

Virtual Worlds and Moral Evaluation  
Jeffrey Dunn  
(please cite published version in *Ethics & Information Technology*)

## 1. Introduction

Consider the multi-user virtual worlds of online games such as *EVE* and *World of Warcraft*, or the multi-user virtual world of *Second Life*. In such worlds, players create avatars, which then have extensive freedom to do what they please in virtual worlds. The players can socialize, set up businesses, buy houses, etc., all in the virtual world.<sup>1</sup> Many interesting questions suggest themselves when considering these virtual worlds. In this paper, I focus on one question in particular. Assume that it is wrong for me to follow a person home and then take his or her property by force. Suppose I perform a similar action, but via my avatar in one of these virtual worlds. Have I then done something wrong?<sup>2</sup>

This is a particular example of the kind of question on which I will focus in this paper. But the question I will focus on is both more general and more specific than the particular example would suggest. It is more general because I will not focus solely on *theft*. It is more specific, because I will focus solely on actions that affect other avatars, which are controlled by real people. That is to say, I will be focusing on actions that meet the following description: actions performed within a virtual world (via an avatar), which affect other avatars in the virtual world and would be wrong *if* the situation in the virtual world were *real*. The question about such virtual actions is: are they wrong?<sup>3</sup>

There are three main types of answers that can be given to this question. What I'll call the *Asymmetry Thesis* maintains that such actions are never wrong.<sup>4</sup> The *Symmetry Thesis* maintains that such actions are always wrong. A third option is to deny both theses. According to this third view, some such actions are wrong and some are not.<sup>5</sup>

This paper will focus on the Asymmetry Thesis. My ultimate conclusion is that the Asymmetry Thesis is false: some actions that can be performed in virtual worlds that would be wrong if performed in the real world are themselves morally wrong. I'll further argue that the Asymmetry Thesis is closely related to the possibility of what Edward Castronova (2004)

---

<sup>1</sup> There are differences between these examples. *World of Warcraft* and *EVE* have built-in objectives that players try to accomplish, whereas *Second Life* has no built-in objectives. Nevertheless, they are similar in having persistent virtual worlds where players have freedom to perform many different actions and interact with other players. These similarities are what will be important in what follows.

<sup>2</sup> A particular example of the sort of event I'm considering is the famous (infamous?) assassination and theft that took place in 2005, within *EVE* by members of The Guiding Hand Social Club. For details about this case, see Francis (2005), Rossignol (2005), and Craft (2007).

<sup>3</sup> Brey (1999) and Gooskens (2010) each address a similar issue, but with respect to games that have only one human player, the rest of which are controlled by the computer. Bradley & Froomkin (2004) consider the relationship between real rules and virtual worlds, but the emphasis is on using virtual worlds to assess the efficacy of proposed legal rules. Strikwerda (2012) has recently focused on virtual theft and asked when real-world legal action should be taken in such cases. For other work on this topic, see Powers (2003), Wolfendale (2007), and Johansson (2009). Some of these will be discussed below.

<sup>4</sup> Powers (2003, p. 198) suggests that he ascribes to the Asymmetry Thesis. This is despite the fact that he argues that some virtual actions—though not the actions in the virtual worlds of things like *Second Life* or *World of Warcraft*—can be wrong.

<sup>5</sup> Craft (2007) argues that actions like the one that occurred in *EVE* in 2005 are wrong. It's unclear, however, whether he ascribes to the Symmetry Thesis or simply the third view. Wolfendale (2007) appears to reject the Asymmetry Thesis, but it is pretty clear she also rejects the Symmetry Thesis. Johansson (2009) also rejects the Asymmetry Thesis. It is less clear what his views are about the Symmetry Thesis.

has called *closed synthetic worlds*. These are virtual worlds designed solely for play, where nothing that occurs in such worlds matters. Castronova argues that it would be valuable to have such virtual worlds and that we can in fact create them. However, I'll argue that if the Asymmetry Thesis is false, then there can be no closed worlds. Finally, I will argue that even if the Asymmetry Thesis is false, there still seems to be *some kind* of asymmetry between actions performed in virtual worlds and their real-life counterparts. I conclude by attempting to explain this limited asymmetry between virtual and real actions.

Here, then, is what's to come. Section 2 considers the best arguments for the Asymmetry Thesis. Section 3 shows how the Asymmetry Thesis is related to Castronova's closed worlds. Section 4 argues positively against both the Asymmetry Thesis and the possibility of closed worlds. Section 5 draws on the arguments from Section 2 to try to explain the limited asymmetry that there seems to be between virtual actions and their real-world counterparts.

Before jumping in, a small defense of the focus on this question and these answers is needed. One way of thinking about this issue is to abstract away from the fact that things are being done *via* virtual worlds. When I crash my car into someone's property I have used a tool to cause an effect—in this case, damage to someone's property. Similarly, one might see these virtual worlds as merely complex tools that enable one to cause certain effects. What is up for moral evaluation, then, is my actual action of making a certain keystroke. We could say that a certain keystroke is morally objectionable if it causes someone harm, or if it does not maximize utility, or if it violates someone's rights, or if it transgresses the Categorical Imperative.

I think there's a lot that's correct about that kind of view. However, taking such a perspective on the issue may perhaps hide certain interesting regularities from our view. We are often interested in morally evaluating *types* of actions. For instance, hitting a person against that person's wishes is a type of action that is usually wrong. By focusing on types of actions, rather than particular action *tokens*, interesting moral patterns come to light. Having some rough-and-ready generalizations about the moral status of types of actions can also be very useful.<sup>6</sup> This motivates the focus on the theses mentioned above, which concern a certain type of salient action: the type of action where you do something in a virtual world via your avatar that would be morally wrong were the virtual world actual.

## 2. Arguments For The Asymmetry Thesis

### 2.1 Virtual Arguments

For concreteness, focus on a particular virtual action that meets the description in the Introduction. For example, the virtual theft of some virtual agent's virtual property (i.e., property within the game) plausibly meets the description. Alternatively, lying to or betraying a virtual friend within the virtual game world would meet the description, too. Here's a natural response to the allegation that such actions are wrong: *but none of it is real!*

Let's see if we can make this kind of response a bit more precise. Here's one way it might go.

