

Free will and moral responsibility in video games

Christopher Bartel¹

Published online: 9 January 2016
© Springer Science+Business Media Dordrecht 2016

Abstract Can a player be held morally responsible for the choices that she makes within a videogame? Do the moral choices that the player makes reflect in any way on the player's actual moral sensibilities? Many videogames offer players the options to make numerous choices within the game, including moral choices. But the scope of these choices is quite limited. I attempt to analyze these issues by drawing on philosophical debates about the nature of free will. Many philosophers worry that, if our actions are predetermined, then we cannot be held morally responsible for them. However, Harry Frankfurt's compatibilist account of free will suggests that an agent can be held morally responsible for actions that she wills, even if the agent is not free to act otherwise. Using Frankfurt's analysis, I suggest that videogames represent deterministic worlds in which players lack the ability to freely choose what they do, and yet players can be held morally responsible for some of their actions, specifically those actions that the player wants to do. Finally, I offer some speculative comments on how these considerations might impact our understanding of the player's moral psychology as it relates to the ethics of imagined fictional events.

Keywords Violence · Video games · Free will · Harry Frankfurt · Moral choice · Moral responsibility

I have committed numerous violent acts in video game worlds—brutal, unspeakable acts. Most of the time, I felt no remorse. Most of the time. But in fact, some of the time,

I did feel uncomfortable, sometimes deeply. But why should I? In single-player video games, my actions—however horrible they might seem—are perpetrated against non-player characters, bloodless virtual beings who do not exist outside of the gamespace and who feel no pain. My virtual actions carry no real-world consequences. So, why should I ever feel uncomfortable? In this paper, I want to address the familiar debate over violence in video games by focusing on the issue of the player's occasional feelings of moral discomfort.

The debate over violence in video games is one facet of a much larger debate that carries fascinating philosophical implications about the scope of morality. Specifically I am interested in the relationship between morality and imagination. What are the limits of morality? Does the scope of morality end where imagination begins?¹ Is it itself morally wrong to imagine or fantasize about immoral things?² Some fantasies really do seem to be harmless—for instance, idle fantasies, which are those thoughts and images that pass through our minds seemingly without our control and sometimes against our will. It would be entirely unfair to think that it is morally wrong to entertain *passive*

✉ Christopher Bartel
bartelcj@appstate.edu

¹ Appalachian State University, 401 Academy Street,
I. G. Greer Hall 114, Boone, NC 28608, USA

¹ There has been some recent discussion of the problem of the morality of fantasies and works of fiction explicitly. Contemporary interest in the debate began with Kendall Walton's (1990) theory of fiction and imagination. Discussion of Walton's view and the possible limitations of imagination regarding morality have led to debate over the possibility of "imaginative resistance". For discussion, see Gendler (2000) and Walton (1994). Some theorists have questioned whether the contents of works of fiction can be inherently immoral. For general discussion, see Cooke (2012, 2014), Gaut (1998), Harold (2002), and Smuts (2013). Finally, the recent debate over the "gamer's dilemma" offers a more specific case of this general problem, and one that relates directly to video games. For discussion, see Luck (2009), Patridge (2011, 2013), and Young (2013).

² I take this way of formulating the question from Cooke (2014).

fantasies. But, what about those fantasies and daydreams that we willingly return to time and again; the ones that make up a significant part of our unspoken mental lives; the ones that we *really* enjoy? Intuitively there seems to be a morally relevant difference between fleetingly imagining some immoral act and repeatedly fantasizing about that immoral act. For instance, it may be one thing to imagine enjoying a casual sexual encounter with a colleague, but another thing to regularly pore over such a fantasy with relish. On the other hand, another common intuition that I suspect many people hold is that there is nothing morally wrong with fantasizing about doing something that would be immoral, just so long as one never actually acts on those imagined fantasies. These two intuitions seem to be in conflict and it is not obviously clear which one should win out. To adjudicate between these intuitions, we need a better understanding of the role of morality in fiction and imagination.

