage in sexual acts was not unheard of in Second Life. Indeed, regular users

---

<sup>1</sup> 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', *The Village Voice*, 23 December 1993. The article also appears as Chapter 1 in Dibbell's *My Tiny Life* (New York: Holt Publishers, 1998).

<sup>2</sup> A voodoo doll program is simply one that allows one user to take control of another user's character or program.

<sup>3</sup> Dibbell (n 1) *Ibid.*

<sup>4</sup> Lynn, R. 'Virtual Rape is Traumatizing But is it A Crime?', 4 May 2007, *Wired Magazine*.

noted that the use of voodoo doll program to engage in acts of ‘rape’ was well-known, though usually some form of ‘consent’ was sought.<sup>5</sup>

In October 2016, another instance of sexual misconduct in a virtual environment was reported. In a widely-discussed article, the journalist Jordan Belamire reported how she had been sexually assaulted while playing the game QuiVR, using the Oculus Rift.<sup>6</sup> The Oculus Rift is an immersive virtual reality system. Users don a headset that puts a virtual environment into their visual field. The user interacts with that environment from a first-person perspective, with the overall goal of the headset being to create an immersive virtual experience (to get you ‘lost’ in the virtual world).<sup>7</sup> QuiVR is an archery game where players fight off marauding zombies. It can be played online with multiple users. Players appear to other users in a disembodied form as a floating helmet and a pair of hands. The only indication of gender comes through choice of name and voice used to communicate with other players. Jordan Belamire was playing the game in her home. While playing, another user — with the onscreen name of ‘BigBro442 — started to rub the area near where her breasts would be (if they were depicted in the environment). She screamed at him to ‘stop!’ but he proceeded to chase her around the virtual environment and then to rub her virtual crotch.

Each of these three incidents involves what might be termed ‘virtual sexual assault’. They each involve actions that take place in virtual environments in which virtual characters, controlled by humans, engage in representations of sexually explicit acts (sexual touching, intercourse etc.) without the consent of at least one of their users. Clearly some of the participants were psychologically harmed by what happened. There are differences between the three cases. They involve different types of virtual environment, ranging from the simple

---

<sup>5</sup> On this, see the discussion of the *Second Life* case by Benjamin Duranske “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/>. I put the terms ‘rape’ and ‘consent’ in inverted commas because, as we shall see below, they are being discussed outside of their ordinary usage. Even the use of ‘rape’ and ‘consent’ in the same sentence is confusing since in many legal systems rape is simply defined as non-consensual sex so you could consent to a rape.

<sup>6</sup> 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>

<sup>7</sup> The feeling of immersion alters the user’s experience of the virtual world. It is no longer separate or distant but phenomenologically real. The extent to which the mind can play these tricks of immersion on us, causing us to believe that we are located in different physical spaces, or that our bodies extend into physical artifacts, is well-documented in the philosopher Thomas Metzinger’s book *The Ego Tunnel* (New York: Basic Books, 2010). I return to this important point later in the chapter.

text-based world of LambdaMOO, to the more complex, immersive visual world of QuiVR. Nevertheless, the three incidents raise similar ethical and legal questions. *Prima facie*, there is something undesirable about the conduct involved. But how serious is it and what should we do about it? Should we really refer to such incidents as ‘virtual sexual assault’? Is this something that needs to be regulated and criminalised? Do existing laws on sexual assault and sexual harassment cover incidents of virtual sexual assault, or do we need a new set of criminal offences to address this issue? These are the questions addressed and answered in this chapter.

The chapter is organized into four main parts. In part one, I propose a definition and classificatory scheme for understanding the phenomenon of virtual sexual assault. I argue that virtual sexual assault is complex due to the different modes of interaction that are possible in a virtual environment, and that there could be as many as six different types of virtual sexual assault that are worthy of our consideration. In part two, I consider how seriously it is worth taking these six different types of virtual sexual assault, addressing the question of whether criminalisation is warranted. I also consider whether existing laws on sexual assault could be interpreted so as to cover cases of virtual sexual assault. I follow this up in part three by looking at responsibility for virtual sexual assault. I ask whether it is right and proper for us to deem someone responsible for actions that they perform in a virtual world, particularly if those virtual actions have no obvious real-world effects. Finally, I address a variety of issues that arise in relation to consent (and its absence) in virtual sexual encounters.

I do not propose any one single view of virtual sexual assault in this chapter. I think the issue is complex, with some forms of virtual sexual assault worth treating as seriously as any real world form of sexual assault, and other forms worth taking less seriously. But my goal is not to offer the last word on any of the claims I make. It is, instead, to clarify and provoke future research and debate on this important topic. The analysis is intended to be general in scope, but on occasion I do rely on legal examples and case studies drawn from the jurisdictions with whose laws I am most familiar, particularly the law of sexual assault in England and Wales. These examples are provided simply to illustrate and explain the arguments being presented.

## **1. Classifying Virtual Sexual Assault**

Before we can answer some of the interesting ethical and legal questions that arise from the phenomenon of virtual sexual assault, we must get a clearer picture of what that phenomenon entails. The legal philosopher Litska Strikwerda is one of the few people who has addressed this topic at any length,<sup>8</sup> although she focuses specifically on virtual ‘rape’. What I propose here is an expansion and update of her definition and classificatory scheme.

I start with the definition. Virtual sexual assault is a species of virtual act. Philip Brey has defined a virtual act as any act that is ‘initiated by a user within a virtual environment and is defined over objects and persons within that environment’.<sup>9</sup> This is a useful starting point. Strikwerda argues that we can combine this definition of a virtual act with the definition of a real world rape in order to create a satisfactory definition of virtual rape. She does so by adopting a definition of real world rape as ‘the act of forcing sex upon an unwilling party’,<sup>10</sup> and then defines virtual rape as the virtual act of forcing sex upon an unwilling party, where the act is defined over persons in the virtual environment. For her, the critical question then becomes what type of virtual environment is used to perform the act of virtual rape. She distinguishes between two such environments: those that require users to interact using avatars (i.e. onscreen textual or visual representations that they manipulate using a controller) or those that facilitate immersive interactions, such as the one described above for the game QuiVR.

There are several problems with this suggested definition. First, Brey’s original definition of a virtual act limits its focus to virtual acts that are performed by human users of virtual characters. In other words, it imagines cases where there is a human inputting commands that influence the actions of onscreen characters. While this may be the most interesting and relevant case of a virtual act it is not the only form that virtual acts can take. Many of the characters in virtual environments may be ‘wholly virtual’ in nature, i.e. computer programmed characters with some degree of artificial intelligence. Since we are living through an era in which artificial intelligence is on the ascendancy, due to advances in

---

<sup>8</sup> See Strikwerda, L. ‘Present and Past Instances of Virtual Rape in Light of Three Categories of Legal Philosophical Theories of Rape’ (2015) 28 (4):491-510

<sup>9</sup> Brey, P. ‘The Physical and Social Reality of Virtual Worlds’ in Grimshaw, M. (ed) *The Oxford Handbook of Virtuality* (Oxford: OUP, 2014), p. 49 - this quote was originally sourced through Strikwerda (n x).

<sup>10</sup> Strikwerda (n 8), p 493

machine learning algorithms,<sup>11</sup> there is every reason to expect that wholly virtual characters will play an increasingly important role in our virtual interactions in the future. This is important because it means that the distinguishing features of virtual sexual assault are not determined solely by the type of virtual environment used for the performance; they are also determined by the type of virtual character performing the act: is it human-controlled or not? Second, the definition of rape used by Strikwerda is both limiting and anachronistic. To describe ‘rape’ as ‘forced sex’ is contentious. Although there are jurisdictions that still insist upon a force requirement when it comes to proving the occurrence of rape, many jurisdictions now favour a definition of rape as non-consensual sex. At the same time, Strikwerda’s definition of rape does not share the anatomical obsessions of some jurisdictions and so will seem misleading to many. For example, in England and Wales, rape is explicitly defined, in Section 1 of the Sexual Offences Act 2003, as non-consensual ‘penetration of the vagina, anus, or mouth’ by the penis. Anything that does not involve this penetrative interference is not counted as ‘rape’. Of course, it makes sense for Strikwerda to drop the anatomical obsessions since virtual acts will never involve the penetration of one biological orifice by another biological organ — indeed one could argue that the absence of direct biological contact is the essence of a virtual interaction — but this merely underscores the point I am making: it is anachronistic to focus solely on virtual rape if your interest is in unwanted sexual actions in virtual environments more generally.

What I suggest then is a broader definition of virtual sexual assault:

**Virtual Sexual Assault:** An unwanted, forced, or nonconsensual sexually explicit behaviour performed by virtual characters, to one another, acting through representations in a virtual environment.