Virtual Argument 1: The actions are not wrong because the theft/betrayal took place in an online environment, and actions done online are not capable of being wrong.

---

<sup>6</sup> It is often hard to know what effects a particular action has, or what one's motivation is in performing it. It is often much easier to simply recognize what type of action one is performing.

We can quickly reject this argument. There are clearly actions that can be done online that are nevertheless morally wrong. If I am clever enough, I can use online banking systems to steal your (real!) money, by transferring it virtually from your account to mine. The actions I take are all online, and yet the theft is real, and clearly wrong.<sup>7</sup>

The problem with Virtual Argument 1 might suggest another, more subtle one. So, consider:

Virtual Argument 2: The virtual theft is not wrong because it involved the theft of virtual property, not real property. If nothing real was stolen then there was no theft, and so no morally wrong action. The virtual betrayal is not wrong because it did not betray a real person, but rather an avatar, a virtual person.

This is a more formidable argument. But, I think that it, too, fails.

When the argument is about virtual theft, it fails because virtual property can be property—and can be stolen—just as much as can “real” property. Take an example. Suppose you have spent time and effort composing pieces of music, which are in—and have always been in—digital format. Suppose you store these digital compositions online. If I hack into your account, and transfer those files to my computer, deleting them from your online storage, I have stolen from you. There is no moral difference between this kind of theft, and theft of a more traditional sort where, say, I steal your CD collection. But this shows that the above argument is mistaken. The wrongness of theft does not depend on the stolen item being a “real” object (whatever that means). And this means that the above argument does not establish the Asymmetry Thesis.<sup>8</sup>

When the argument is about virtual betrayal, it fails because although the avatar is in *some* sense virtual, it is also the representation of a real person. As Jessica Wolfendale (2007) argues, participants in these virtual worlds identify with their avatars. Avatars are not just objects; they are the way in which a person presents herself to the online community. She sums up the idea as follows: “Avatars are therefore far more than mere online objects manipulated by a user. They are the embodied conception of the participant’s self through which she communicates with others in the community.” (Wolfendale 2007, p. 114)<sup>9</sup> Virtual Argument 2, it seems, fails.

Consider one final argument in favor of the Asymmetry Thesis for virtual worlds, which brings out the main problem with this class of arguments. The argument draws an analogy between the Asymmetry Thesis and a similar thesis that seems to be true of dreams. That is, consider whether dreaming that one does a morally repugnant action is actually wrong. Many think the answer is *no*. Julia Driver (2007) defends such a position, thus defending something like the Asymmetry Thesis for dreams. She summarizes her argument as follows:

In the end, my view is that the dream should be viewed as another and very different context. Without systematic positive or negative effects they have no actual moral significance. ...dream immorality may have epistemic significance, as a sign of

---

<sup>7</sup> Hunter & Lastowka (2004) note this, too.

<sup>8</sup> Indeed, as others have noted (Hunter & Lastowka 2004; Fairfield 2005; Craft 2007), the objects in virtual worlds meet the standard definitions of things that have economic value: the players who accumulated those in-game artifacts invested time in acquiring them, they regret their loss, they have utility within the virtual world, they can be bought and sold using real currency, etc. For a recent argument for the claim that virtual theft can be real theft, see Strikwerda (2012).

<sup>9</sup> In a different way, (Munn *forthcoming*) emphasizes the way in which an avatar is more than just an online object. He argues that the fact that participants have avatars that represent themselves allow them to *do* things with each other in a virtual world, rather than just discuss things with each other. He sees this as important to understanding how true friendship can develop in virtual worlds, in a way that it cannot in mere chat rooms.

something wrong with a person. But there is no dream immorality *per se*. And the same carries over to the other non-veridical contexts. (p. 21)

Driver's view is that immoral "actions" in dreams are not morally significant, because morally wrong "actions" in dreams are not regularly connected with negative effects. Driver extends this claim to other non-veridical contexts. Online virtual worlds might seem to be a prime example of such a context. This can be leveled into an argument for the Asymmetry Thesis as follows:

Virtual Argument 3: Virtual actions in virtual worlds are not wrong because these "actions" are like the "actions" one takes within one's dreams. Just as there is no dream immorality, there is no virtual world immorality either.

However, the analogy is not apt. Although dreams and virtual worlds are both non-veridical, there is a key difference between them. One important feature of virtual worlds is that virtual actions are systematically connected up with virtual consequences in a regular way. If they *weren't* like this, it is hard to imagine anyone wanting to spend any time in such worlds. But since things in the virtual worlds are valuable to those that participate, virtual actions become systematically connected with real consequences. This is an important difference between dreams and virtual worlds.<sup>10</sup>

Although I certainly haven't considered every argument for the Asymmetry Thesis based on the virtual nature of virtual worlds, I think it is plausible that the mere fact that an action takes place in a virtual world does not provide grounds for the conclusion that the action is not morally wrong. The claim, "But it's not real!" doesn't establish the Asymmetry Thesis.

## 2.2 Game Arguments

The three arguments above try to show that something about the virtual nature of the virtual worlds is sufficient to show that actions within those worlds are not morally wrong. I think these arguments fail. But consider another natural response to the allegation that such actions are wrong: *but it's just a game!*<sup>11</sup>

Consider an argument that defends the acceptability of virtual theft by appealing directly to the fact that it occurred within a game:

Game Argument 1: Theft or betrayal in a virtual world is not morally wrong because virtual worlds are associated with games and actions within games aren't right or wrong.

This argument fails. For it is certainly possible to do morally wrong things within games. For instance, suppose you and I are playing basketball and I decide to punch you in the face to get the ball. Not only have I broken the rules of basketball, I have also done something that is (plausibly) morally wrong. Just because an action takes place within the context of a game, it doesn't follow that it is not morally wrong. This argument is much too quick.