Debates about violence in video games tend to focus on possible correlations between virtual violence and real-world crime. There has been much discussion of what real-world harm (if any) could be associated with virtual violence,³ and many researchers have sought evidence of some correlation between virtual violence and real-world negative behaviors.⁴ This body of work is certainly important, but researchers working on these topics often adopt a narrow understanding of the moral relevance of virtual violence. Specifically we can observe this narrowness in the way that many theorists understand the concept of *harm*: many seem to assume that “harm” equates to observable, quantifiable, real-world crimes: things that either draw blood, result in bruises, or result in a loss of property. If this is what we take “harm” to mean, and if the enjoyment of violence in video games does not cause any noticeable increase in real-world crimes—that is, if no real-world blood has been drawn—then it causes no harm. This is a narrow view of harm because there are obviously harms that are not equated with quantifiable crimes. For instance, it may not be a crime to be unfaithful to one’s lover, but infidelity is still a kind of harm. Could a broader conception of harm offer us a way to think about immoral fantasies? Some fantasies might not be innocent. Consider this case: Joe and Sally are not married, but they have agreed to a monogamous relationship. Joe never acts unfaithfully, but he regularly and willfully fantasizes about cheating on Sally. Has Sally been harmed? She certainly might feel like she has been.

³ See for instance, Dunn (2012), Gooskens (2010), McCormick (2001), Patridge (2011), and Tavinor (2009).

⁴ The body of literature on violent video games is enormous. The reader would do well to begin with a review of the literature, such as Barlett et al. (2009) or Ferguson (2007).

Now consider the case of violent video games: is it ever morally wrong to virtually enact a violent fantasy in a video game? Certainly video game violence causes no real-world blood to be drawn; but if we move beyond a narrow conception of harm, then we might wonder whether our enjoyment of violent video games can ever amount to the kind of non-innocent fantasy that would be analogous to Joe’s fantasies of infidelity.

These are big questions that I cannot hope to fully answer here. Instead, I want to suggest that we can make progress toward answering these questions by considering those cases where players become distinctly uncomfortable with the violence in some video game. A notorious case to consider is the “airport massacre” mission in *Call of Duty: Modern Warfare 2* (2009)—the “No Russian” mission. In this scene, the player assumes the identity of an undercover American CIA agent who is attempting to infiltrate a group of Russian terrorists. During the mission, the group enters a crowded airport and massacres scores of unarmed civilians. The player is given the option to participate in the massacre or to refrain with no penalty to their progress or achievements in the game. The player is also given the option to skip the mission at any point. However, the player is not given the option to turn her weapon on the terrorists and attempt to save the civilians. The inclusion of this mission in the game was highly controversial and makes many players uncomfortable. Some players simply cannot bring themselves to open fire on virtual civilians. But why not? If the game merely represents the virtual depiction of violence and it does not contribute to any real-world harm, then why should we feel any discomfort about participating in imaginary violence?

In this paper, I will argue that some of these cases of moral discomfort can be explained by a conflict concerning the player’s sense of free will. Specifically, I will argue that games that offer the player a moral choice can cause a sense of moral discomfort when the player finds none of the available choices to be morally acceptable as in the case described above. In these cases, the player may feel coerced into making a moral choice that she does not want to make and, at that moment, come to realize the limitations of her in-game free will. If this analysis is correct, then this may also give us an interesting way to think about those moral choices that we make in video games that we feel perfectly comfortable with—in those cases, our choices really do reflect our free will. In what follows, I will first describe the problem of free will and moral responsibility, and Harry Frankfurt’s (1971) account of compatibilism. I will then examine how Frankfurt’s compatibilist account can be fruitfully applied to moral choices in video games. Finally, I will end with a discussion of the possible ramifications of this account and briefly indicate a possible direction for future research.

Before moving on, one caveat is required. The argument of this paper is part of a larger project, which is to argue for the legitimacy of some moral criticism of video game violence by developing a variation of Aristotle's *virtue ethics*.⁵ On my account, the legitimacy of the moral criticism of video games must overcome two obstacles: (a) it must identify which virtual actions a player can be held accountable for, and (b) it must explain what could be morally wrong about virtual actions. These are two questions that tend to be taken together, but I believe that it can be fruitful to treat them separately, as I aim to demonstrate here. I hope to show that Frankfurt's account of moral responsibility allows us to draw important distinctions about player culpability. Of course, to say that a player is morally responsible for some action does not thereby explain why that action is morally wrong—that is, solving (a) does not solve (b). In this essay, I will only address (a). I will have nothing to say about (b) here. Finally, while I would wish to answer (b) by defending virtue ethics, nothing that I say here should depend on the plausibility of virtue ethics. Indeed, it is my hope that my account of (a) could accommodate other ethical frameworks.