This is a deliberately vague definition and requires some unpacking. It covers all manner of sexually explicit behaviours, not just those that we would classify as ‘rape’. That said, the definition allows for sub-types of virtual sexual assault, such as virtual rape. In this regard, I would suggest that virtual rape be defined as a sexually explicit behaviour performed by virtual characters through representations in a virtual world that, if performed in the real

---

<sup>11</sup> Alpaydin, E. *Machine Learning: The New AI* (Cambridge, MA: MIT Press, 2016); Domingos, P. *The Master Algorithm* (New York: Basic Books, 2015); and Bostrom, N. *Superintelligence: Paths, Dangers, Strategies* (Oxford: OUP, 2014)

world, would count as rape. This avoids placing any reliance on specific anatomical representations, which could vary across jurisdictions, but allows for this should it be required in a particular jurisdiction.<sup>12</sup>

The concept of ‘representation’ is important to this definition of a virtual assault. This importance derives from the observation made previously about the absence of direct biological contact being the essence of the virtual interaction. The suggestion here is that what makes a virtual act different from a real world act is that instead of the action being performed ‘flesh on flesh’ (so to speak) it is performed through representations, i.e. on-screen avatars or the like. There are complications to this, as we shall see below, nevertheless, acting via a representation is, I submit, central to what makes a virtual act ‘virtual’.

Finally, the definition also hedges on the normative condition that turns a virtual sexual act from something that is acceptable into something that counts as assault. Although my general preference is that discussions of virtual sexual assault should focus on the idea of ‘consent’, there may be problems applying this to all apparent incidents of virtual sexual assault, particularly if the characters involved, due to their being wholly virtual, are deemed incapable of consenting. This is why I hedge to include the possibility of unwanted or forced contact. It is important to realise that when it comes to interpreting and applying these normative conditions, representation still plays a key role. Although the virtual acts may involve a real-world human being who does not consent or want their performance (or upon whom they are ‘forced’), that is merely a *sufficient* condition for something counting as an incident of virtual sexual assault. A representation that depicts an act that *appears* to be nonconsensual, unwanted or forced will also count.

The deliberate vagueness of the definition can be tempered by developing a classificatory scheme that makes sense of the different forms that virtual sexual assault can take. This helps us to draw important ethical and legal distinctions between the possible instances of virtual sexual assault, and allows us to hone in on the ones that raise the most significant concerns. We can start this exercise by distinguishing between virtual sexual assaults on the basis of the virtual characters involved in their performance: do they involve characters controlled by a human user or are they wholly virtual? We can also distinguish between virtual sexual

---

<sup>12</sup> In other words, in the UK a virtual rape would only arise where a virtual character performs an act that represents the penile penetration of the vagina, anus or mouth of another virtual character.

assaults depending on who occupies the role of the victim and the perpetrator. Were human users both the victims and perpetrators? Or did a human user assault a wholly virtual character? Or vice versa? We can easily construct a two-by-two matrix for categorising the different possibilities here, with one dimension for the *human user-virtual character* distinction and another for the *perpetrator-victim* distinction.

But before we do so, we must add another layer of complexity to the classificatory scheme. As noted already, Strikwerda argues that we should categorise the different types of virtual rape depending on the *mode of virtual interaction* they involve. This seems appropriate. *Prima facie*, there is something different about an interaction that takes place through avatars that are controlled via a mouse or other controller — as is the case in a virtual world like *Second Life* — and one that takes place via an immersive device like the Oculus Rift — as is the case in a virtual world like QuiVR. There is greater ‘distance’ between the real world and the virtual world in the former case and this distance could be ethically and legally significant.<sup>13</sup> This is particularly true when you realise that immersive interactions could involve the use of haptic technologies — i.e. technologies that allow you to transmit touchlike sensations via a network. Companies have already started to create haptic devices that facilitate long-range virtual sexual interactions. The Dutch company Kiiroo,<sup>14</sup> for instance, has a range of ‘teledildonic’ devices that enable people to have virtual sexual experiences that replicate touchlike sensation. The users wear immersive helmets, watch their partners or pornographic actors, and use vibrators and artificial vaginas for sexual stimulation. The devices allow for interactivity, whereby one character can stimulate another through the haptic device, at a distance. The users still act or are acted upon via representations — the onscreen depictions of their partners or the actors — but the gap between the real and the virtual is narrowed. The actions performed via the virtual representations are immediately transduced into effects on real human bodies: the users can feel completely psychologically immersed in the experience. This seems very different from the text-based interactions via avatars in LambdaMOO or the visual interactions in Second Life. In those cases the user never completely dissolves the gap between themselves and the virtual representations through which they act. So I propose to follow Strikwerda by distinguishing between virtual sexual assaults involving *avatar-interactions* and those using *immersive-interactions*.

---

<sup>13</sup> The concept of ‘distance’ and its ethical and legal importance is discussed at greater length below.

<sup>14</sup> For information <https://www.kiiroo.com>, accessed 8/14/2017

This makes it more difficult to construct a simple two-by-two matrix for classifying the different kinds of virtual sexual assault since we are now dealing with three dimensions of difference (the *human user - virtual character*; *perpetrator-victim*; and *avatar-immersive* dimensions). We could solve the problem by constructing a cube to represent the three dimensions, but it is also still possible to work with the two-by-two grid by simply compromising on how you identify one of the dimensions of difference. I have done this in the following diagram.

	<b>Avatar-interaction</b>	<b>Immersive-interaction</b>
<b>Human Perpetrator</b>	<p><b>Virtual victim:</b> human user virtually sexually assaults a wholly virtual agent using an avatar.</p> <p><b>Human victim:</b> human user virtually sexually assaults an avatar controlled by another human being, using an avatar.</p>	<p><b>Virtual victim:</b> Human user virtually sexually assaults a wholly virtual agent using immersive tech.</p> <p><b>Human victim:</b> Human user virtually sexually assaults a virtual agent controlled by another human being, using immersive tech.</p>
<b>Virtual Perpetrator</b>	<p><b>Virtual victim:</b> wholly virtual avatar is virtually sexually assaulted by a wholly virtual agent.</p> <p><b>Human victim:</b> avatar controlled by a human being is virtually sexually assaulted by a wholly virtual agent.</p>	<p><b>Virtual victim - n/a</b> this interaction is not distinct from an interaction between two wholly virtual avatars since the immersive tech only applies to fuse the real-world with the virtual world.</p> <p><b>Human victim:</b> human user is virtually sexually assaulted by a wholly virtual agent controlled by another human being while wearing immersive tech.</p>

This diagram depicts what I call the ‘logical space’ of virtual sexual assault.<sup>15</sup> It shows all the possible manifestations of virtual sexual assault (given the three specified dimensions). Not every location within this logical space is worthy of serious consideration or analysis. In particular, I would argue that the interactions that involve wholly virtual characters as both perpetrator and victim do not merit consideration. Why not? I work with the assumption that

wholly virtual characters are not, and will not for quite some time, count as moral victims or moral agents. That is to say, I assume that they cannot be harmed by what is done to them (at least, not ‘harmed’ in any morally significant way) and they cannot be morally responsible for what they do. Consequently, wholly virtual incidents of sexual assault are purely representational in nature. They are best analysed and treated as a sub-category of pornography, not as a type of ‘assault’.

Omitting this type of virtual sexual assault still leaves six potentially significant locations within the logical space. Each of these locations involves either a human perpetrator, human victim, or both. For ease of reference, I will label all six of these locations in the following manner:

**Type A** - Human user virtually sexually assaults a wholly virtual character using an avatar (Worthy of consideration because the human perpetrator is a moral agent and so the act may say something about their moral character)

**Type B** - Human user virtually sexually assaults a virtual avatar that is controlled by another human being, using their own avatar (Worthy of consideration because both the perpetrator and victims are moral agents and moral patients.)

**Type C** - Human user virtually sexually assaults a wholly virtual character using immersive technology. (Worthy of consideration because the human perpetrator is a moral agent and so the act may say something about their moral character.)

**Type D** - Human user virtually sexually assaults a virtual avatar controlled by another human (wearing immersive technology) while also using immersive technology. (Worthy of consideration because both the perpetrator and victims are moral agents and patients.)

---

<sup>15</sup> I adopt the terminology from the work of Christian List. See List, C. ‘The Logical Space of Democracy’ (2011) 39(3) *Philosophy and Public Affairs* 262-297 and List, C. and Vallentini, L. ‘Freedom as Independence’ (2016) 126(4) *Ethics* 1043-1074

**Type E** - Wholly virtual character virtually sexually assaults a virtual avatar controlled by a human user. (Worthy of consideration because the human victim is a moral patient and so can be harmed by the interaction.)

**Type F** - Wholly virtual character virtually sexually assaults an avatar controlled by a human user who is using immersive technology. (Worthy of consideration because the victim is a moral patient and so can be harmed by the interaction).