---

<sup>10</sup> A better analogy to virtual worlds, perhaps, would be the "shared dreams" that are featured in the film *Inception* (Nolan 2010). But it's not clear that actions done to others in such shared dreams really fail to be immoral.

<sup>11</sup> Not all multi-user virtual worlds are best thought of as games. For instance, Warburton (2009) does not describe *Second Life* as a game. Nevertheless, *Second Life* does have rules (*Community Standards* n.d.) and many people enter the virtual world of *Second Life* voluntarily and for enjoyment. Thus, it is appropriate to see several of the arguments in this section (especially Game Argument 2 and 3) as applicable to the virtual world of *Second Life* even if *Second Life* should not be counted as a game.

But a more subtle version of that argument is available. One might claim that an action within a game can only be morally wrong if it is against the rules of the game. In the basketball example, punching a player in the face breaks the rules of basketball. Perhaps this explains why it is wrong to punch someone during a basketball game. In fact, something like this idea has been enshrined in Korean law. In Korea, acts of virtual theft within MMOGs are criminally punishable but only if the theft is achieved by violating the rules of the game (Castronova 2004). This law suggests that there is nothing wrong with virtual actions if, in performing these actions, one adheres to the rules of the game. This, then, suggests:

Game Argument 2: If an action within a virtual world is in accordance with the rules of the that world, then the action is morally acceptable. Thus, if virtual theft or virtual betrayal is in accordance with the rules of the world, then the virtual theft or virtual betrayal is not morally wrong.<sup>12</sup>

This argument suggests something interesting. It suggests that we can create spaces (actual or virtual) where the standard moral norms do not apply. Instead, the only obligations are to follow the explicitly stated rules. There is precedent for something at least similar to this. Many have argued that in *war* some of the standard moral norms do not apply.<sup>13</sup> However, unlike the example of morality in war, this argument suggests that by crafting the right set of rules for the virtual world, we can make it the case that an action that is normally wrong, fails to be wrong.

Before evaluating this argument, note that it may fail as stated, because it seems to rely solely on explicitly stated rules. There are surely examples of actions that one can take within games that do not violate any explicit rules and yet are still wrong. For instance, if a defensive football player has concealed a gun in his uniform and then proceeds to shoot an opponent, this violates no rules of football. The rules of football just don't say anything about this. However, such an action is clearly wrong, even if the defensive player shot his opponent solely for an in-game reason, e.g., to stop his opponent from scoring. But perhaps there is way out of this problem. Perhaps we could say that all the normal rules of morality obtain in a game *unless* the rules of the game state that certain of these normal rules fail to obtain. The argument would then maintain that an action in a game is not wrong so long as it is in accordance with the rules of the game.

So, what should we make of this argument? Does it show that we could have a virtual world with the appropriate rules where no actions are morally wrong? I don't think it does, for it seems that games can have immoral rules—rules that permit immoral actions. For example, consider a the type of game portrayed in Stephen King's novel *The Running Man* (written under the pseudonym 'Richard Bachman'). In this game the contestant's goal is to stay alive and is given a 12-hour head start before the contestant is chased by "hunters" who attempt to kill him. The contestant is given extra monetary rewards the longer he stays alive and is also given bonuses for killing law enforcement officers or hunters. Consider the action of killing a law enforcement officer by the contestant, or the action of killing the contestant by the hunter. Such actions are permitted by the rules of the game. Yet it seems clear that such actions are

---

<sup>12</sup> Both Kerr (2008) and Hunter & Lastowka (2004) note that something similar to this argument is upheld by US courts. Kerr sums it up nicely, saying that from a legal perspective, "the rules of the game trump" (p. 421). That is, if an action is performed within a game that would be illegal outside the game, the courts do not treat this action as illegal so long as the action violates no rules of the game. Of course, the courts are concerned with *legal* rightness and wrongness, not *moral* rightness and wrongness. But the form of the argument is the same.

<sup>13</sup> For some contemporary examples of this, see Walzer (2006) and Orend (2006).

still morally wrong. Though such actions are in accordance with the rules of the game, this is irrelevant since the game itself is a morally bad one.

In response to this, one might note that I have failed to take into consideration one very important feature of games. In games, the players *consent* to play and to play by the rules. Perhaps this can provide an argument for the Asymmetry Thesis.

Game Argument 3: In games players consent to play by a certain set of rules, and if a player consents to a certain set of rules, then there is nothing morally wrong with someone else adhering to those rules in their treatment of the player. Thus, if a virtual world permits theft or betrayal and the participants consent to these rules, there is nothing morally wrong with such betrayal or theft.

Game Argument 2 focused on the rules of the game, but left it mysterious how the rules could magically transform the actions performed in accordance with them. Consent is supposed to fill the gap in that argument. In other situations consent seems capable of transforming the moral status of actions, so this explanation appears to be of the right sort to account for the seemingly magical way in which the rules of a game could transform the moral status of actions.<sup>14</sup>

So, can consent to certain rules do all that it's being asked to do here? First note that something like this certainly seems to work *sometimes*. Boxers do things to each other in the ring that would not be acceptable outside the ring. And this is because they both agree and consent to be bound by the rules of boxing. But can consent to the rules of the game establish the Asymmetry Thesis?

To consider this question, first note that although many have been willing to grant that consent has some transformative power, most have thought it to be limited in its power.<sup>15</sup> That is, perhaps mutual consent to a rule that says you can hit me in the face can make hitting me in the face no longer morally wrong. But it is consistent with this that mutual consent to rules that say you are allowed to eat me does not make it the case that eating me is morally acceptable.<sup>16</sup> Now, the fact that consent can't make *every* action morally acceptable doesn't show that it cannot make every action *within a virtual world* morally acceptable. But it does suggest that consent to the rules of a virtual world may not render every in-game action morally acceptable. To make this clearer, consider again the example of the game in *The Running Man*. We can perfectly well imagine that game being played where all the players consent to the rules of the game. Nevertheless, it still seems to be a morally bad game where immoral actions are performed. For this reason, it seems that consent cannot render every virtual action performed in a virtual world morally acceptable.

