

Virtual Worlds and Moral Evaluation

Jeff Dunn

Introduction

Consider the virtual worlds of Massively Multiplayer Online Games (MMOGs) such as *EVE*, *World of Warcraft* or *Second Life*. In such games, 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. Many interesting questions suggest themselves when considering these MMOGs. 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 MMOGs. Have I then done something wrong?¹

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 an MMOG (via an avatar), which affect other avatars in the MMOG and would be wrong *if* the situation in the MMOG were *real*. The question about such virtual actions is: are they wrong?²

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.³ 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.⁴

The ultimate goal of this paper is to show how the Asymmetry Thesis is closely related to what Edward Castronova (Castronova 2004) has called *closed synthetic worlds*. These are virtual worlds designed solely for play, where nothing that occurs in the worlds matters. Castronova argues that it would be valuable to have such virtual worlds and that we can in

¹ 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).

² (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. For other work on this topic, see (Powers 2003), (Wolfendale 2007), and (Johansson 2009). Some of these will be discussed below.

³ (Powers 2003: 198) suggests that he ascribes to the Asymmetry Thesis. This is despite the fact that he argues that some virtual actions—though not the kinds in MMOGs—can be wrong.

⁴ (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.

fact create such worlds. I'll suggest that either we cannot have such closed worlds, or if we can, that they will be impoverished in a way that limits their value. Here, then is what's to come. I'll first consider some arguments in favor of the Asymmetry Thesis and argue that these fail. After that, I'll explain how the Asymmetry Thesis is related to the possibility of closed worlds, arguing both that if the Thesis is false, then there are no closed worlds and that if closed worlds are impossible, then this is a strong reason to doubt the Asymmetry Thesis. In the third section I'll give positive arguments against the Asymmetry Thesis. Finally, I consider what seems the most promising way of arguing for the Asymmetry Thesis and thus the possibility of closed worlds. However, I point out that if successful, this defense points towards a way in which these worlds will be importantly impoverished.

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.⁵ 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.

Arguments For The Asymmetry Thesis

But It's Not Real!

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—which is an attempt to establish the Asymmetry Thesis—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. 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

⁵ 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.

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.⁶

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.⁷

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 argues (Wolfendale 2007), players of MMOGs identify with their avatars. For players, 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: 114)⁸ Virtual Argument 2, it seems, fails.

Consider one final argument in favor of the Asymmetry Thesis for MMOGs, 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 (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 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)

⁶ (Hunter & Lastowka 2004) note this, too.

⁷ 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.

⁸ 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 players of MMOGs 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 MMOGs, in a way that it cannot in mere chat rooms.

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. The virtual worlds of MMOGs 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 MMOGs 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 MMOGs are both non-veridical, there is a key difference between them. One important feature of the virtual worlds of MMOGs 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 play the game, virtual actions become systematically connected with real consequences. This is an important difference between dreams and virtual worlds.⁹

Although I certainly haven't considered every argument for the Asymmetry Thesis based on the virtual nature of MMOGs, 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.

But It's Just A Game!

The two 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!*

Consider first an argument that focuses on the fact that games are voluntary; one need not play.

Game Argument 1: Theft or betrayal in an MMOG is not morally wrong because the players voluntarily enter into the virtual world of the game.

Notice that this argument does not argue from the fact that players *consent* to play by certain rules, to the conclusion that certain actions fail to be morally wrong. Such an argument will be considered later. Rather, this argument maintains that the fact that players voluntarily enter the virtual worlds, the actions such as theft or betrayal are not morally wrong in such worlds.

I think that this argument fails. To see why, consider an analogy. Suppose that there is a dangerous neighborhood where it is known that muggings have occurred. Suppose I voluntarily walk through the neighborhood and am mugged. It might be true that I was foolish to walk through such a neighborhood. Perhaps I am partly to blame for the misfortune I have suffered. Nevertheless, this does not on its own exonerate the people who mugged me. It was wrong for them to do that. If a person voluntarily enters an area where mugging occurs, this does not thereby make it morally acceptable to mug that person. More generally, if an action

⁹ 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.

of type *A* against person *P* is wrong, it is not thereby morally acceptable in virtue of the fact that *P* voluntarily enters a place where *A* is known to occur.

Let's try a different tack. Consider an argument that defends the acceptability of virtual theft by appealing directly to the fact that it occurred within a game:

Game Argument 2: Theft or betrayal in an MMOG is not morally wrong because actions within games aren't right or wrong.

This argument, too, 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. So, 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.

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 3: If an action within a game is in accordance with the rules of the game, then the action is morally acceptable. Thus, if virtual theft or virtual betrayal is in accordance with the rules of the MMOG, then the virtual theft or virtual betrayal is not morally wrong.¹⁰

This argument suggests something very 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.¹¹ However, unlike the example of morality in war, this argument suggests that by crafting the right set of rules for the virtual world of an MMOG, 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

¹⁰ Both (Kerr 2008) and (Hunter & Lastowka 2004) note that something similar to this argument is upheld by US courts. Kerr (p. 421) sums it up nicely, saying that from a legal perspective, "the rules of the game trump". 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.