Free will and moral responsibility

The Elder Scrolls V: Skyrim (2011) is an enormous game. With numerous cities, landscapes, and deep, dark places to explore and side-quests to complete, the game is simply overwhelming. Players of *Skyrim* can easily spend hundreds of hours playing the game without making any progress in the main storyline. Additionally, the many ways in which the player can develop her player-character are mind-boggling. At the beginning of the game, the player is asked to select the race, gender, and physical appearance of the main character. Most fantasy RPG's also require the player to select the character's class at this point, which then pigeonholes the player into being a fighter, thief, mage, or whatever from the start. However, *Skyrim* departs from this convention. Instead of selecting one's class at the beginning, the player is free to develop her fighting style as she plays, which offers even more freedom to choose exactly how to develop the player-character in response to the circumstances that one finds in the game.

It seems like open world games offer the player a considerable degree of choice, but it is also obvious that the player's choices are not unlimited. For instance, *Skyrim* allows the player to wander the countryside at will, murder innocent civilians, and steal anything that is not nailed down; but the player is not given the option to become a pacifist, or to pursue a scientific study of the biology of

frostbite spiders, or to give up the adventuring life and open a bed and breakfast. Obviously the reason for such limitations simply has to do with the technological limitations on how much can be programmed into the game. The degree of choice that the player is afforded in such games may be immense, but it has its limits.

This raises important questions. It is a widely held intuition that a person can only be held morally responsible for things that they willingly choose to do. With so many choices at my fingertips in *Skyrim*, including moral choices, am I somehow morally responsible for my in-game virtual actions? Or are the limitations on my choices significant enough to distance me from moral responsibility? Of course, we should recognize that games like *Skyrim* are exceptional cases that offer a wide range of choices. Many games offer very little significant choice, or none at all. For instance, in *BioShock Infinite* (2013) the protagonist is drawn into a battle between the fascist Founders and the rebel Vox Populi; however by traveling through alternative realities, the player is required to fight on both sides of this conflict at different times in the game. The player is never given the choice to align herself with one side or the other permanently. When I have no choices, or if none of the choices that I am given are what I really want to do, then how can I be held morally responsible for my virtual actions? The answer to this question depends partly on whether players have genuine free will within video game worlds. So, let us start with a brief discussion of why philosophers are interested in free will and why it would be relevant to our understanding of morality.

Does anyone ever act freely? Or are our actions somehow predetermined? And can I really be held morally responsible for the things that I do if I am not in control of my own actions? These are questions that have concerned philosophers for millennia because they cut to the heart of some of the deepest concerns about our existence. While this is not the right place to examine these issues in detail, a broad outline will be helpful.

Imagine that the past is an unbreakable chain of events that extends back into time from the present moment all the way to the very origins of the universe. There are not multiple chains of the past, there is only one. The past is rigid and no amount of willing that things were otherwise can change the past. We could say that the past is "determined". What about the future? Is it also one unbreakable chain of events stretching forever onward ahead of us—that is, is the future also determined? Or perhaps the future is made up of an infinite number of chains that lead in an infinite number of different possible directions? *Determinism* is broadly defined as the inability to choose to do otherwise. If the future is structured like the past as one unbreakable chain of events, then no one can genuinely choose to do something other than they are predetermined

⁵ See McCormick (2001) and Sicart (2009) for similar accounts.

to do. Alternatively, a person only has free will if she could have chosen to do otherwise—that is, if there genuinely were alternative choices available to the individual and it is within our ability to freely, willfully select some of those possibilities. Of course, there are some things that are genuinely impossible for me to choose—for instance, I cannot choose to turn into a dragon—but the fact that my choices are limited in this respect does not constitute a lack of free will. Setting aside my inability to choose the impossible, I have free will if I have the ability to choose among possible options that are genuinely available to me.⁶

It is not my intention here to try to defend one particular view of the debate over free will and determinism. Instead, the above description is merely intended to identify what the problem is. A more interesting concern, for our purposes, is to think about the moral consequences of determinism. I mentioned previously that there is a widely held intuition that, when it comes to moral responsibilities, we cannot hold an individual responsible for an event that she had no control over. My moral responsibilities only extend as far as acts that I willingly commit or duties that I willingly neglect. For instance, if lightning hits the tree in my yard and the tree then topples over and crushes your car, I cannot be held morally responsible for the damage to your car. The damage to your car had nothing to do with any choice of mine—I neither willfully damaged your car nor did I neglect some duty to protect your car—so I cannot be held responsible.