There are also some specific worries about the *kind* of consent to the rules that can be given by those who participate in virtual worlds. It is, for instance, not always clear what the rules of a particular virtual world are. This is especially problematic if some of the rules of a virtual world are unwritten. And if you do not know what the rules of a virtual world are, it is hard to see how you could consent to such rules. Perhaps more troubling than this is the *kind* of consent that plausibly takes place within virtual worlds. Consider a particular heinous action that could occur within a virtual world. Perhaps it is some sort of abuse by one avatar against

---

<sup>14</sup> For example, (Hurd 1996) talks about the “moral magic” of consent.

<sup>15</sup> For instance, Wertheimer (2000, p. 560) writes: “It is a mistake to think that consent always works “to make an action right when it would otherwise be wrong,” if “right” is equivalent to “morally worthy” or “justified.” ”

<sup>16</sup> For example, consider the reactions to the case of Armin Meiwes, the German citizen who in 2001 killed and ate a man who had apparently consented to such an arrangement. For some details, see Lander (2003).

another. One way in which this action could be consented to is if the two avatars consented to this particular instance of abuse within the game. But that's not the kind of consent that we are considering here. The kind of consent that we are considering is where both players consent to the rules of the game (which in this case permit this kind of abuse), and then where the specific action itself need not be consented to. It is one thing to consent to a specific action, but it is quite another to consent in this general way. It's not clear that this general, open-ended kind of consent always works to transform the moral status of actions, even if one thinks that consent to specific actions does.<sup>17</sup> Thus, the above considerations suggest that even when players consent to the rules of a virtual world we do not have a virtual world that satisfies the Asymmetry Thesis.

### 2.3 Play Arguments

We have just tried arguing for the Asymmetry Thesis based on the fact that virtual worlds are games. A related, though distinct, strategy is to argue for the Asymmetry Thesis in virtue of the fact that virtual worlds involve *play*.<sup>18</sup> The basic idea is that play is itself a valuable thing. Perhaps this is because play is an instrumental good: it has positive consequences. Or perhaps it is because play itself is an essential aspect of a human life that is going well. Further, virtual worlds have opened up a new way for humans to engage in play.<sup>19</sup> With this basic idea granted, we can argue for the Asymmetry Thesis in various ways.

Play Argument 1: Play is of significant value. Further, every action performed within a virtual world helps to sustain the contributes to the value of play. Since this value is so great, no matter the negative consequences of a particular virtual action, the net value of the consequences of the virtual action are positive. Thus, every virtual action has positive consequences and so no virtual action is morally wrong.

This argument clearly adopts a consequentialist picture of rightness and wrongness, where the rightness and wrongness of an action is determined by the value of the action's consequences. Some may not grant that picture. But we can put this aside since the argument has a much clearer flaw: however great the value of play is, it just isn't right that this value always trumps the negative consequences of a particular virtual action. Further, it is doubtful that every virtual action contributes to the value of play.

However, this clearly flawed argument leads to a more nuanced one:

Play Argument 2: Play is of significant value. Particular virtual actions that contribute to play may have net negative consequences, but are morally acceptable because such actions are in general valuable. Similarly, a particular act of telling the truth may have net negative consequences but is morally acceptable because truth-telling in general has positive consequences. The same thing is true for virtual actions that contribute to play. Thus, such actions are always morally acceptable.

If Play Argument 1 assumes a straightforward consequentialist picture of rightness and wrongness, this argument assumes something more akin to a rule consequentialist picture.<sup>20</sup> According to this picture, an action's moral status depends not on the consequences of that particular action, but rather on the consequences of a rule that permits that kind of action. In

---

<sup>17</sup> In a discussion of BDSM, Athanassoulis (2002) maintains that consent can transform only in specific instances. Nielsen (2010), on the other hand, maintains that consent can be efficacious even if it is given in a more general, open-ended way.

<sup>18</sup> Thanks to an anonymous referee for suggesting this line of argument.

<sup>19</sup> Perhaps not all multi-user virtual worlds are meant to sustain play. Certain parts of *Second Life* are perhaps not meant to sustain play in this way. However, many virtual worlds are meant to sustain play and one could see these arguments as applying to those virtual worlds or parts of virtual worlds.

<sup>20</sup> See, for instance, Brandt (1979), Hooker (2000).

this argument, the idea is that virtual actions as a general type contribute to the positive value of play, and in light of this virtual actions are morally acceptable, even if in a particular situation they have negative consequences.

There are two problems with this argument. First, as noted in response to Play Argument 1, it is unlikely that all virtual actions contribute to the value of play. For example, verbally abusing another player does little to contribute to the play experience; were that verbal abuse removed, the play experience would still have value. If that's right, then Play Argument 2 cannot establish the Asymmetry Thesis, since some virtual actions could still be morally wrong. More importantly, however, this argument runs into a classic problem with rule consequentialism. Consider a particular virtual action, *a*, which does contribute to the value of play, but nevertheless has net negative consequences. As noted in response to Play Argument 1, such actions surely exist. If Play Argument 2 is to be successful, we have to say that *a* is morally acceptable, even though it has net negative consequences. But this is unmotivated. For the whole reason to think that virtual actions are morally acceptable is that they contribute to the value of play. So, it is the positive value of play that is making particular virtual actions morally acceptable. However, in this case we have an action, *a*, whose net effects are *negative*. It seems, then, that such an action should not be counted as morally acceptable, certainly not if the acceptability of virtual actions derives from their positive consequences with respect to play. So, Play Argument 2 fails.

This suggests that if arguments centered on play are to be successful in establishing the Asymmetry Thesis, they must appeal to something more than just the positive consequences of play. Consider, then, the following:

Play Argument 3: Intentions and not consequences are what determine the moral status of an action. Since play is a good, if an action within a virtual world is done with the intention of sustaining play, then it is morally acceptable. Thus, if an action is performed within a virtual world with the intention of sustaining play, then even if it would be wrong if performed outside the virtual world, it is still morally acceptable.