¹¹ For some contemporary examples of this, see (Walzer 2006) and (Orend 2006).

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 establish the Asymmetry Thesis, at least for virtual worlds that have the necessary set of rules? Here's a *prima facie* reason to think it does not. Suppose that there is a neighborhood—a real neighborhood—where the rules do not prohibit stealing others' goods. In fact, the rules of the neighborhood state that it is okay to steal others' goods here. So, why would anyone move in? Suppose it's not that any resident particularly wants his or her goods stolen, it's just that each finds that the thrill of living in such a dangerous neighborhood outweighs the potential loss of goods. Suppose you've been living in the neighborhood for a while and come home to find everything in your house has been stolen. Was that action of stealing morally wrong? Opinions can split on this. Suppose the action is more heinous: suppose it is hate speech or rape. Do the rules change the moral status of those actions? My intuition is *no*. The actions are wrong, although perhaps blame would be assigned differently than normal. This suggests that we cannot voluntarily create places where the moral norms hold no sway, places devoid of moral value. And if that's right then Game Argument 3 fails to establish the Asymmetry Thesis.

Castronova and Closed Worlds

I've argued that some plausible arguments for the Asymmetry Thesis 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 (Castronova 2004) has said about virtual worlds.¹² 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.¹³ 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: 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'être*. (Castronova 2004: 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, I think Castronova is clear that he intends his comments to be about *moral* issues, too. For instance, early in the paper (Castronova 2004: 188) he approvingly cites Huizinga's (Huizinga 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.

¹² Castronova prefers the term 'synthetic worlds'.

¹³ 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.

Castronova argues that although the virtual worlds of current MMOGs 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.¹⁴ He suggests that there are steps we can take 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: 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 the virtual world of an MMOG 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 virtual worlds of which the Asymmetry Thesis is true. 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. If the Asymmetry Thesis is true, then actions normally wrong are not wrong within the MMOG. If, in addition, actions that are normally *not* wrong were not wrong within the MMOG, 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. It seems to me, then, that 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.

Against Closed Worlds and Asymmetry

Against Closed Worlds

I won't challenge the idea that it may be desirable to have such 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 like those in MMOGs cannot be completely closed: some things really do matter. We could keep this from happening only at

¹⁴ 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 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: 196)

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.

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.¹⁵

Marcus Johansson (Johansson 2009) briefly offers two arguments against the Asymmetry Thesis.¹⁶ 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.¹⁷ This would show the Asymmetry Thesis to be false. For the acts of virtual cruelty on which the

¹⁵ Of course, this would then cast doubt on the possibility of closed worlds.

¹⁶ 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.

¹⁷ Notice that this mimics the Kantian argument for why it is wrong to be cruel to animals (Kant 1996: 192-3). This is noted by (Brey 1999).

argument focuses 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 some I have already presented. First, it notes that players of MMOGs 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.¹⁸

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.

The Rules of the Game and Consent

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.

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 games that are not wrong, but would be if it did not occur within a game. I don't think I've done anything wrong, for instance, if I invade your country while playing a game of *Risk*. And this seems to raise a potential way to reinstate the Asymmetry Thesis. Maybe, the thought goes, we can do

¹⁸ 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.

whatever we do in *Risk* to make invasion acceptable, but do it wholesale, for all actions. This would seem to assure us that the Asymmetry Thesis is true of such a game and also allow us to build such games into the closed virtual worlds that Castronova thinks it would be desirable to have.

So, why is it not wrong to invade your country when we're playing *Risk*? One might say this is solely due to the fact that invading a country in *Risk* has no negative consequences. But that's not quite right. Of course it doesn't have all the consequences that come with invading an actual country. But if a player (or even a particularly devoted observer) is heavily invested in the game, the invasion can have negative consequences. Here's a thought, then, about why such an invasion is not wrong. First, the invasion is in accordance with all the rules of the game, and second, the players have consented to play by those rules. The first part of this explanation draws on the thought behind Game Argument 3. But the simple fact that an action is in accordance with the rules is not enough. The second part of the explanation attempts to deal with this, appealing to the notion of consent. 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.¹⁹

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 establish the Asymmetry Thesis?

To consider this question, first note that the arguments against the Asymmetry Thesis considered above don't make an allowance for consent. For instance, Wolfendale (Wolfendale 2007) presents a version of the argument from psychological harm. She first considers the argument by Thomas Powers (Powers 2003) to the conclusion that the infamous virtual rape in *LambdaMOO* was morally wrong. Unlike the graphical virtual worlds of modern MMOGs, *LambdaMOO* was a text-based online world. The virtual rape occurred when a user described sexually explicit acts being done to several members of the *LambdaMOO* community. Powers argues that this kind of action was both socially unacceptable to the users of *LambdaMOO* and that the action psychologically harmed those it was directed against. Interestingly, Powers argues that although the virtual rape in *LambdaMOO* was morally wrong, similarly heinous virtual actions in MMOGs are not morally wrong because deviant behavior is often socially acceptable in such games. Wolfendale takes issue with this last claim, noting that many forms of deviant behavior in MMOGs do seem to be socially unacceptable, just as they were in *LambdaMOO*.