If determinism is true, then it would seem to follow as a natural consequence that we cannot be morally responsible for any of our actions. The truth of this consequence would be a terrible blow to our understanding of morality because one of the fundamental concepts—the concept of *moral responsibility*—would need to be abandoned. All that we would be left with is cruel fate: if I am destined to steal your car tomorrow and it is beyond my control to choose otherwise, then how can you blame me? I am simply caught in the tide of fate just like you.

Yet, some philosophers argue that, even if determinism is true, we can still be held morally responsible for at least some of our actions. This is a debate between so-called “compatibilists” and “incompatibilists”. *Incompatibilism* is the straightforward belief that, if determinism is true, then moral responsibility really is out the window in just the way that I previously described. Alternatively, *compatibilism* is the belief that a robust sense of moral responsibility is compatible with determinism.

⁶ In the debate over free will, philosophers—from the most ancient to the present day—have sought to analyze precisely what is the source of the problem and how it should be resolved. For an introduction to the debate over free will and some indication of the range of views available, see Pereboom’s (2009) collection of some of the classic essays.

Is it more intuitive to hold to compatibilism or incompatibilism? This is an open question.⁷ Regardless, to accept compatibilism, one must give up on the intuition that individuals can only be held responsible for the actions that they willingly undertake. For compatibilism to make sense, we need a good explanation of how it is that an individual can be held responsible for an action that they did not and could not choose. There are many different versions of compatibilism that have been the subject of much debate; but I want to focus on Frankfurt’s particularly interesting version.

Frankfurt’s version of compatibilism is officially neutral about whether or not determinism is true (1971, p. 20); but for the sake of argument, let us suppose that some version of determinism is true. Frankfurt distinguishes between the *freedom to act* and the *freedom to will*: even if our *actions* are predetermined, our *will* is not (Ibid, pp. 14–15). A person can willfully choose to *want* something even if that person cannot willfully choose to act on that wanting. On this account, our free will does not allow us to choose how we *act*; instead it allows us to choose whether our actions are what we *want*. The actions that I want to commit are the ones that I identify myself with even if it is determined that I must carry them out. As illustration, Frankfurt offers the example of the unwilling drug addict, which is a person whose physical desire for a drug drives her to act in certain ways, but whose *will* desires to act in other ways. This is a person who does not identify herself with her actions. Her will is to avoid taking drugs, but she acts in accordance with her physical addiction. Whether or not you find this example to be convincing, I believe that Frankfurt does point to a kind of experience that most people would be familiar with: the experience of feeling detached from ourselves, of feeling like we are out of control of our actions, of being consciously aware that we do not want to be a part of what we are doing and yet feeling like we are unable to stop. Sometimes we might find ourselves doing things that we think are uncharacteristic of ourselves—like, joining a group of friends in gossiping about a close friend—or we might find ourselves doing things that we wish we would not do—like, shouting at a loved one who we do not genuinely wish to hurt. Think of those moments when a person gets swept up by the crowd and participates in some event even while thinking to herself, “This isn’t me.” According to Frankfurt, these are cases—regardless of how rare they might be—where our actions and our will come apart.

Frankfurt’s account might seem like a shallow form of free will—and compared to the freedom to act, it is shallow—but Frankfurt’s point is that the freedom to choose

⁷ For discussion of whether compatibilism or incompatibilism are intuitive, see Nichols and Knobe (2007) and Nahmias et al (2006).

what we want is sufficient to secure our moral responsibility even if we do not have the freedom to choose how we act. Even if all of my actions are somehow predetermined, I can exercise my will to choose whether or not I identify myself with those actions; and that is all that is needed for genuine moral responsibility. Indeed, when we act with a free will in Frankfurt's sense, we are doing exactly what we would do even if our actions were not predetermined (Ibid, pp. 18–20). There would be no difference between the freely willed actions of an agent in a deterministic universe and the actions of a genuinely free agent in an indeterministic universe.

Finally, an important point to consider in Frankfurt's account is that our freedom to will is intimately connected to the concept of a person and ultimately also to our development of a sense of self. Part of what makes us human—part of what is “essential to being a person” (Ibid, p. 10)—is our capacity to will. Our sense of self is partly constituted by our will. Most of the time, I identify myself with my actions; but sometimes I do not. When my actions correspond with my will, we can say that my actions are fully my responsibility, and that I can be held morally responsible for those actions.⁸

This is an intriguing proposal as an analysis of the problem of free will, even if it is not a widely accepted one. But that should not concern us. Whether or not Frankfurt's theory offers us a helpful way of thinking about moral responsibility in reality, his theory seems to fit in the case of video games superbly. It is to that issue that I now turn.