This argument will not be able to establish the Asymmetry Thesis in full generality. At best it will establish the thesis for virtual actions performed by players whose intention is to sustain play by performing those actions. It is not clear how many participants in virtual worlds have such an intention. But put this aside; perhaps we can imagine virtual world populated only with participants who have such intentions. Are all actions performed in the virtual world morally acceptable?

An example suggests that the answer is negative.<sup>21</sup> Suppose that a father is pretending to be a monster to engage in play with his young daughter. We can imagine the father taking this pretense too far and portraying himself as a monster in a way that is too realistic, and ends up doing significant harm to his daughter. The father could be aware that the pretense was causing significant harm, and yet engage in that pretense with the intention of sustaining play, by making the play scenario more realistic. In such a case, I think we would judge the father to have done something wrong, despite the fact that the action is done with the intention of sustaining play (recall: he knows that the pretense is harming his daughter). But if that's right, then the intention of sustaining play does not guarantee that an action is morally acceptable and Play Argument 3 fails.

---

<sup>21</sup> I am grateful to an anonymous referee who suggested this very helpful example.

### 3. Castronova and Closed Worlds

I have considered three classes of arguments in favor of the claim that actions done in virtual worlds that would be wrong if real, are not wrong. The first class of arguments focuses on the *virtual* status of such actions; the second class focuses on the fact that the virtual actions are performed within *games*; the third class focuses on the fact that the virtual actions often contribute to *play*. I have argued that these initially plausible arguments fail.

I now want to point out one interesting reason why this might be important. I'll do this by showing how the Asymmetry Thesis is related to what Edward Castronova (2004) has said about virtual worlds.<sup>22</sup> Castronova distinguishes between open and closed virtual worlds. An open virtual world is really just an extension of the real world. It is like a new neighborhood, albeit one that does not take up any *actual* space.<sup>23</sup> In open virtual worlds, all of the norms and rules and responsibilities of the real world carry over:

As open worlds, they will host events that really do matter. By definition, these will be places that are not play. As worlds that matter, these open worlds will deserve exactly the same legal treatment as the real world receives. (Castronova 2004, p. 205)

A closed virtual world, on the other hand, is a world that is explicitly designed to be one where things don't mean anything and where there is no real value:

“Closed” worlds, by contrast, are intended to provide a different array of benefits... These are play spaces; nothing matters there. Assets there have no value. Losses there are unimportant. Crimes committed create no claims of redress. Lost hours are simply lost. No act is actionable. The complete lack of consequence is, in fact, a declaration and imposition of the state. Indeed, it is the closed world's *raison d'etre*. (Castronova 2004, p. 207)

It may appear from these quotations that Castronova is only concerned with whether things legally matter in closed worlds. Indeed that is one of his primary focuses. Despite this, Castronova is clear that he intends his comments to be about *moral* issues, too. For instance, early in the paper (p. 188) he approvingly cites Huizinga's (1955) claim that nothing can be a game if it involves *moral* consequence. And throughout the paper, Castronova is concerned with maintaining virtual worlds as play spaces where nothing really matters in *this* sense, not just in the more restricted legal sense.<sup>24</sup>

Castronova argues that although the current multi-user virtual worlds are not closed, it is desirable to have such closed virtual worlds, because they allow people a space in which to play, free from the constraints of the real world.<sup>25</sup> He suggests that there are steps we can take

---

<sup>22</sup> Castronova prefers the term 'synthetic worlds'.

<sup>23</sup> Or, at least, it is one that does not take up actual space in the normal sort of way. There is some sense in which a virtual world takes up actual space on a hard drive or server.

<sup>24</sup> As an anonymous referee notes, Castronova's views on this may have changed. In a post the *Terra Nova* blog (Castronova, 2005, December 24), Castronova considers the moral status of avatar choice. He writes: “But when the avatar is a considered an expression of self, in a social environment, then deliberately choosing a wicked character is itself a (modestly) wicked act.” This seems to imply that behavior within virtual worlds is morally evaluable and so *matters*, morally speaking. Of course, it is consistent with this claim that this behavior is significant precisely because such worlds are not closed. On the other hand, Castronova might maintain that such worlds are closed and yet have moral (though not legal) importance. Nevertheless, in print Castronova has defended the view described in the text, and the connection between the Asymmetry Thesis and Castronova's (2004) view is interesting independent of whether or not Castronova himself currently holds his previous view.

<sup>25</sup> He writes: “As meaning seeps into these play spaces, their status as play spaces will erode. As their status as play spaces erodes, the laws, expectations, and norms of contemporary Earth society will increasingly dominate the atmosphere. When Earth's culture dominates, the game will be over, the fantasy will be punctured and the illusion will be ended for good. Taxes will be paid. The rich and poor will dance the same macabre dance of

to create these closed worlds: “It is fairly easy to create conditions under which games do or do not matter. It's a choice we make as a society.” (Castronova 2004, p.190) As an example of how this might go, Castronova notes that if we were to allow football players recourse to the legal system to redress a particularly hard hit from an opposing player, we could make football matter *more*. On the other side, by de-localizing teams, so that they are not tied to geographical locations, we could make football matter *less*. With respect to virtual worlds, prohibiting the sale of virtual items for real currency may be one step towards making them matter *less*.

In a moment I'll evaluate whether there could be closed worlds like this, but for now I want to point out that there is a strong connection between the possibility of such closed worlds and the Asymmetry Thesis. First, note that if a virtual world were a closed world, then we would have established the Asymmetry Thesis for that world. If a certain game is one where nothing *means* anything, then any actions performed in the game cannot hold any *moral* meaning. So, if closed virtual worlds are possible, then there are possible worlds for which the Asymmetry Thesis is true. In other words, if a virtual world is closed, then the Asymmetry Thesis is true of it.