What I propose to consider over the next two sections is exactly how seriously each of the six types ought to be treated. I do that by first considering the harm that might be involved in these cases and then by considering responsibility for virtual acts.

## **2. Should we criminalize virtual sexual assault?**

Let's return to our opening examples. When you read about the 'rapes' in LambdaMOO and Second Life, or the 'groping' in QuiVR, what was your reaction? Did you sense that something wrong was happening? Was it wrong enough to warrant criminalisation? Before answering those questions, I would suggest considering a real-world case that shares certain similarities with the cases of virtual sexual assault identified above.

The case in question is the English case of *R v. Devonald* [2008] EWCA Crim 527. The case involved a man named Stephen Devonald who posed on the internet as a 20-year-old woman named 'Cassey'. He did so in order to make contact with the victim, who was a 16-year-old boy, and the former boyfriend of Devonald's daughter. Using the false persona, Devonald encouraged the victim to masturbate in front of a webcam. The goal, apparently, was to film the victim and post the video online in order to cause him some embarrassment. This was to exact revenge for the way in which the victim had treated Devonald's daughter. Devonald was convicted under Section 4 of the Sexual Offences Act 2003, which makes it an offence to cause another to engage in sexual activity without their consent.

The case is interesting for a number of reasons, not least of which is the novel way in which the court interpreted what it meant to consent to sexual act. (The court held that the Devonald had deceived the victim as to the nature and purpose of the act in question and thereby vitiated his consent to the act. We will return to issues relating to consent and

deception later in this chapter.) For now, I want to highlight three features of the case that are instructive for present purposes. The first is that the case involved a type of virtual interaction between a perpetrator and a victim: they interacted via on-screen representations, with Devonald using the fake avatar of ‘Cassey’ to persuade the victim to engage in masturbation. Second, despite acting via virtual representations, the interaction clearly had real-world consequences and repercussions: the victim’s masturbatory act occurred in the real world, showing how virtual interactions can sometimes spill over into reality. Finally, the circumstances of the case, despite having novel features, already fell under the scope of an existing criminal offence, namely the offence of causing another to engage in sexual activity without consent. There was no need to create a new offence simply because the actions were mediated through technological artifacts. What’s more, I doubt anyone felt as though it was wrong to apply that offence to this scenario: it seems perfectly appropriate to do so.

This holds some important lessons for the present inquiry. In asking whether the incidents of virtual sexual assault identified above warrant criminalisation, we need to be cognisant of the blurry boundary between the virtual and the real, and the fact that existing criminal offences already cover some virtual actions. We also need to be cognisant of the fact that we already live much of our lives online, and that many of our simple day-to-day interactions are mediated through virtual environments. Consequently, although it may seem initially odd to many people, the idea of subjecting virtual acts to criminal sanction is neither novel nor exceptional. Granting this, we may still worry that some incidents of virtual sexual assault slip through the gaps of existing laws, and that even if they are covered by existing offences, these offences are not serious enough to address the gravity of the harm caused by the relevant incidents. This suggests that instead of simply looking to the different possible forms of virtual sexual assault and how they might match up with the existing law, we should first develop a principled framework for answering questions pertaining to the appropriate criminalisation of virtual acts. Over the remainder of this section I will try to develop such a framework and then use it to make two arguments about how we might go about criminalising certain instances of virtual sexual assault.

(i) *When should we criminalise virtual acts?*

I think there is a simple test we can apply to determine the appropriateness of criminalising virtual acts. The test is this: a virtual act X can be criminalised when it is *real*

and when its properties and effects satisfy the conditions that generally warrant criminalisation. This simple test masks some complexities. Two follow up questions must be asked: When are virtual acts (or their effects) real? And what are the conditions of an act that would ordinarily warrant its criminalisation?

In answer to the first question, we can be guided by the work of the aforementioned Philip Brey. He has developed a sophisticated framework for understanding the reality of the virtual world. As he points out, there is much confusion about the status of objects and events in virtual environments. Some people think that if you say that an action or object exists virtually you are, effectively, saying that it is not real. This makes the phrase ‘virtual reality’ sound strangely oxymoronic. But this, of course, is incorrect. Virtual entities clearly are real *qua* virtual entities. To use one of Brey’s examples, a virtual apple (i.e. a representation of an apple in a virtual environment) does not exist *qua* real apple (you cannot bite into it and taste its flesh), but it does exist *qua* virtual apple. It really does exist within the virtual environment. This is to say that the virtual apple really exists as a *simulation* or *representation*. But, and this is the important point, that’s all the virtual apple’s existence will ever be: a simulacrum. The same goes for any virtual object or event that represents a real-world object or event that necessarily has physical properties or dimensions. Tables, chairs, rocks, water, and natural gas can only ever be simulated in virtual reality; they cannot exist *qua* real tables, chairs, rocks, water and natural gas in the virtual environment.

Objects and events that are necessarily physical are easy cases. Other virtual objects and events share a more confusing relationship with reality. Consider something like a clock. I have a virtual clock on my laptop right now. I can read the time off it (it’s 14:35, in case you were wondering). This virtual clock is different from the clock upon my wall: it has no physical or material instantiation (apart from its symbolic representation in the microchips of my computer). But this doesn’t make it any less real than the clock upon my wall. There is no significant ontological difference between the two. This is because the status of clocks is determined solely by their functionality, not by the presence or absence of specific physical properties.

Now consider something even more complex, like money. The naive view — one that many have had over the course of history (though this may not be true for future generations) — is that money is a physical object: it is some valuable commodity like gold and silver, or it

is the notes and coins that we carry in our pockets. This naive view is obviously false. Money is a functional kind, like a clock. Its reality is determined by the function it performs in recording debts and obligations, and in paying for goods and services.<sup>16</sup> But money is a different functional kind from a clock. Although the clock's status is determined by its functionality, the success or failure of that functionality depends on its ability to track a real world phenomenon: the passage of time. There is no similar real-world phenomenon determining the functionality of money. Money is a *social construction*. It exists and has the functions it has through an exercise of collective imagination and will. Through collective action we can confer the status of money on pretty much anything, and nowadays we confer the status of money on digitally recorded bank balances and cryptocurrency ledgers. Indeed, most money that exists today — that you regularly use to pay for goods and services— is purely virtual in nature. It never has, and never will, exist in the form of physical currency.

This complicates our understanding of the relationship between the real and the virtual quite a bit. It seems that there are some virtual objects and events that can only ever exist as simulations or representations; but there are others that, although represented in a virtual environment, are every bit as real as their physical counterparts. Brey argues that we can use John Searle's theory of social reality to understand what is going on here.<sup>17</sup> Searle's theory holds that there is a distinction between physical facts (i.e. physical objects, events, states of affairs) and socially constructed facts. The former are ontologically objective and epistemically objective: they exist in the absence of human minds and we can all come to know of their existence through mutually agreed upon methods of inquiry. The apple from earlier on would be a classic example. Socially constructed facts are different. They are ontologically subjective and epistemically objective. They only exist in the presence of human minds (because they are created through acts of collective will) but this does not mean that we can simply imagine into existence whatever social objects we like. There has to be collective agreement on their status and this gives social objects much of the intersubjective persistence and stability that we associate with 'real' objects. Hence they are epistemically objective. Money really exists and we all agree that it exists in certain amounts and distributions. My bank balance is an epistemically objective fact: I cannot imagine extra

---

<sup>16</sup> For lengthy defences of this conception of money, see Ingham, G. *The Nature of Money* (London: Polity Press, 2004); Graeber, D. *Debt: The First 5000 Years* (New York: Melville Press, 2011); and Martin, F. *Money: The Unauthorised Biography* (London: The Bodley Head, 2013)

money into my account, no matter how hard I try. Searle argues that this happens because we confer the status of money on certain representations through the collective creation of a *constitutive rule*, i.e. a rule of the form ‘X counts as Y in context C’. A significant number of the objects, events and activities with which we engage everyday are social in nature and are made possible through collective constitutive rules. Getting married, being promoted, passing an exam, receiving an award, and being qualified/credentialed are all examples of this.

The upshot of this is that we have a principled way of determining whether a virtual object or event is real or not. It is real, according to Brey, if it is a purely functional object or if falls within the definition of Searle’s social facts; it is merely a simulation if it falls within the definition of Searle’s physical facts.

There is one additional wrinkle to this. Brey notes that virtual acts require a somewhat different analysis. Whenever a virtual action involves at least one human participant — and as noted above we are only interested in these cases — it can have intravirtual and extravirtual effects.<sup>18</sup> The intravirtual effects are those that only affect the virtual representations. The extravirtual effects are those that affect the human beings who control or participate in the virtual actions. The intravirtual effects can only be real if the virtual representations fall within the scope of Searle’s social facts. The extravirtual effects are always, uncontroversially, real, either because they have physical and mental effects on the human user or because they have effects on the extravirtual social reality in which the user operates.