She then goes on to explain why virtual actions like virtual rape, murder, betrayal, or verbal abuse are morally wrong:

...it is the relationship between the controller and her character mediated by an environment of shared expectations that makes virtual harm a genuine moral wrong. Virtual harms constitute real moral harms against real people. What is different in virtual harm is not the nature of the harm but how the harm is inflicted, which is through participants' attachment to and identification with their online characters. It is

¹⁹ For example, (Hurd 1996) talks about the "moral magic" of consent.

the nature of the relationship between the avatar and the controller that makes virtual harm possible... (Wolfendale 2007: 114)

What is important to note, however, is that the fact that the actions are socially unacceptable in a world plays no role in her argument. That is, this argument, if successful, works independent of whether certain deviant behaviors are socially acceptable. Some virtual action could cause psychological harm to a player even if that action is in accordance with what is socially acceptable. Likewise, some virtual action could cause psychological harm even if it is in accordance with consented to rules. Thus, if one is convinced by the argument from psychological harm, then one must reject the idea that consented to rules can render the Asymmetry Thesis true of some virtual world.

Of course, an equally consistent response is to see this as pointing out a flaw with the arguments against the Asymmetry Thesis. But even if one thinks those arguments are flawed, and so that *sometimes* consent is relevant to whether or not an action that caused harm was wrong, it doesn't follow that consent to a certain set of rules can always do this. Although many have been willing to grant that consent has some transformative power, most have thought it to be limited in its power.²⁰ 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.²¹ 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 an MMOG may not render every in-game action morally acceptable.

There are also some 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. Presumably, however, this is largely a practical problem. If consent to known rules *can* render no actions wrong within a virtual world, then the problem is merely to get those who join the worlds to know the rules. There is no in-principle barrier to having a closed virtual world if this were the only problem.

Perhaps more troubling is the kind of consent to the rules that plausibly takes place. Consider a particular heinous action that could occur within an MMOG. Perhaps it is some sort of abuse by one avatar against 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 general consent always works to transform the moral status of actions, even if one thinks that specific consent does.²² Consider again the example above of the neighborhood where

²⁰ For instance, (Wertheimer 2000: 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." "

²¹ 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).

²² In a discussion of BDSM, (Athanasoulis 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 way.

one of the rules is that stealing is permitted. Suppose I move in to the neighborhood, consenting to the rules. But suppose a year later I find that I don't want my things stolen. Presumably no one is to *blame* for stealing my things in such a situation—they have every reason to believe I have consented to such behavior. But despite this, we might still think that a (blameless) moral wrong is done. My things are taken without my consent. If that's right, then initial consent to the rules of a game does not guarantee that all in-game actions in accordance with the rules of a game are morally acceptable.

The above considerations suggest that even when players consent to the rules we cannot have a virtual world that satisfies the Asymmetry Thesis. But though they suggest this, they do not completely establish the point. Suppose, then, that the complications above can be navigated. Suppose we can explain how rules can be written for a game that render any in-game action permissible so long as the players consent to those rules. If so, it seems that we could have a virtual world for which the Asymmetry Thesis is true. Though in itself this is not *sufficient* for such a world to be the kind of closed virtual world that Castronova (Castronova 2004) thinks would be valuable, it would go a long ways towards establishing the possibility of such worlds. However, it is important to note that such worlds would exhibit an important limitation. This is because consent sometimes changes not only the moral status of certain actions but also the nature of those actions. That is, it is different to betray someone who has consented to that betrayal than it is to betray someone who has not so consented. It is different to steal from someone who has consented to such stealing than it is to steal from someone who has not so consented. In fact, we may not even want to call the consented-to actions instances of betrayal or stealing at all. But this means that securing the truth of the Asymmetry Thesis for a world involves a trade-off. There are certain things in those worlds you just can't do. And this impoverishes those worlds in certain ways. I won't take a stand on whether this lowers the value of these worlds sufficiently to undermine their desirability (from the perspective of Castronova). But it does point to a trade-off: we can have a very realistic virtual world, or we can have a world that satisfies the Asymmetry Thesis, but we cannot have both. In the virtual as well as the real, it seems, everything comes with a cost.²³

Games

EVE ONLINE, CCP, PC, 2003
RISK, Parker Brothers, 1957
SECOND LIFE, Linden Lab, PC, 2003
WORLD OF WARCRAFT, Blizzard/Vivendi, PC, 2004

References

Athanassoulis, N. (2002). The role of consent in sado-masochistic practices. *Res Publica*, 8, 141–155.
Bradley, C. & Froomkin, A. (2004). Virtual worlds, real rules. *New York Law School Law Review*, 49, 103–146.

²³ An earlier version of this paper was presented at the conference *Ethical Inquiry through Video Game Play and Design* hosted by DePauw University. Thanks to the participants for helpful comments. Thanks also to Erik Wielenberg and two anonymous referees for this conference for their comments.

- 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.
- 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.
- 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.
- 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.
- Rossignol, J. (2005). A deadly dollar. *The Escapist*, 19, 18–22.
- Walzer, M. (2006). *Just and unjust wars: A moral argument with historical illustrations*. Basic Books.
- 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.