Further, if the Asymmetry Thesis were true of a world, then we would have gone a long way to showing that there could be closed virtual worlds. First note that it is plausible that actions that are normally morally acceptable are still acceptable if performed in the virtual world. This is true independent of the Asymmetry Thesis. But if the Asymmetry Thesis is true, then actions normally wrong are also morally acceptable within the virtual world. Thus, if the Asymmetry Thesis were true then we would have a virtual world where nothing one does could be morally wrong. This would be a world devoid of at least one important kind of meaning: moral meaning. This doesn't yet show that closed worlds—worlds devoid of *all* meaning—are possible. But it suggests that they are, and so the truth of the Asymmetry Thesis is evidence in favor of the possibility of closed worlds.

The possibility of closed worlds and the Asymmetry Thesis are closely related. Arguments against the Asymmetry Thesis show that there could be no closed virtual worlds. And arguments against the possibility of closed virtual worlds tell against the Asymmetry Thesis.

#### 4. Against Closed Worlds and Asymmetry

##### 4.1 Against Closed Worlds

I won't challenge the idea that it may be desirable to have closed virtual worlds. However, I think it is questionable whether closed virtual worlds, at least of the sort that people want to spend time in, are possible. Why think this? To put it simply, they are impossible because people really value things in these virtual worlds—they value their virtual goods, their friendships, and their experiences there. In addition, they have significant freedom to act in these worlds. When you have freedom to act in ways that affect things that people value in this way, this seems sufficient for morally significant action. Since moral significance is a kind of meaningfulness, virtual worlds cannot be completely closed: some things really do

---

mutual mistrust that they do on Earth, with no relief, no re-writing of beginnings, and no chance to opt out and start over. The art that once framed an immersive imaginary experience will be retracted back to the walls of the space, and the people will go back to looking at it rather than living it. Living there will no longer be any different from living here, and a great opportunity to play the game of human life under different, fantastical rules will have been lost.” (Castronova 2004, p. 196).

matter. We could keep this from happening only at the cost of restricting the players' freedom or by keeping people from valuing anything in those worlds. But the second option is possible only by making the world so impoverished that no one wants to enter it, and the first makes it a kind of virtual reality, but not an interactive *world*.

What about Castronova's examples of how we can make various games matter *more* or *less*? There are two responses. First, note that the examples Castronova gives—allowing football players to take legal action for in-game actions, and de-localizing teams—would indeed seem to make people care more or less about the games in question. But the important issue is whether or not these things really alter how much in-game actions matter. Now, one might think that whether or not in-game actions *really* matter is purely a consequence of whether or not people think such actions matter (Castronova can be read as suggesting this idea). But if so, then the cost of getting closed virtual worlds where nothing matters is that no one cares about them. But if no one cares about them it's hard to see how they are the kinds of valuable places that people would want to spend time in. So, *if* we can get closed worlds in this way, they won't have the attendant benefits that Castronova wants.

If, on the other hand, in-game actions can matter independently of whether people think they matter, Castronova's examples fail to establish his point. It's not clear from those examples that we can make in-game actions matter or not, all we can do is make people fail to care about such actions. There is a second response to Castronova's examples, however. Suppose that the examples do succeed in showing that we can take certain steps to make in-game actions *really* matter more or less. It doesn't follow from this that we can make them *completely* fail to matter. But it is this strong claim that we need for there to be closed virtual worlds.

So, it seems to me that we haven't been given good reasons to think that the kinds of closed virtual worlds Castronova says are valuable are possible. I've also given a positive reason to think they're *not* possible. If, as I argued above, the possibility of closed worlds and the Asymmetry Thesis stand or fall together, then the impossibility of closed worlds tells against the Asymmetry Thesis.

#### 4.2 Against Asymmetry

My strategy has been to bind the possibility of closed worlds together with the Asymmetry Thesis and then argue against the former. But one can also argue directly against the Asymmetry Thesis.<sup>26</sup>

Marcus Johansson (2009) briefly offers two arguments against the Asymmetry Thesis.<sup>27</sup> The first is the *argument from moral development*. The idea behind this argument is that if I perform some cruel action via my avatar in a virtual world, this strengthens my disposition to be cruel in real life. And, since the latter is wrong, so is the former.<sup>28</sup> This would show the Asymmetry Thesis to be false. For the acts of virtual cruelty on which the argument focuses

---

<sup>26</sup> Of course, this would then cast doubt on the possibility of closed worlds.

<sup>27</sup> He attributes these arguments to Brey (1999), though fails to mention that Brey offers the arguments with respect to virtual worlds that have only one human user rather than interactive worlds with many human users. Thus, Brey's arguments are actually a bit different than those presented by Johansson and attributed to Brey. Johansson, however, offers the arguments with respect to interactive virtual worlds, and so I focus on those in the text.

<sup>28</sup> Notice that this mimics the Kantian argument for why it is wrong to be cruel to animals (Kant 1996, pp. 192-3). This is noted by Brey (1999).

are such that, if they were really performed, they would be wrong. It would, however, show it to be false in an odd way. An act of cruelty in the real world is wrong primarily because of how it affects the person to whom you are cruel. But according to this argument, actions would be wrong in virtual worlds solely because of how they affect the *actor*. So although some actions that would be wrong if performed in the real world will come out as wrong in the virtual world, they will be wrong for very different reasons. Thus, although the Asymmetry Thesis would be false, this argument leaves open that there is an important asymmetry between virtual worlds and the real world.

In any case, I have doubts about this argument. First, it is uncontroversial that an action can have an effect that is bad without itself being wrong. More specifically, it seems as though an action could cause one's dispositions for cruelty to increase without itself being wrong. For instance, going to serve in a just war may increase one's disposition to be cruel. But this, on its own, doesn't show that serving in the war is itself wrong. A second problem with this argument is that it is based on an empirical premise that is not clearly established. It is questionable whether or not performing cruel actions in virtual worlds really *does* strengthen one's dispositions to be cruel in real life. It's not, of course, absurd to claim that this is the case, but it is not absurd to claim that it's not, either.