So much for the nature of virtual reality. We now must confront the second sub-question. I said earlier that virtual sexual assault might warrant criminalisation if it is real and if it meets the conditions that ordinarily warrant criminalisation. We have some sense of the former now, but what about the latter? What are the conditions that warrant criminalisation? Without seeking to defend my views at length, I here adopt a modified version of a position known as ‘legal moralism’. This might seem like an inauspicious starting point. Legal moralism is often associated with historically regressive attitudes towards the criminalisation

---

<sup>17</sup> See Brey (n 9) and Brey, P. ‘The Social Ontology of Virtual Environments’ (2003) 62(1) *American Journal of Economics and Sociology* 269-282. See also Searle, J. *The Construction of Social Reality* (London: Penguin, 1996); and Searle, J. *Making the Social World* (Oxford: OUP, 2010).

of private conduct, often sexual, in nature. For instance, the most significant debate about legal moralism in the 20<sup>th</sup> century was the debate between the legal theorist H.L.A Hart and the judge Patrick Devlin about the criminalisation of homosexuality.<sup>19</sup> In that debate, Devlin defended what has been referred to as the moralistic view, according to which the criminalisation of homosexuality was legitimate. Hart defended a narrower, liberal view, according to which it was not. He held that harm to others was an essential condition for the criminalisation of any act. Since homosexuality did not cause harm to others, it should not be criminalised.

If legal moralism is to be understood along the lines set by the Hart-Devlin debate, then I quite agree that it would be an inauspicious place from which to begin. If it were, I would be bound to the notion that actions which cause disgust or give offence warrant criminalisation, and this is not a position I hold. Fortunately, it does not have to be understood along these lines. Indeed, if anything, the Hart-Devlin debate is misleading as to the true nature of legal moralism. Devlin's view was really one of legal populism, not moralism, since it was about ensuring social cohesiveness through the law; and Hart's view was itself a species of legal moralism, since he presumably believed that causing harm to others was immoral.

My understanding of legal moralism is much narrower.<sup>20</sup> It is as follows: a necessary condition for the criminalisation of any act is that it be immoral in some way. This is quite a modest view and, I would argue, a sensible one. Immorality comes in many flavours. If we follow the work of social psychologists such as Jonathan Haidt it could come in as many as six different flavours, including loyalty to authority, fairness, harm, disgust, freedom, and sacredness.<sup>21</sup> Whatever the case may be, the essence of legal moralism is that an act must be some flavour of immoral in order for it to warrant criminalisation.

Legal moralism, as defined, only identifies a necessary condition for criminalisation, not a sufficient one. Identifying a sufficient condition is much trickier and, indeed, there are likely to be many sufficient conditions. Few if any defenders of legal moralism hold that all

---

<sup>18</sup> The terminology comes from the work of Johnny Soraker and his dissertation *The Value of Virtual Worlds and Entities – A Philosophical Analysis of Virtual Worlds and their Impact on Well-being* (Inskamp: Enschede 2010). It is referenced in both Brey (n 9) and Strikwerda (n 8)

<sup>19</sup> Devlin, P. *The Enforcement of Morals* (Oxford: OUP 1965); and Hart, HLA *Law, Liberty and Morality* (Stanford, CA: Stanford University Press, 1963).

<sup>20</sup> For details on this understanding of legal moralism, see Simester, A.P. and Von Hirsch, A. *Crimes, Harms and Wrongs* (Oxford: Hart Publishing, 2011), pp. 22–30

immoral conduct warrants criminalisation. The more common view is that certain sub-types of immoral conduct warrant criminalisation. Following Anthony Duff, I hold that one feature of these sub-types is that they must implicate a public wrong, i.e. an immoral act that demands public attention, scrutiny and accountability.<sup>22</sup> Lying to your friends about where you were last night might be immoral, but it wouldn't be a public wrong. Lying on your tax return is a different matter. Furthermore, and following the classic liberal view, I hold that the most suitable candidate immoral actions are those that cause harm to others, provided that the harm crosses a certain threshold of seriousness and is a matter of public concern. I would also allow for the possibility of some immoral actions that warrant criminalisation and yet do not cause harm to others. So, for example, desecrating the bodies of the dead or destroying natural objects might warrant criminalisation, even if no one is harmed in these actions.<sup>23</sup> Of course, the most contentious issue is whether actions that cause harm to the self warrant criminalisation. My own personal view is that they do not, though some actions may cause harm to one's moral character, which could have downstream effects on others, and so could, in certain circumstances, warrant criminalisation.

Where does that leave us in relation to virtual sexual assault? Recall, the position being defended here is that virtual sexual assault warrants criminalisation if it is real and if it meets the conditions that ordinarily warrant criminalisation. So we need to ask ourselves are virtual sexual assaults real and do they cause harm or implicate public wrongs? The answer to the latter seems relatively straightforward – they obviously can cause harm to others since users can be psychologically traumatised or even physically interfered with, depending on the mode of virtual interaction – but the answer to the former is more contentious.

There are effectively three positions you can take on the nature of sexual assault:

**I. Sexual assaults are essentially physical:** An action only deserves to be called 'sexual' if it involves sexually motivated or stimulated physical touching. This would mean that virtual sexual assaults are not real sexual assaults; they are mere simulations of sexual assault.

---

<sup>21</sup> Haidt, J. *The Righteous Mind* (London: Penguin 2013).

<sup>22</sup> Duff, A. "Towards a Modest Legal Moralism" (2014) 8 *Criminal Law and Philosophy* 217–235.

<sup>23</sup> On these examples, see Scoccia, D. "In defense of "pure" legal moralism" (2013) 7 *Criminal Law and Philosophy* 513–530

**II. Sexual assaults are essentially social:** An action acquires the status of being ‘sexual’ through acts of collective will and intention. This would mean that virtual sexual assaults could be every bit as real as real-world sexual assaults.

**II. Sexual assaults can be both physical and social:** Some sexual acts are essentially physical in nature (e.g. sexual intercourse) others perhaps largely social in nature (e.g. sexual harassment or obscene communications), and many include a mix of both physical and social properties.

The first position tallies with the law in many countries. As mentioned previously, under the Sexual Offences Act of 2003 in England and Wales, rape is defined as the intentional penile penetration of the vagina, anus or mouth; and the related offence of assault by penetration is defined as the penetration of the vagina or anus with something other than the penis. Both offences require that certain physical conditions be met: particular biological orifices must be penetrated by particular objects. The same goes for the more general offence of sexual assault. This is defined under Section 3 of the 2003 Act as intentional sexual touching, which again seems to require some physical contact. If you adopted the first position, then virtual sexual assault could only be said to occur when a human user was assaulted by another virtual agent (human controlled or not) whilst wearing a haptic suit that enabled them to satisfy the relevant physical conditions (such as penetration or touching). This doesn’t mean that the mere simulations of sexual assault could not be criminalised — if it causes extravirtual harm to others, or to the individual performing it, that are tantamount to public wrongs, then it might be properly criminalised — but it does mean that it could not be criminalised *qua* sexual assault. It would have to be criminalised as something else.

The second position is more at home with certain traditional and contemporary feminist understandings of sexual assault. As Strikwerda points out,<sup>24</sup> sexual offences like rape were traditionally understood as property-related crimes: the crime being sexual interference with the property of a husband or father. Property is a social construct: an object, person or piece of land is only your property if society confers that status upon it through collective will. One could view a virtual avatar as a type of property and hence virtual sexual assault as a real

---

<sup>24</sup> Strikwerda (n 8)

form of sexual assault. But even if you do conceive of sexual assault in these property-related terms, you would still be tasked with distinguishing sexual interferences with property rights from non-sexual ones. For example, what, traditionally, distinguished trespass from rape? Most countries answered this by appealing to certain essentially physical characteristics of sexual acts, which suggests that this traditional, conservative view is, in practice, similar to the first position when it comes to viewing sexual acts as essentially physical. It would need to be updated in order to cover virtual sexual assault.

A more promising way to defend the second position is to use certain ideas and concepts drawn from feminist theory.<sup>25</sup> It is, of course, a mistake to view feminism as a single coherent theoretical framework. There are many feminisms and they sometimes contradict. Nevertheless, there is a prominent strain of feminist theory, with adherents in legal theory, that argues that many of the characteristics and concepts we hold to be ‘sexual’ are social constructs. Why is it sexually provocative and immodest for a woman to bare her arms in some cultures and not in others? The answer lies in socially constructed norms of sexuality. This constructionist view also carries over into how certain feminists think we should understand sexual assault. One suggestion is that sexual assault is best understood as a non-consensual interference with someone’s *sexual agency*, where agency is understood as a social construct.<sup>26</sup> This is not an uncommon view, with several feminist scholars arguing that sexual agency is a construct that emerges from competing discourses, prejudices and beliefs in society about what an individual can and cannot permissibly do. The typical view is that female sexual agency is undermined and limited by patriarchal ideologies — that we live in a ‘rape culture’ that normalises interferences with female sexual agency and constantly polices its boundaries. If you adopt this view, then there is no reason why virtual sexual assault couldn’t be every bit as serious an interference with sexual agency as real world sexual assault.