Frankfurt's compatibilism in the virtual world

Many of the violent acts that players commit in video games—perhaps even most of them—are not freely chosen by the player. The *Grand Theft Auto* games offer some excellent examples here. I will focus on two. In *Grand Theft Auto IV* (2008), the player controls the character of Niko Bellic, who undertakes the non-optional mission to help the Irish mob rob a bank—the “Three Leaf Clover” mission. Predictably, the robbery goes badly. When the police show up, Niko must shoot his way out in order to make his escape. However it is played, the mission is chaotic and scores of police officers are killed in the gunfight.

When I first played through this mission, I was horrified. I felt awful about shooting police officers, even if I was merely pretending to shoot at virtual representations of

police officers. The first time I played through this mission, I only lasted a few minutes before my player-character was killed. The reason why I played so poorly is because I refused to shoot at the police! I tried to sneak out of the violence without firing a shot. Unfortunately, I soon realized that this strategy was not an option. To complete the mission, I had to lead my gang members safely out of the violence, and they were not willing to go without a fight. If I wanted to complete the game, then I had to resign myself to the fact that I had to shoot my way out. It took me a few attempts to finally complete this mission—the gang members that I had to protect kept stupidly running into harm's way—but finally, I did it.

I felt deeply uncomfortable with that mission; but am I morally responsible for the virtual murders that Niko committed? According to Frankfurt, the answer to this question comes in two parts. First, it has to do with whether or not I *acted* freely when carrying out those murders. So, could I have chosen to do otherwise? For obvious reasons, I think not. First, the mission is a non-optional one. The player must complete the “Three Leaf Clover” mission in order to complete the game. Once the player has committed herself to completing the game, then she must also tacitly commit herself to completing this mission. Second, it is not possible to control the behavior of the NPC gang members who all too willingly want a fight with the police. As part of the mission is to lead the gang members out of danger, the player is required to eliminate the danger—so, the player is not given the choice to do otherwise.

In the case of the “Three Leaf Clover” mission, the player genuinely does not have any other alternative than to kill the police officers. So, if determinism is broadly defined as the inability to choose to do otherwise, then this mission was determined to happen. Of course, one could simply choose *not* to play *GTA 4*. But that solution does not settle the moral question. My purpose in this essay is to examine the morality of *virtual actions*. If you choose not to play the game, then there are no virtual actions to talk about. The real question is whether the player is morally responsible somehow for the actions that she performs *in the game*. So, I ask the reader to set aside the option to refuse to play. Moreover, I suspect that most players resort to the refusal to play a game only when they feel pushed to the limit—however far that might be. Most gamers will suffer through a challenging mission for any number of reasons. Some players may push through an uncomfortable mission because the player's overall enjoyment of the game is still quite high, or because the player wishes to see how the game will resolve this difficult mission, or because the player is otherwise invested in the story. While some players will refuse to continue playing a game because of a challenging mission, many do not refuse; and it is the actions of these gamers that I am interested in explaining.

⁸ Certainly there is more that can be said about Frankfurt's account of moral responsibility—for instance, Frankfurt cautions that we need to distinguish between being “fully” responsible and being “solely” responsible (p. 20, fn. 10)—but this level of detail is not necessary for my purposes here.

For this reason, I will set aside the option to refuse to play the game.

Given that my actions in the game were non-voluntary and that I could not have chosen to do otherwise, there is strong reason to believe that I did not *act* freely. But, according to Frankfurt's theory, it is still an open question whether I am morally responsible for my actions or not. The second important consideration is, do I identify myself with those actions? Did I carry out those actions because they were what I *wanted* to do? In my case, the answer was strongly, No. I really did not want to carry out that mission. I really would have preferred to select some other non-violent resolution. I was an "unwilling player". When I played through that mission, I told myself, "This isn't me shooting these officers, it is Niko Bellic."