The second argument offered by Johansson strikes me as more promising. Call this the *argument from psychological harm*. It draws on ideas similar to those presented in Section 2. First, it notes that participants tend to identify with their avatars and value the things in these virtual worlds. So, if those avatars are not treated with respect, or if the valued things in the virtual world are taken, this causes psychological harm to real people. And, if causing psychological harm to real people is wrong, then the virtual actions that do this are themselves wrong. This would show the Asymmetry Thesis to be false. Failing to show respect to an avatar (say, by stealing its virtual property or betraying it) is to do an action that would be wrong if that action were performed in the real world. And since it causes psychological harm, it is wrong.<sup>29</sup> Of course, it is consistent with this argument that some or even most virtual actions cause *positive* psychological consequences and so are morally right.

If the argument from psychological harm is successful, then the Asymmetry Thesis is false for virtual worlds. And, since if there were a closed virtual world, the Asymmetry Thesis would be true of that world, these arguments would show that there could be no closed virtual worlds.

## 5. Limited Asymmetry

In the previous section I argued that there is a close connection between the possibility of closed virtual worlds and the truth of the Asymmetry Thesis for a virtual world. I then cast doubts on whether closed virtual worlds are possible and also considered arguments directly against the Asymmetry Thesis.

---

<sup>29</sup> Notice that according to this argument, these virtual actions are wrong for roughly the same reasons that their real counterparts are wrong. This is an important difference between it and the argument from moral development. Of course, just because an action causes *some* harm, it does not follow that it is wrong. The harm could result in some greater benefit. However, it is hard to see what this greater benefit could be in these cases. The most plausible thought is that although I harm you by doing something to your avatar in a virtual world, the benefit is the pleasure I get from doing this. But we don't think that this sort of benefit offsets harm caused in non-virtual situations, so it is not clear why things would be different here (for more on this idea, however, see section 6).

There is, however, something of a puzzle here. For we do seem able to create something like *partially* closed worlds. After all, there are some things we can do in some games that are not wrong, but would be if it did not occur within the context of the game. I don't think I've done anything wrong, for instance, if I invade your country while playing a game of *Risk*. Similarly, even if it is not the case that there are virtual worlds where *every* virtual action is morally acceptable, there does seem to be an asymmetry between *some* virtual actions and their real-world counterparts. What accounts for this?

One promising way forward here is to try to combine the various features in the failed arguments for the Asymmetry Thesis (Section 2) to see if *together* they might help explain the partial asymmetry we see in virtual actions compared to their real-world counterparts.<sup>30</sup> Recall that the main argument against the Asymmetry Thesis is the argument from psychological harm. Since virtual actions, just like real-world actions, can cause harm, virtual actions that cause significant harm can be wrong. This suggests that if there are virtual actions that are not wrong, even though the real-world counterpart of the action would be wrong, it is because either the virtual action does not cause the same amount of harm, or the harm it causes is mitigated by something else. The virtual, game, and play arguments all have something to say on this score.

Consider first the fact that virtual actions are virtual. The virtual nature of these actions perhaps lessens the harm that can be done by them. The reason for this concerns how closely players can identify with their avatars. The idea is that although players often identify quite closely with their avatars, the identification is *weaker* than it is with their own physical bodies. One reason to think that this is so is that, given current technological ability, we depend for our existence on our bodies, but not on our avatars. Roughly: if our avatars die, we keep existing; if our bodies die, we do not.<sup>31</sup> This is one reason to think that identification with one's avatar, although potentially significant, is at our current state of technology less close than identification with one's physical body. But if that's right, then it is plausible that the harm that you can do to me by doing something to my avatar is *less* than the harm you can do to me by doing something to my body. For example, although virtual rape and real rape both cause harm, the harm caused by the latter is greater.

One has to be careful with this generalization. It is ultimately an empirical matter how closely players do identify with their avatars. Though it is somewhat plausible that this identification is weaker than the identification a player makes with her own body, this need not be the case. Further, future technological advances could surely change the dependence relations between persons and their avatars and thus render the identification between a player and her avatar as strong (or stronger than) the identification between a person and her body. More than that, even if the identification between person and avatar is *now* weaker than identification between person and body, it might not be *much* weaker. Nevertheless, it seems we do currently have an approximate regularity: actions in virtual worlds cause less harm than actions in the real world because players identify less with their avatars than with their real bodies.

Consider next the fact that many of those who play in these virtual worlds plausibly give some minimal kind of consent. Many players have some familiarity with the game they are going to play and by voluntarily entering into the game give some minimal consent to by

---

<sup>30</sup> Thanks to Erik Wielenberg for prompting me to think about this strategy.

<sup>31</sup> I don't mean to commit myself to a controversial thesis about personal identity here. The claim is not that it is metaphysically necessary that I cannot exist without my body. Rather, the claim is that, given current technology, I depend on my body more than I depend on my avatar in a virtual world.

bound by the rules of the game. We saw above that consent does not have unlimited power to transform every action that would be wrong to one that is morally acceptable. Nevertheless, consent can transform some actions. The example of boxing seems to establish this. Given this, harm done in a virtual world, which is often lessened because of the virtual nature of these worlds, is further mitigated by the fact that most players have consented to the possibility of having such harming actions done to them.

Consider finally the fact that play itself may be a good. If there is a real value to play then actions within a game that sustain play themselves have positive consequences. Not every action within a game helps to sustain the play. And even of those that do, some may cause more harm than the benefit of sustaining play (consider the example of the father playing monster from Section 2.3). Nevertheless, the fact that actions within a game that sustain play have this positive consequence of sustaining play mitigates against the harm such actions may cause because the positive consequences must be balanced against the negative consequences.

Putting these three ideas together, we can gain insight into the limited asymmetry there is between virtual actions and real-world actions. Because the actions are virtual, the harms they cause tend (currently) to be less than the harms caused by real actions. Because many players give consent to engage in games where harmful actions may be done to them, some of the harm caused is not morally wrong. Finally, because play itself can be a good, there are positive consequences of actions that sustain play but nevertheless cause harm. This all suggests that the actions that will be wrong in virtual worlds are actions that (i) occur in worlds where players strongly identify with their avatars and so the virtual nature of the world does not lessen the harm, (ii) there is little or no consent given to being harmed, perhaps because players are ignorant about possible harms that could be done to them in the virtual world, and (iii) the actions that harm do not contribute to play, but are gratuitous harms.