The third position is probably closest to my own. This holds that sexual acts can have different mixes of physical and social properties. Some sexual offences are properly viewed as essentially physical – penetrative sexual assault and rape being the most obvious cases. Some are essentially social in nature – sexual harassment being a possible example since,

---

<sup>25</sup> Strikwerda (n 8) also examines this line of argument.

although it is usually defined in terms of *unwanted* or inappropriate sexual advances, the definition and determination of what counts as an unwanted or inappropriate sexual advance is largely a product of social convention. And, finally, some are just mixes of physical and social properties – simple sexual assault is arguably of this form because, although typically defined as ‘sexual touching’, the definition and determination of what counts as ‘sexual’ touching is again a matter of social convention. This third position is easily reconciled with the idea that there is a range of sexual offences, differentiated not only in terms of the seriousness of the harm they cause, but also in terms of their metaphysical constitution (physical, social or a mix of both). According to this view, some forms of virtual sexual assault could be captured by existing offences and others might require the creation of new offences, which are still properly classified as being ‘sexual’ in nature.

*(b) How should we go about criminalising virtual sexual assault?*

To this point, I have tried to outline a framework for thinking about virtual sexual assault and its possible criminalisation. Is there anything more concrete and prescriptive to be said? How exactly should we treat the six types of virtual sexual assault identified in the previous section? I have two arguments to make in relation to this.

First, I think it is probably a mistake to criminalise virtual sexual assault of types A and C. These are cases where a human user sexually assaults a wholly virtual character, either through an avatar or using immersive technology. In neither case is there a human victim who could be harmed, physically or mentally by the action. The only person being harmed is the perpetrator themselves, perhaps by corrupting their moral character. As such, these instances of virtual sexual assault would be equivalent to classic vice crimes, where the criminalisation is justified by its effect on the perpetrator’s moral virtue. My reason for being sceptical about the wisdom of criminalising this conduct is not because I think the conduct in question is perfectly morally acceptable. On the contrary, I think someone who revels in committing acts of virtual sexual assault is doing something that is contrary to moral norms. But I think criminalisation is not an effective or appropriate strategy for addressing this kind of immoral conduct, for much the same reasons as I do not think that throwing drug addicts

---

<sup>26</sup> Abrams, K. ‘From Autonomy to Agency: Feminist Perspectives on Self-Direction’ (1998) 40 *William and Mary Law Review* 805-846; and Turkheimer, D. ‘Rape on and off campus’ (2015) 65 *Emory Law Journal* 101-146

in prison is an effective strategy for addressing their problems. What might be more appropriate would be to criminalise the makers or manufacturers of virtual reality technology for creating worlds that facilitate or encourage these types of virtual sexual assault, but in many instances this would be an overreach. Prohibition might conflict with freedom of expression and it might be unfair to the manufacturer if the user was the one who took advantage of the relative freedom/autonomy they were provided in the virtual environment to commit the relevant acts. In other words, unless the virtual environment was explicitly designed to facilitate fantasies of rape and sexual assault, and its doing so is completely gratuitous and devoid of any artistic or public merit, criminalisation of makers and manufacturers would also seem a step too far.<sup>27</sup> Regulatory interventions might be more appropriate, with the design of industry safeguards and standards being an obvious first step towards addressing problems that may arise. This is true not only in relation to cases involving wholly virtual victims, but all other cases too.

Second, I think that virtual sexual assault cases involving human victims (i.e. Types B, D, E, and F) could be justifiably criminalised, but the approach toward each type could vary. Some could be adequately captured by existing laws, others might require the creation of new laws. The reason for being more confident about the criminalisation of cases involving human victims is because they would seem to fall under the classic ‘harm to others’ rationale for criminalisation. Furthermore, we can be reasonably confident that people *are* harmed by virtual sexual assault. The harm would be primarily psychological as opposed to physical in nature,<sup>28</sup> but this is no different from real-world forms of sexual assault. Legal provisions on sexual assault do not typically require physical harm. Furthermore, when the assault is carried out on a victim wearing haptic technology, there would be physical effects on the body, making this type of virtual sexual assault almost indistinguishable from real-world sexual assault. This also suggests that forms of virtual sexual assault that are mediated through the use of haptic technologies would already be dealt with by existing laws on sexual assault.

---

<sup>27</sup> The computer game *Rapelay* might be a good example of a game that does go too far in this respect and is deservedly banned. But note that this raises something of a ‘dilemma’ that is widely discussed in the ethics of computer games, namely: why do we seem to have such different attitudes toward sexual violence in video games as opposed to non-sexual violence? See 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; and Luck, M. and Ellerby, N. ‘Has Bartel Resolved the Gamer’s Dilemma?’ (2013) 15 *Ethics and Information Technology* 229-233

<sup>28</sup> Distinguishing between the physical and psychological is philosophically problematic, but we can set aside debates about mind-body dualism for now.

In the absence of haptic technologies, things become a little bit more difficult. Sexually assaulting a person's on-screen avatar, might be covered by sexual harassment laws or obscene communication laws. Sexually explicit text messages, emails and videos that are shared and communicated online can, for example, already form the basis of a sexual harassment claim, so it is no great imaginative leap to suggest that sexually assaulting someone's avatar could do the same. The problem is that treating it as a type of harassment or obscene communication may not do justice to the seriousness of the offence. Sexual harassment is usually required to have taken place in the workplace, and so a good deal of virtual sexual assault might be missed unless the legal definition of harassment was expanded. Furthermore, you could argue that there is an important qualitative moral difference between making sexually explicit advances/communications via a virtual medium and actually sexually touching or penetrating someone's on-screen persona. If so, then we may need to create a new class of offence — virtual sexual assault — that does not require physical contact before it can be said to take place. Alternatively, we could take existing legal definitions of sexual assault and broaden them so that they no longer have a physical contact element to them.

### **3. Criminal Responsibility for Virtual Sexual Assault**

The previous section considered whether virtual sexual assault ought to be criminalised. Students of criminal law theory will have noted that in attempting to address this issue I focused on the so-called *actus reus* elements of an offence, i.e. what makes for a criminal act. I did not consider what it would take for someone to be justifiably held criminally responsible for those actions, the so-called *mens rea* or mental side of things. It is best to treat these as analytically distinct. The justification for criminalisation is grounded largely in the external properties of an action (i.e. its objective character and its effects on the world, including the minds of others). Nevertheless, there is little point in criminalising an action if you cannot justifiably hold someone responsible for it. Is this possible in the case of virtual sexual assault? To answer that, it is worth distinguishing between types of virtual sexual assault in which a human acted as the perpetrator and cases in which it was a wholly virtual agent that perpetrated the act.

#### *(i) Responsibility when there is a Human Perpetrator*

In cases involving human perpetrators, responsibility might seem like a non-issue. Humans have minds. They can act knowingly, intentionally, and recklessly, through their virtual representations, and hence they can satisfy the conditions of blameworthiness that are typically applied in the law. If they intentionally or recklessly perpetrated a virtual sexual assault, it would be right and proper to hold them responsible for doing so.

But we have to be careful with this. The responsibility for virtual acts is more complex than responsibility for real world acts. There are two reasons for this. The first is that the virtual environment itself might be structured and constrained by other agents (i.e. it might limit your possibilities for action), and this might call into question how free or autonomous the virtual action really was. The second is that people often distance themselves from their virtual personas. This psychological distancing might call into question their responsibility for what happens in the virtual world. Geert Gooskens has addressed both of these issues in his work<sup>29</sup> and I want to assess his arguments in what follows.