Of course, that rationale is not true strictly speaking. It certainly was me who was pushing the buttons on the controller directing Niko through the killing spree. If I had put down the controller, Niko would have stopped shooting. So, my *actions* were certainly implicated in the event. But, importantly for Frankfurt, my *will* was not. I felt truly detached and distant from what was happening. From that detached point of view, my experience of the game had changed. I was not playing the game *as myself*, instead I was playing the game *as Niko*. I was able to throw myself into the violence and carry off the mission successfully only because I was directing Niko to behave in a way that I thought was authentic for that character—but those actions were not authentic for me. Further, I imagined that Niko might have felt the same way about the mission that I had. On my interpretation of the game, I felt detached from the violence because Niko felt detached too. Niko is being unwillingly sucked into a world of crime that he does not want and he is being forced to protect his gang members with a sense of guilt and regret that is similar to my own. On Frankfurt's account, I cannot be held morally responsible for those virtual murders because those actions did not reflect my will—and possibly, those actions do not reflect Niko's will either.

By contrast, imagine another player—imagine that it is Joe again—who plays through the "Three Leaf Clover" mission and who fully identifies with the actions of Niko. Joe directs Niko to shoot the police because this *really is* what Joe wants to do—Joe is a "willing player". It is his will and his desire that Niko should shoot scores of police. In this case, there is no distinction between Joe's actions and his will: they are one and the same thing. Insofar as Joe is doing what he wants to do in the game, then Joe can be held morally responsible for his actions in the game according to Frankfurt.⁹

⁹ Cf. Frankfurt's comment on the moral responsibility of the "willing addict" (pp. 19–20).

It is important to notice the similarities and differences between Joe and I. For both of us, our *actions* were not freely chosen: the "Three Leaf Clover" mission is non-optional and the nature of the mission requires that scores of police officers must be shot in order to complete it. Regarding our actions, Joe and I behave in the same way. In fact, we can go further with this idea: imagine that Joe and I employ the same strategy in the game with the same success rate. In that case, both of us play the game with the same level of violence and intensity and we achieve the same results. From the point of view of our actions, we are identical. But importantly, despite our similarities, there is still a morally relevant difference between Joe and I, which is the matter of our wills: Joe is doing what he wants to do in the game, but I am not.

Another example, this time from *Grand Theft Auto V* (2014), will help to illustrate the stark contrast between a willing player and an unwilling player. In this installment of the *GTA* series, the player is able to switch between three main playable characters: Michael, Franklin, and Trevor. All three characters commit numerous crimes; however they do so for different reasons within the narrative of the game. Michael and Franklin are motivated partly by a sense of hubris and partly by a desire to build criminal empires for themselves. Trevor, on the other hand, is motivated by darkly sadistic forces. He really enjoys violence for its own sake and shows little remorse at having to commit some of the worst crimes. In one notorious mission—"By the Book"—Trevor is required to torture a bound captive in order to gain information from him. The methods of torture that the player is asked to choose from—waterboarding, electric shock, tooth extraction—are brutal. The mission is non-optional and the scene continues until the victim eventually cries.

Gamers often defend violence in video games as merely harmless fun that carries no meaning beyond the fictional world of the game. Yet, this mission is deeply uncomfortable for many gamers to play—even some of the most hardened. Frankfurt's account of free will offers a way to understand that discomfort: in these situations, the player feels a conflict between her freedom to act and her freedom to will. While the player has little freedom to act, the player still has the free will to either identify herself with the actions that are committed within the game or not.

Now consider the difference between a player who unwillingly forces herself to complete this mission even though it requires her to do something that she does not want to do and a player who willingly, gleefully plays through this mission because it *is* what he wants to do. The unwilling player feels an uncomfortable sense of conflict—she wants to complete the game, but she does not want to *do this*. Like watching a movie in which the viewer is made to witness an event that she does not want to witness, the

player is carried along by the tide of the deterministic game, unable to genuinely choose to act otherwise. But her freedom to will provides her with a sense of detachment from the actions that her player-character is required to commit. Trevor is the monster, not me. The unwilling player does nothing more than witness his monstrosity. By contrast, the willing player does not merely witness Trevor's monstrous acts; he also cheers them on. The willing player wants Trevor to act monstrously. He wants the scene to go exactly as it does. The distance between his will and Trevor's actions breaks down—they are one and the same.

Virtual acts and moral psychology

With this account of free will, Frankfurt offers us a way to maintain a robust sense of moral responsibility: an agent can be held morally responsible for her actions only to the extent that she identifies her sense of self with the perpetration of those actions. In the case of video games, it seems obvious that many of our actions are not freely undertaken; and yet we still have the freedom to identify ourselves with those actions in Frankfurt's sense. But does that mean that we can be held *morally responsible* for our in-game virtual actions? Not quite yet. Before we can say that a player can be held morally responsible for her virtual actions, we would need to say an awful lot more about the morality of virtual actions; and I do not have the space to take up that discussion here. But having come this far, we have one important conclusion to note: to determine whether a player's actions should be open to moral scrutiny or not, we must look at more than the game's content. The player-actions that are relevant to moral consideration are those actions that are freely willed in Frankfurt's sense. It is the player's choices that ought to be the object of moral concern, and not (only) the game's contents itself.