## 6. Conclusion

I have argued that the Asymmetry Thesis is false for modern multi-user virtual worlds. That is, there are no virtual worlds where any action is morally acceptable. I have argued that this is not a mere contingent fact, rather, such worlds are impossible. This tells against the possibility of closed worlds, too. If worlds cannot be morally closed, then they cannot be closed *tout court*. This shows us something interesting. Actions in that players take in these games cannot be disregarded from a moral perspective simply because they occur within virtual worlds.

Nevertheless, it is consistent with these arguments that there are actions within virtual worlds which would be wrong if real, but are not morally wrong when performed within the virtual world. It is an interesting question why and how this works. The penultimate section of this paper investigated this interesting question, suggesting that the virtual nature of these worlds, the consent of the players who enter them, and the value of play itself, may be helpful in explaining how and why we get this limited asymmetry when we do.

## Acknowledgements

Earlier versions of this paper were presented at *Ethical Inquiry through Video Game Play and Design* at DePauw University, and at the 6<sup>th</sup> *International Conference on the Philosophy of Computer Games* in Madrid, Spain. Thanks to the participants of both conferences for helpful comments. Thanks also to Erik Wielenberg, Marcia McKelligan, and the members of the DePauw University Ethics Bowl Team for discussion and comments. I also thank two anonymous reviewers for this journal.

## References

- Athanassoulis, N. (2002). The role of consent in sado-masochistic practices. *Res Publica*, 8, 141–155.
- Bachman, R. (1983). *The Running Man*. Signet Books.
- Bradley, C. & Froomkin, A. (2004). Virtual worlds, real rules. *New York Law School Law Review*, 49, 103–146.
- Brandt, R. B. (1979). *A Theory of the Good and the Right*. Oxford University Press.
- Brey, P. (1999). The ethics of representation and action in virtual reality. *Ethics and Information Technology*, 1, 5–14.
- Castronova, E. (2004). The right to play. *New York Law School Law Review*, 49, 185–210.
- Castronova, E. (2005, December 24). The Horde is Evil. *Terra Nova* [Weblog]. Retrieved June 4, 2012, from [http://terranova.blogs.com/terra\\_nova/2005/12/the\\_horde\\_is\\_ev.html](http://terranova.blogs.com/terra_nova/2005/12/the_horde_is_ev.html)
- Community Standards*. (n.d.). *Second Life*. Retrieved June 18, 2012, from <http://secondlife.com/corporate/cs.php>
- Craft, A. (2007). Sin in cyber-eden: Understanding the metaphysics and morals of virtual worlds. *Ethics and Information Technology*, 9, 205–217.
- Driver, J. (2007). Dream immorality. *Philosophy*, 82, 5–22.
- EVE Online [computer software]. (2003). CCP.
- Fairfield, J. (2005). Virtual property. *Boston University Law Review*, 85, 1047–1102.
- Francis, T. (2005). Murder incorporated. *PC Gamer UK*, 126–129.
- Gooskens, G. (2010). The ethical status of virtual actions. *Ethical Perspectives*, 17, 59–78.
- Hooker, B. (2000). *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*, Oxford University Press.
- Huizinga, J. (1955). *Homo Ludens: A Study of the Play-Elements in Culture*. Beacon Press.
- Hunter, D. & Lastowka, G. (2004). Virtual crimes. *New York Law School Law Review*, 49,

211–229.

Hurd, H. (1996). The moral magic of consent. *Legal Theory*, 2, 121–146.

Johansson, M. (2009). Why unreal punishments in response to unreal crimes might actually be a really good thing. *Ethics and Information Technology*, 11, 71–79.

Kant, I. (1996) *The Metaphysics of Morals*, Mary Gregor (trans.) Cambridge: Cambridge University Press.

Kerr, O. (2008). Criminal law in virtual worlds. *University of Chicago Legal Forum*, 415.

Lander, M. (2003, December 26). Kassel Journal; Eating people is wrong! But is it homicide? Court to rule. *New York Times*. Retrieved December 8, 2011, from <http://www.nytimes.com/2003/12/26/world/kassel-journal-eating-people-is-wrong-but-is-it-homicide-court-to-rule.html>.

Munn, N. J. (*forthcoming*). The reality of friendship within immersive virtual worlds. *Ethics and Information Technology*.

Nielsen, M. E. J. (2010). Safe, sane, and consensual—consent and the ethics of BDSM. *International Journal of Applied Philosophy*, 24, 265–288.

Nolan, C. (2010). *Inception*. Warner Bros. Studios.

Orend, B. (2006). *The morality of war*. Broadview Press.

Powers, T. (2003). Real wrongs in virtual communities. *Ethics and Information Technology*, 5, 191–198.

Lamorisse, A. (1957). *Risk* [board game], Parker Brothers.

Rossignol, J. (2005). A deadly dollar. *The Escapist*, 19, 18–22.

Second Life [computer software]. (2003). Linden Lab.

Strikwerda, L. (2012). Theft of virtual items in online multiplayer computer games: an ontological and moral analysis. *Ethics and Information Technology*, 14, 89–97.

Walzer, M. (2006). *Just and unjust wars: A moral argument with historical illustrations*. Basic Books.

Warburton, S. (2009). Second Life in higher education: Assessing the potential for and the barriers to deploying virtual worlds in learning and teaching. *British Journal of Educational Technology*, 40, 414 – 426.

Wertheimer, A. (2000). What is consent? And is it important? *Buffalo Criminal Law Review*, 3, 557–583.

Wolfendale, J. (2007). My avatar, my self: Virtual harm and attachment. *Ethics and Information Technology*, 9, 111–119.

World of Warcraft [computer software]. (2004). Blizzard/Vivendi.