In relation to the first point, Gooskens argues that some freedom of action in a virtual environment is necessary if we are going to ascribe responsibility and blame to the actions of people in that environment. He points out that the virtual environments associated with traditional videogames don't always afford their players a great degree of freedom. Indeed, classically, games only allowed players to exercise a very narrow form of instrumental rationality. The games would have clearly-defined goals or objectives, and a limited set of pathways (maybe even only one pathway) to achieving those goals or objectives. Players had to follow those pathways in order to journey further in the virtual world. You could argue this doesn't allow for the necessary kind of freedom and so it would be wrong to ascribe responsibility for actions in that environment. This has repercussions for how we view responsibility for virtual sexual assault. For instance, in the controversial early-1980s videogame, *Custer's Revenge*, players had to commit a racially motivated sexual assault in order to succeed in the game. The game had one clearly defined objective — rape a Native American woman tied to a stake — and to succeed in the game you had to manipulate your character to avoid attacks and reach the woman. As such, this game only allowed the player to exercise a restricted form of instrumental rationality. Applying Gooskens analysis, you could argue that this wouldn't be sufficient for ascribing responsibility to the player for

---

<sup>29</sup> Gooskens, G. "The Ethical Status of Virtual Actions" (2010) 17(1) *Ethical Perspectives* 59–78

that virtual sexual assault. This example involves a virtual environment in which the ‘victim’ of the virtual assault is not ultimately controlled by a human, but presumably the same reasoning would apply if the victim was ultimately human: if the virtual environment allows for a restricted form instrumental rationality — e.g. if you have to sexually assault the other player’s avatar to progress in the virtual world — then you cannot rightly be held responsible for it.

The problem with this analysis, however, is that it ignores the decision by the player to play a game with such a clearly-defined goal in the first place. Surely freedom over that choice would meet the necessary condition for responsibility? If you sign up to a virtual world with such limited conditions of engagement, surely you can rightly be held responsible for following through on them? This seems right to me, but then it raises the question of whether there could ever be a sexual assault in such an environment since presumably the other human player would have signed up to those conditions of engagement too, which might suggest that they consented to the outcome. I discuss this issue in more detail in the final section.

In any event, much of this argument is moot since more recent video games and virtual environments allow for a much more robust type of ethical rationality to be exercised by the players. Increasingly, there are multiple ways of achieving a given game objective. For example, in a game like Grand Theft Auto you don’t usually have to mow down pedestrians or kill sex workers in order to complete the missions. You exercise freedom of choice to perform those actions. Furthermore, the kinds of virtual worlds discussed at the outset of this chapter, and the kinds that are foremost in people’s minds when they imagine ‘virtual reality’, are open-ended environments, ones in which you choose your own objectives. These environments would seem to meet the necessary condition of freedom that is needed for responsibility.

In relation to the second issue — that of psychological distance — Gooskens makes use of certain ideas from phenomenology, particularly the ideas of Edmund Husserl, to argue against responsibility for virtual actions. Phenomenology is a branch of philosophy that deals with the contours of subjective experience. Husserl was one of the leading phenomenologists and he developed a particular phenomenology of the “image world”, i.e. the world of representations, both real and fictional. These representations would include photographs or

paintings of real or fictional events and people, and also video game representations involving real or fictional agents and worlds. Gooskens argues that acting in a virtual world is a kind of *image-consciousness* (roughly: an intention to do things through the medium of an image). This is consistent with my earlier characterisation of virtual action as action-via-a-representation. Image-consciousness, according to Husserl, is characterised by the presence of an “as-if”-modifier, i.e. the situation is such that you are actually perceiving or intending something of type X, but it is ‘as-if’ you are perceiving or intending something of type Y. Consider the following example:

**Grandfather’s Photograph:** I am holding in my hand a small black-and-white photograph of my grandfather and looking directly at it. What am I perceiving? I am actually perceiving a grey two-dimensional figure on a piece of card; but it is as-if I am perceiving my grandfather.

That example should be pretty straightforward. The point is that the things that are perceived or intended through image-consciousness are “beings-as-if”, i.e. they are not really there, they are only virtually there.

The key Husserlian move is to argue that image-consciousness not only covers the objective correlates of our perceptions and intentions — i.e. the things we are looking at and the actions we are performing — but also the mental states underlying those experiences. In other words, when looking at my grandfather’s photograph, it is not just that the photograph is only an “as-if” representation of my grandfather, it is also that my subjective state of perception is an “as-if” state of perception. Husserl illustrates the extreme implications of this view with the example of a picture of human suffering. He argues that we cannot really feel real pity by viewing this photograph, we can only feel “as-if” pity. Gooskens argues that the same analysis can be applied to the mental conditions of responsibility for actions in a virtual world. In other words, he holds if you intend to commit a virtual rape or sexual assault through a virtual representation, it is not the case that you actually intend to do this, it is only ‘as-if’ you intend to do so. Consequently, we cannot really ascribe responsibility or blame for actions in a virtual environment.

This strikes me as being a deeply implausible view, particularly when you take into account the complex relationship between the real and the virtual outlined earlier in this

chapter. If it's the case that some virtual actions and events are every bit as real as their real-world equivalents, it is difficult to see how the as-if modifier can be rightly applied to their occurrence. If I take bitcoin from someone's online wallet, it seems fair to say that I have really stolen from them. It is not just 'as if' I have stolen from them. There is also ample evidence from psychology and neuroscience to suggest that Husserl's understanding of image consciousness is wrong. People's sense of ownership over their bodies can, and often is, projected out into the world of representations and artifacts. The famous 'Rubber Hand' illusion, for example, shows that people can incorporate a fake rubber hand into their phenomenological representation of their own body.<sup>30</sup> And as the philosopher Thomas Metzinger observes in his review of the evidence:

*"Owning" your body, its sensations, and its various parts is fundamental to the feeling of being someone. Your body image is surprisingly flexible. Expert skiers, for example, can extend their consciously experienced body image into the tips of their skis. Race-car drivers can expand it to include boundaries of the car; they do not have to judge visually whether they can squeeze through a narrow opening or avoid an obstacle – they simply feel it.*<sup>31</sup>

All of this suggests that the 'as if' modifier propounded by Husserl may not be present. If so, any attempt to deny responsibility for a virtual sexual assault by making the Husserlian move is likely to fall flat.

Interestingly Gooskens seems to be aware of the troubled relationship between the virtual and the real when it comes to the moral assessment of at least some kinds of virtual acts. He argues we can rightly feel moral discomfort at certain kinds of virtual acts, but we can't properly assign moral blame for their performance. In broad outline, his argument is very straightforward. He argues that we can rightly feel moral discomfort whenever a performer of a virtually immoral act blurs the boundary between their real-self and image-world-self. In other words, when the "as-if" modifier begins to lose its grip.

Gooskens uses some nice examples to flesh this out. First, he takes us away from the hi-tech world of video games and back to the more genteel and low-tech world of the stage play.

---

<sup>30</sup> Botvinick, M. and Cohen, J. 'Rubber Hand "Feels" Touch that Eyes See' (1998) 391 *Nature* 756.

<sup>31</sup> Metzinger (n 7), 75

Actors are the quintessential performers of virtual acts and usually we have no problem with them performing acts that would otherwise be highly immoral. Thus, we do not balk at Anthony Hopkins playing the sadistic cannibal Hannibal Lector. But there is some complexity to this. Usually, upon learning that an immoral act is part of a play or performance, we feel a great sense of relief. To use Gooskens example, imagine walking into your flat to find your actor roommate apparently forcing himself on a woman, but then seeing the director of the play issuing various directions from the couch. In this case, what would initially be a feeling of grave moral concern would shift to one of relief. Can the emotions ever run in the other direction, i.e. from a feeling of relief (or ease) to a feeling of discomfort? Gooskens argues that it can. Imagine if you are watching the stage performance of your roommate's play. Everyone is captivated by his intense and realistic performance during the rape scene. Afterwards, you congratulate him on this and he tells you that the reason why he was so intense and realistic is that he was actually sexually aroused during the scene. Gooskens suggests that in this case you would feel moral discomfort. Your roommate appears to have blurred the boundary between his real-self and his image-world-self in a most disturbing manner.

If we accept this analysis, however, don't we end up with the view that it is often going to be right and proper to hold people responsible for virtual sexual assault? After all, the distinction between the real-self and the image-world-self is an increasingly fuzzy one in most people's minds. In many instances, I would argue, it simply collapses. I don't feel separation from the online version of myself and any attempt to disclaim responsibility for my online actions would ring hollow to my peers. Gooskens claim is that we can only feel moral discomfort because the 'as-if' modifier prevents us from thinking that anything intrinsically wrong has occurred in virtuality. But this conclusion seems untenable, at least in the cases of virtual sexual assault in which we are interested, because they have more than a whiff of reality about them. As Gooskens himself concedes, if an action has some real world consequences it may still count as being intrinsically wrong. This would seem to be true of virtual sexual assault in which there is a human victim.

In sum, I would argue that the ascription of responsibility for human perpetrators poses no significant problems.

*(ii) Responsibility for Virtual Perpetrators*

What about cases in which the perpetrator of the sexual assault is wholly virtual? In the cases in which we are interested, these will involve a human victim, so there is the possibility of genuine harm and hence a justification for criminalisation, but the issues surrounding responsibility are more complex. If the wholly virtual agent committing the assault is specifically programmed by a human to perform such an act, then it will be easy to trace a line of responsibility back to that human programmer. If the wholly virtual agent has some degree of autonomy or artificial intelligence, if it is programmed by an organisation or team of humans, and if its behaviour is an unanticipated or unexpected outcome of its autonomy and intelligence, the situation might be a little more difficult. In that case, we may have a ‘responsibility gap’ opening up between the acts of the virtual agent and the decisions made by its programmers and creators.