This conclusion is interesting enough, but I think we can go one step further. If the analysis that I have offered here is correct, then this would provide some evidence that players can and often do make moral decisions within game worlds by employing their actual-world sense of morality. Indeed, the virtual actions that a player identifies with her sense of self can relevantly enter into a consideration of that player's actual moral psychology. This requires some unpacking.

An individual's moral psychology is made up of all of the cognitive apparatus—the concepts, decision-making strategies, heuristics, and affects—that are employed in her moral decision-making. If an individual knowingly and consistently makes certain decisions that cause a considerable amount of needless suffering among those who are effected by her decisions, then this tendency is likely to be reflected somewhere in her moral psychology. Maybe this

decision is due to the way that she conceptually misunderstands the relationship between her decisions and the suffering of others, or due to the way that she conceptualizes the value of other peoples' suffering, or due to some faulty inference that she has a tendency to draw. Whatever the case may be, an individual's moral psychology is the complex web of cognitive factors that play a role in her moral actions and her ability to make morally relevant choices.

My suggestion here—which is in response to the debate between Cooke (2012, 2014) and Gaut (1998)—is that we should consider the possibility that the things that we fantasize about make up part of our moral psychology. To begin, consider the important role that imagination plays in our moral psychology. Imagination is a powerful human attribute that allows us, for example, to consider counterfactual possibilities, plan for future contingencies, and to consider how we might feel about certain scenarios and situations if they were to become actual. Before making a moral decision, we often imagine how certain scenarios might turn out in order to decide what we can morally live with. This practice suggests that, when we imaginatively run through possible scenarios, we employ our *actual* moral values, concepts, and sensibilities. It is not as if we possess distinct moral concepts and values that we employ in imagination, which are separate from those that we employ in our real-lives. If that were the case, then imagination would be a useless tool in moral decision-making. Additionally, our affective and aesthetic responses to works of narrative fiction often depend (in part) on our ability to recognize the moral significance of the events and scenarios that make up the fictional work.¹⁰ For instance, we feel outraged by John Marston's unfair treatment at the hands of the government agents in *Red Dead Redemption* (2010) because we are employing our actual moral conception of fairness. Marston's mistreatment might be fictional, but our moral response to his treatment is actual. Thus, our moral psychology is also employed in our engagement with works of fiction.¹¹

Remember that an important part of Frankfurt's conception of free will concerns the way in which individuals maintain a sense of self, the actions that we identify with becoming part of that conception of self. On my interpretation, Frankfurt's conception of that sense of self is partly made up by our moral psychology. So, when we identify with some action, whether it is virtual or not, and thus when that action becomes part of our sense of self, it is our own actual moral psychology that is being employed. What I am denying is the idea that we develop multiple moral

¹⁰ For discussion of this point, see Carroll (1998) and Walton (1994).

¹¹ I take this to be a point that both Walton (1994) and Gendler (2000) would accept even despite their disagreements.

psychologies—one that we employ in reality and others that we only employ in imagination. For illustration consider again the difference between the willing and the unwilling player. The unwilling player goes through the motions because she has no other genuine options available to her. In playing the “By the Book” mission, the unwilling player may think of her actions within the game as being authentic for the character of Trevor; but not authentic for her. In that case, the unwilling player might develop a *fictional* moral psychology that she applies to Trevor, and she acts within the game in a way that is consistent with his fictional moral psychology.¹² It is in this sense that a player can play a game as a villain, just as an actor can play the part of a villain without thereby coming to hold the same moral viewpoint as the villain.¹³ However, what is missing in these cases is the player’s (or actor’s) *endorsement* of the immoral actions of the villainous character. While we may construct a fictional moral psychology to account for the actions of a villainous character, we do not endorse that moral psychology; and therefore it does not enter into our own moral psychology. We imaginatively maintain a distance between our sense of self and that of the fictional villain. I can play *GTA 5* as Trevor, and I can make moral decisions that are authentic for his character; but I cannot bring myself to endorse his actions by willingly identify my sense of self with them. By contrast, the willing player does not need to develop a fictional-Trevor-psychology because Trevor’s actions are the player’s actions. If our sense of self is partly made up by our moral psychology, which in turn is partly made up by the actions (fictional or non-fictional) that we identify with, then the willing player is making moral choices within the game based on his own actual moral psychology.

