



Real wrongs in virtual communities *

Thomas M. Powers

Division of Technology, Culture, and Communication, School of Engineering and Applied Science, University of Virginia, P.O. Box 400744, 351 McCormick Road, Thornton Hall, Charlottesville, VA, USA
E-mail: tmp4p@virginia.edu

Abstract. Beginning with the well-known cyber-rape in LambdaMOO, I argue that it is possible to have real moral wrongs in virtual communities. I then generalize the account to show how it applies to interactions in gaming and discussion communities. My account is supported by a view of moral realism that acknowledges entities like intentions and causal properties of actions. Austin's speech act theory is used to show that real people can act in virtual communities in ways that both establish practices and moral expectations, and warrant strong identifications between themselves and their online identities. Rawls' conception of a social practice is used to analyze the nature of the wrong and the stage-setting aspect of engaging in a practice.

Key words: Austin, LambdaMOO, moral realism, Rawls, speech act theory, virtual communities

Introduction: Virtuality and reality

Nelson Goodman opens his book *Ways of World-making* by asking, "What distinguishes genuine from spurious worlds? What are worlds made of? How are they made? What role do symbols play in the making?" These questions, central to philosophy, are also keys to the moral understanding of the online world. In this paper I will try to illuminate some morally relevant aspects of virtual, online communities by reference to more basic philosophical concepts in theories of moral realism, speech acts, and social practices.

Can someone commit a moral wrong against another person, even though their interactions take place entirely in cyberspace? This question is raised by the infamous "rape" in the online community LambdaMOO, but the attempt to answer it immediately stretches the boundaries of traditional theories of morality and ontology. This event in cyberspace, first discussed by Julian Dibbell¹ (1993), involved a real life controller who, through his online char-

acter (Bungle) and some subprograms of the LambdaMOO, managed to assault two other characters (legba and Starsinger) in the MOO. Dibbell's account and subsequent analyses suggest strongly that the controller of Bungle harmed the two people who had created and controlled legba and Starsinger. Moreover, the cyber-rape seems to have shocked the other members of LambdaMOO, and they responded by means that are common in close-knit communities: they condemned and ostracized Bungle, they called on an authority to issue a punishment, and they established stricter rules for behavior in the future. Still, it would seem easy to dismiss these actions and reactions as morally insignificant, due to the play-like nature of the online community and the mediation of events by "make-believe" characters. The circumstances of virtuality, it seems, substantially mitigate our moral concerns. I will argue, on the contrary, that the actions of Bungle's controller constitute real moral wrongs, though not wrongs on the level of rape. Further, I will give an analysis of these wrongs that leads us to a more general account of morality online, one that is useful for understanding practices such as flaming and other forms of online mischief.

Allegations of moral wrongs in cyberspace present new ground for traditional philosophy, involving both conceptual studies of communities and social practices, and empirical studies of chat rooms, discussion boards, and virtual worlds such as multiple-user dimensions (MUDs), multiple-user object-oriented worlds (MOOs), and massive multiplayer online role-
polity" of cyberspace, see David Jacobsen. Contexts and Cues in Cyberspace. *Journal of Anthropological Research*, 52(4): 461–479, 1996.

* This work was first presented at the CEPE Conference in Boston, June 2003. I would like to thank Deborah G. Johnson, Keith W. Miller, and Jameson Wetmore for helpful comments on various drafts. Support for this research has come from the National Science Foundation.

¹ Julian A. Dibbell. A Rape in Cyberspace: How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database in to a Society. *The Village Voice*: 36–42, December 23, 1993. http://www.juliandibbell.com/texts/bungle_vv.html, at 24.01.03. Since the publication of the original article, some scholars have cast doubt on Dibbell's account of what happened in the LambdaMOO community. For a general account of interpretive difficulties in the "anthro-

playing games (MMORPG). The very application of moral language to interactions in cyberspace begs two more fundamental philosophical questions. First, are interactions in cyberspace real events, or are the domains of the virtual and the real mutually exclusive? Second, are the harms that some people claim to have suffered in cyberspace real moral wrongs? I will argue that, in some virtual communities, the moral qualities of the actions are genuine because they are embedded in authentic social practices. Also, they are qualities of real *acts*, on the view that text-based virtual communities such as the MOOs, like typical "real life" communities, function by means of performative utterances.² These virtual performances have the relevant features of robust action as described by speech act theory. For these reasons, it makes sense to speak of moral patients as having suffered real moral wrongs, and accordingly to assign blame to moral agents for having committed these wrongs. Moreover, the agents and patients ultimately are real people, and not merely the characters they control.