This ‘responsibility gap’-problem has been widely discussed in the debate about autonomous robots. Indeed, the possibility of such a gap is probably the leading objection to the creation of robotic weapons systems.<sup>32</sup> Many legal and ethical scholars think that the use of such systems would undermine accountability and responsibility for acts of war. Similar discussions have also begun in relation to responsibility for the misdeeds of autonomous cars.<sup>33</sup> I would argue that when it comes to determining the responsibility for acts of virtual sexual assault committed by wholly virtual agents, we will end up following and learning from this debate.

Although I cannot do justice to all the positions defended in the existing ‘responsibility gap’-literature, I can at least give a sense of the solutions proposed to the problem. Many of them point out that the gap can be plugged by existing legal doctrines and concepts. It is not as though the law has never had to deal with problems arising from corporate agency or non-human agency before. There are already laws that allow for corporations to be held criminally responsible for certain actions, as well as laws that impose duties on humans to control non-human agents (such as animals). Recklessness and negligence are well-

---

<sup>32</sup> For the classic discussion of this problem, see Sparrow, R. ‘Killer Robots’ (2007) 24(1) *Journal of Applied Philosophy* 62-77; for a lengthy analysis of the issues, see the contributions on responsibility in Bhuta, N. Beck, S. Geiß, R. Liu, H.-Y. and Kreß, C. (eds). *Autonomous Weapons Systems: Law, Ethics, Policy* (Cambridge: Cambridge University Press, 2016).

<sup>33</sup> I discuss some of these issues in Danaher, J. ‘Robots, Law and the Retribution Gap’ (2016) 18(4) *Ethics and Information Technology* 299-309; see also Nyholm, S. ‘Attributing Agency to Automated Systems: Reflections

recognised standards of criminal responsibility for offences as serious as manslaughter. The recklessness standard states that someone can be held responsible for an act if they recognised the risk of its occurrence but took the risk anyway. The negligence standard states that someone can be held responsible for an act if they ought to have recognised the risk of its occurrence. Neha Jain has recently argued that both standards could be applied to the creators and/or deployers of autonomous weapons systems if and when those systems do something illegal.<sup>34</sup> A similar argument could be made in relation to creators and programmers of virtual agents who commit virtual sexual assault. As a last resort, it might also be possible to apply a strict liability standard to the creators and manufacturers, thereby obviating the need to provide any foresight or lack of care on their part. That said, we should be careful here. The use of strict liability standards in criminal law is highly controversial, and usually reserved for so-called regulatory offences. That may not be much of an issue, however, since the perpetration of sexual assault by a wholly virtual agent would probably require the creation of an entirely new category of offence. And, furthermore, it may be the case that civil law is better remedy to the kind of harm that is caused by such an act.

#### **4. Consent and Virtual Sexual Assault**

It is worth closing by reflecting on consent and its importance in our approach to virtual sexual assault. Consent is, according to many, a kind of ‘moral magic’.<sup>35</sup> It is what converts morally impermissible acts into permissible ones. This is particularly true in relation to sexual offences. In many instances, the absence of consent is baked-in to the definition of sexual assault and rape. It is what differentiates acceptable sexual activity from immoral sexual activity.

But what does it mean to consent to sexual activity? This is a topic that could fill many books.<sup>36</sup> In broad outline, consent is a subjective attitude toward the performance of an act (roughly: an agreement that the act be performed which is not necessarily equivalent to a desire) that is communicated through objective behaviour. So you consent to sexual activity

---

on Human-Robot Collaborations and Responsibility-Loci’ (2017) *Science and Engineering Ethics* DOI: 10.1007/s11948-017-9943-x

<sup>34</sup> Jain, N. ‘Autonomous Weapons Systems: New Frameworks for Individual Responsibility’ in Bhuta et al (eds) *Autonomous Weapons Systems* (n x)

<sup>35</sup> I take this from Hurd, H. ‘The Moral Magic of Consent’ (1996) 2(2) *Legal Theory* 121–146.

<sup>36</sup> See, for example, Wertheimer, A. *Consent to Sexual Relations* (Cambridge: Cambridge University Press, 2003); Westen, P. *The Logic of Consent* (London: Ashgate, 2004);

when you subjectively agree to its performance *and* you communicate that subjective agreement through your behaviour. When it comes to determining the law and ethics of sexual consent, the communication of consent is all important. We can never know what a person is really thinking. We can only go by their outward behaviour. What objective behaviour counts as an effective communication of consent? Must you clearly and unambiguously state your agreement to the sexual activity? Or will a raised eyebrow and alluring glance be enough? Many of the historical problems with sexual assault have centred on the tendency to infer consent from outward behaviours that are highly ambiguous and frequently misinterpreted. And much of the modern debate about sexual assault is concerned with the desirability of moving away from such misinterpretations to affirmative consent standards, i.e. a legal norm of consent that requires clear and unambiguous signals of assent to sexual relations.<sup>37</sup>

Given the centrality of consent to our modern conception of sexual assault, it stands to reason that consent will take a similarly central location in our understanding of virtual sexual assault. People already engage in sexual behaviours in virtual environments and will continue to do so in the future. And, just as is the case in the real world, developing a set of norms around which objective behaviours and signals ‘count’ when it comes to consent to virtual sexual activity will be crucial when it comes to distinguishing the permissible forms of virtual sexual activity from the impermissible forms.

For the most part, the debates about real-world consent and virtual-consent consent will and should follow similar tracks. It is possible for people to communicate their agreement to sexual activity through their virtual characters and avatars, either by text or speech, and it is desirable to create a set of norms that encourages them to do so in clear and unambiguous terms (if they wish to engage in virtual sexual activity). In this respect, I would argue that if we are encouraging a shift toward affirmative consent standards in the real world (and it seems as if we are) we should encourage the same shift the virtual world.

Nevertheless, despite my enthusiasm for treating the two domains in the same way, there may be ways in which consent-relevant communications between sexual partners in virtual

---

<sup>37</sup> Little, N. ‘From no means no to only yes means yes: the rational results of an affirmative consent standard in rape law’ (2005) 58(4) *Vanderbilt Law Review* 1321-1364; Dougherty, T. ‘Yes means yes: communication as consent’ (2015) 43(3) *Philosophy and Public Affairs* 224-253

worlds could be more problematic than consent-relevant communications in the real world. I want to close by discussing two such problems.

The first is the problem of deception. Many legal systems accept that certain forms of deception undermine consent to sexual relations, and there are ethicists who argue that nearly all forms of deception should.<sup>38</sup> In the legal system with which I am most familiar – England and Wales – deception as to the ‘nature and quality of the act’, and deception as to identity of the perpetrator (if this involves impersonation of someone known to the victim), are so-called ‘irrebuttable presumptions’ of non-consent in cases of sexual assault.<sup>39</sup> In other words, if it can be shown that a defendant deceived a victim as to the sexual nature of the activity in which they engaged, or if they can be shown to have impersonated someone known to the victim, it is impossible for them to argue that they had a reasonable belief in the consent of the victim. Other forms of deception may also undermine consent, but do not provide an irrebuttable presumption of non-consent.

The two forms of deception that provide irrebuttable presumptions obviously do occur in the real world. Deception as to the nature and quality of a sexual act typically arises where some doctor or medical official deceives someone into undergoing an examination that provides the perpetrator with sexual gratification but is presented to the victim under the pretence of medical necessity.<sup>40</sup> It can arise in other contexts too. The case of *R v. Devonald* (discussed previously) was successfully prosecuted because the court held that there was deception as to the nature of the act in question: the victim thought his masturbation was for mutual gratification when it was actually for embarrassment and revenge. Impersonation of someone known to the victim is relatively rare in the real world, and the reported cases usually involve someone climbing into the bed of a victim when it is dark and difficult to see who is really there. Nevertheless, it does happen.<sup>41</sup>

I mention these examples because one concern we may have about virtual sexual activity is that it is more susceptible to consent-undermining forms of deception than real world forms. Since people are not interacting ‘flesh on flesh’, but instead through virtual

---

<sup>38</sup> Dougherty, T. ‘Sex, Lies and Consent’ (2013) 123(4) *Ethics* 717–744

<sup>39</sup> See section 76 of the Sexual Offences Act 2003

<sup>40</sup> See, for example, the cases of *R v. Tabussum* [2000] 2 Cr App R 328 and *R v. Green* [2002] EWCA Crim 1501

<sup>41</sup> See, for example, the case of *R v. Elbekkay* [1995] Crim LR 163 (CA).