In closing, I want to briefly indicate a possible avenue to pursue in future research. To understand the morality of virtual actions, we should pay more attention to those aspects of our moral psychology that motivate our fantasies and imaginings. I began this essay by asking, is it ever morally wrong to fantasize about certain things? Thinking of Frankfurt’s compatibilist resolution of the problem of free will, those fantasies that we willingly return to time and again would appear to be ones that we identify with our sense of self. When we think about the morality of violence in video games, we should not solely fixate on whether there is something intrinsically wrong with violent content in games or whether the enjoyment of violent games correlates to real-world harms. Those are

certainly important questions, and many researchers have sought to explore those ideas. But if we want a full picture of the morality of violence in video games and a better understanding of the role and limitations of the imagination in our moral psychology, then we should consider why some acts of virtual violence are important to some players, what motivates players to seek out virtual violence, and what it means to identify with the virtual representation of violent acts.

References

- Barlett, C., Anderson, C., & Swing, E. (2009). Video game effects—confirmed, suspected, and speculative. *Simulation and Gaming, 40*, 377–403.
- Carroll, N. (1998). Art, narrative, and moral understanding. In J. Levinson (Ed.), *Aesthetics and ethics: Essays at the intersection* (pp. 126–160). New York: Cambridge University Press.
- Cooke, B. (2012). On the ethical distinction between art and pornography. In H. Maes & J. Levinson (Eds.), *Art and pornography* (pp. 229–253). New York: Oxford University Press.
- Cooke, B. (2014). Ethics and fictive imagining. *Journal of Aesthetics and Art Criticism, 72*, 317–327.
- Dunn, J. (2012). Virtual worlds and moral evaluation. *Ethics and Information Technology, 14*, 255–265.
- Ferguson, C. (2007). The good, the bad and the ugly: A meta-analytic review of positive and negative effects of violent video games. *Psychiatric Quarterly, 78*, 309–316.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy, 68*(1), 5–20.
- Gaut, B. (1998). The ethical criticism of art. In J. Levinson (Ed.), *Aesthetics and ethics: Essays at the intersection* (pp. 182–203). New York: Cambridge University Press.
- Gendler, T. (2000). The puzzle of imaginative resistance. *Journal of Philosophy, 97*, 55–81.
- Gooskens, G. (2010). The ethical status of virtual actions. *Ethical Perspectives, 17*, 59–78.
- Harold, J. (2002). Flexing the imagination. *Journal of Aesthetics and Art Criticism, 61*, 247–257.
- Luck, M. (2009). The Gamer’s dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology, 11*, 31–36.
- McCormick, M. (2001). Is it wrong to play violent video games? *Ethics and Information Technology, 3*, 277–287.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2006). Is incompatibilism intuitive? *Philosophy and Phenomenological Research, 73*, 28–53.
- Nichols, S., & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Nous, 41*, 663–685.
- Patridge, S. (2011). The incorrigible social meaning of video game imagery. *Ethics and Information Technology, 13*, 303–312.
- Patridge, S. (2013). Pornography, ethics, and video games. *Ethics and Information Technology, 15*, 25–34.
- Pereboom, D. (Ed.). (2009). *Free will*. Indianapolis, IN: Hackett.
- Sicart, M. (2009). *The ethics of computer games*. Cambridge, MA: MIT Press.
- Smuts, A. (2013). The ethics of singing along: The case of ‘mind of a lunatic’. *Journal of Aesthetics and Art Criticism, 71*, 121–129.

¹² Cf. Sicart (2009).

¹³ Thanks to an anonymous reviewer for this journal for the example and for pushing me to clarify this point.

- Tavinor, G. (2009). *The art of videogames*. Malden, MA: Wiley.
- Walton, K. (1990). *Mimesis and make-believe*. Cambridge, MA: Harvard University Press.
- Walton, K. (1994). Morals in fiction and fictional morality/I. *Proceedings of the Aristotelian Society*, 68, 27–50.
- Young, G. (2013). Enacting taboos as a means to an end; but what end? On the morality of motivations for child murder and paedophilia within gamespace. *Ethics and Information Technology*, 15, 13–23.