representations, it could be much easier, for example, for one party to impersonate someone known to a victim. The usual techniques of identity-checking that apply in the real world are more easily circumvented in virtual environments. The account associated with your virtual avatar could be hacked or you might carelessly share the relevant information or log-in details with someone else. If the poorly-named practice of ‘fraping’ is anything to go by, impersonating another may end up being relatively common. Similarly, deception as to the nature and quality of a virtual sexual act might be more common in virtual interactions than real-world interactions. The case of *R v. Devonald* already suggests as much, showing how easy it is for there to be deception as to the purpose of what is taking place when the parties to the act are not in direct physical contact with one another. More fancifully, some people have suggested that sexual versions of the ‘Chinese Room’ thought experiment could be realised in virtual environments.<sup>42</sup> The Chinese Room thought experiment was originally proposed by John Searle as a way of arguing that computers could never think in the same way as humans (never have consciously directed thoughts). The experiment asks you to imagine a man inside a room with a big book of instructions. Chinese symbols are fed into the room and the big book of instructions tells the man how to respond to each and every one of these symbols. Unbeknownst to the man inside the room, the net effect of this is for somebody outside the box to be able to carry out a conversation (in Chinese) with the room. One party – the man inside the room – is completely ignorant as to what is taking place. A similar level of ignorance could arise with virtual sexual activity. One human user of a virtual environment could input instructions and commands to their avatar that, unbeknownst to them, have the net effect of sexually assaulting somebody else’s character, or vice versa.

The susceptibility of virtual sexual activity to these forms of deception probably does not have any major ethical or normative significance. If we think that certain forms of deception vitiate consent in the real world then they should also to vitiate consent in virtual worlds. That said, the relative ease with which deception might occur in virtual worlds suggests that there is added risk to virtual sexual encounters, and perhaps a need for higher standards or norms of consent, e.g. additional checks and verifications of identity.

---

<sup>42</sup> Sanberg, A. ‘The Chinese pleasure room: ethics of technologically mediated interaction’ *Practical Ethics Blog*, 20 September 2016 – available at <http://blog.practicaethics.ox.ac.uk/2016/09/the-chinese-pleasure-room-ethics-of-technologically-mediated-interaction/> (accessed 8/14/2017)

A second problem with consent in virtual worlds arises from the possibility of resisting or avoiding unwanted sexual contact. Historically, many jurisdictions have included a ‘force’ requirement in their definitions of sexual assault and rape. The idea was that a sexual act was only non-consensual or unwanted if the victim actively resisted its performance and so required the perpetrator to use force to complete the act. Active resistance could manifest in many ways, e.g. by trying to escape the perpetrator’s clutches or by actual physical violence. The force requirement placed a very onerous duty on victims of sexual assault, and was not, according to many critics of the law, fair to the lived experience of assault victims, many of whom could freeze or become psychologically paralysed by the trauma of the assault. Fortunately, there has been a general shift away from the force requirement in the past half century – though it lingers in some places – and an acceptance of the fact that victims of sexual assault do not need to actively resist unwanted sexual advances in order to convince a court that they are victims of a sexual crime.

There is an interesting normative and socio-cultural question to be asked as to why this shift has happened. Is it simply because we have become convinced that consent should be an active, affirmative agreement and that it should never be inferred from outward passivity? Or is it because we feel it is impractical (and hence unfair) to place a duty on victims of sexual assault to actively resist their attackers? The answer may make some difference when it comes to assessing the moral seriousness of virtual sexual assault.

One common rebuttal to the claim that virtual sexual assault should be taken seriously – ‘common’ in the sense that whenever I have discussed the topic it has been brought up – is that it is much easier to resist or avoid virtual sexual assault than it is to resist or avoid real world sexual assault. In the real world, it may be hard to get away from your attacker. Their physical presence may cause you to freeze up or shut down and hence become passive and non-resistant. In a virtual world, it will be much easier to get away: just take off your headset or switch off your computer. You are instantly free from the environment in which the assault was taking place.<sup>43</sup> Consider once more the story of Jordan Belamire, discussed in the introduction to this chapter. While we may sympathise with what happened to her and agree that the actions of her perpetrator were abominable, she could have quickly escaped from his

---

<sup>43</sup> Jaclyn Seelagy notes that the degree of customizability of the user experience and the manufacturer’s safeguards against unwanted interactions, would also be relevant factors in thinking about both consent and the

clutches by taking off her headset. The upshot of this is the following: If we think that the practicality of resistance is what matters when it comes to non-consent, and if we agree that it is much easier to resist in a virtual environment, we might think that a force requirement should be included in the definition of what counts as a virtual sexual assault. We might go even further and start to think that there are virtual equivalents of ‘asking for it’. So if someone enters a virtual world (e.g. whilst, say, wearing a haptic sexsuit), in which virtual sexual assault is known to be common, or if they fail to absent themselves from an environment when such behaviour becomes apparent, they were ‘asking for it’ and do not deserve our sympathy. In other words, we might argue that it is only those cases where someone is locked into a virtual environment against their will, that are worth taking seriously.

There are three reasons why I think we should reject such a view. The first is simply that even if the practicality of resistance is the factor that has moved us away from a force requirement, there is no reason to think that the psychological forces that make resistance impractical (and hence unfair) in the real world would not apply to the virtual world. Jordan Belamire’s description of her virtual reality groping points out how realistic, and traumatising, it was to be immersed in the virtual world, suggesting the psychological distance between her virtual self and real self was not all that great. Second, and more importantly, the practicality of resistance should not be the decisive factor in any event. The reason for rejecting a force requirement should, in my opinion, lie in the preferability of open, affirmative consent standards. These should apply to the virtual world as well and we shouldn’t read consent into lack of virtual resistance. Finally, because the boundary between the virtual and the real is increasingly blurred, and because we now live much of our lives through virtual world representations of ourselves, it seems unfair to insist that victims of virtual sexual assault should absent themselves from virtual worlds when things start to get unpleasant. Just as women should not be required to modulate their dress for fear of being sexually assaulted on city streets, so too should people not be required to flee from a virtual assaulter. The duties and responsibilities should fall on the perpetrator, not the victim.

## **5. Conclusion**

---

appropriate subject of responsibility in the event of a non-consensual interaction. See Seelagy, V. ‘Virtual Violence’ (2016) 64 UCLA Law Review Discourse 412-433

Virtual sexual assault is an emerging phenomenon and is likely to become more prevalent in the future as we continue to live more of our lives in virtual environments. It is important that we start to think about the seriousness of this type of assault, its appropriate classification and the desirable legal and ethical responses to its occurrence. Building on initial work done by others, particularly the work of Litska Strikwerda, I have tried to provide a framework for thinking about these important issues. I have not tried to defend any particular view, rather I have presented a series of arguments about how we should understand and respond to this phenomenon. The main conclusions of these arguments can be summarised as follows.

First, I have argued that virtual sexual assault needs to be classified along three dimensions (victim-perpetrator dimension; avatar-immersive interaction dimension; and human user-virtual character dimension). Using this classificatory scheme we can identify six major types of virtual sexual assault that are worth taking seriously either because they involve a human perpetrator, or human victim, or both.

Second, I have argued that virtual sexual assault may warrant criminalisation if (a) it has properties that make it 'real' and (b) those properties would ordinarily warrant criminalisation. Following the work of Philip Brey, I argued that there is a complex relationship between the virtual world and the real world, and that certain actions and events that take place in a virtual world are every bit as real as their real-world equivalents. This is true when the virtual acts have real world consequences and/or when they are social, as opposed to physical, in nature. This means that some instances of virtual sexual assault are real, because virtual sexual assault can have real world consequences and there are some grounds for thinking that certain aspects of sexual activity are social, as opposed to physical, in nature. What is more, the properties of virtual sexual assault that are real are ones that make it apt for criminalisation, because it can involve harm to others. The actual approach to criminalisation is complex, with certain forms of virtual sexual assault likely to be covered by existing laws, and others necessitating new laws.

Third, I have argued that it is possible to hold people responsible for virtual sexual assault. Where the virtual assault is perpetrated by a human acting in a virtual environment, there is a responsible moral agent to whom responsibility can be ascribed. Any claims to the effect that there is some psychological/moral distance between the real-self and the virtual-self that make such responsibility ascriptions impossible should be resisted. Furthermore, in

cases where there is no immediate human perpetrator, it may be possible to hold the designers and manufacturers of the virtual environment liable for the actions that take place within it. In this respect, there is much to be learned from the ‘responsibility gap’ debate in relation to robotic weapons systems.

Finally, I argued that our approach to consent in the case of virtual sexual assault should mimic our approach to consent in the case of real world sexual assault, although we should acknowledge a greater potential for deception in virtual sexual interactions, while at the same time rejecting the claim that the ease with which someone can resist virtual sexual assault is a relevant factor in its normative assessment.