Any philosophical discussion that begins with instances of 'real' and 'wrong' invites what Kant referred to as the *quid juris* of all critical inquiries: how is one justified in employing these terms? I intend my account to be mostly agnostic with respect to these deeper issues, in part because, despite the long-standing prominence in philosophy of such inquiries, there is no widespread agreement over the proper use of the terms. Nonetheless, we can sketch the primary doctrines. First, realism is the view that the world is the way it is, independent of our thoughts, beliefs, and attitudes concerning it. Commitment to realism does not settle the more precise ontological question of *what* is real, and philosophers have long disputed over whether "the real" includes facts, classes, sense data, relations, universals, numbers, or just simply particular material objects and their properties. Second, moral theory is the body of competing accounts of the right- or wrong-making qualities of actions and the general nature of obligation. We can combine realism and moral theory at a high level of abstraction. To say that some action X is a *real moral wrong* is to say nothing more than that X really was part of the world of events and that, under some conception of morality, X had a wrong-making quality. For reasons of space and modesty, I will not give a general account of all that is real, nor argue for a general account of what makes an action

² See A. Cicognani. On the Linguistic Nature of Cyberspace and Virtual Communities. *Virtual Reality: Research, Development and Application*, 3: 16-24, 1998, at <http://www.arch.su.edu.au/anna/papers/language.pdf>, and A. Cicognani and M.L. Maher. Design Speech Acts. 'How to do Things with Words' in *Virtual Communities. Proceedings of CAAD Futures*, 1997.

right or wrong. Nonetheless, I want to note that there are well-established areas of philosophy that do engage precisely these kinds of issues in the "real world," so the analysis of the cyber-rape in LambdaMOO will in some way cover standard philosophical territory.

The presumption that virtual communities are not real, in the relevant sense for moral analysis, must be faced head on. Antonyms for 'real' abound in philosophical, scientific, and popular writing, and offer competing boundaries for the concept of reality. Things that are fake, inauthentic, imaginary, illusory, fanciful, fictitious, or simply non-existent, are taken not to be real. Likewise, in moral theory the notion of wrongness is never defined relative to actions or entities that are not real. Wrongness is spelled out variously as that quality of an action that results in net unhappiness (utilitarianism), violates duty (Kant), steers one away from the goal of well-being (Aristotle), or violates community norms of *Sittlichkeit* (Hegel and, to some extent, MacIntyre). On some possible combinations of realism and moral theory, practically every act (or alternately, almost no act) would count as a real moral wrong. Instead of defining *a priori* the realm of real moral wrongs, I would like to examine possible cases and work backwards, as it were, from examples embedded in social practices, to the case of the alleged moral wrong in LambdaMOO. Ultimately, I hope to show that actions such as those that constituted the LambdaMOO "rape" are serious and real moral wrongs, even though the wrong in that case was clearly not one of rape itself. If I am right, this opens the door for the moral analysis of other online behaviors.

Real-making qualities of action

Let us begin with a general thesis about reality, which I will call the Causal Principle of Reality (CPR): real causes have real effects, and vice-versa. Far from expressing a tautology, this principle contains a metaphysical commitment to a certain kind of world. What it posits is that events that we experience as cause and effect are not only conceptually but also *nomologically* connected as instances of general natural laws. There are no imaginary causes of real effects, and no real causes that go without an effect. Even mystics, in granting reality to a supernatural realm of gods, genies, or spirits, thereby proclaim the reality of the causes in their explanations of effects that they find in everyday experience. The Argument From Design is one kind of causal account. For moral theory, the import of the CPR is that it connects actions (as causes) with various morally-relevant effects: pain, harm, and ultimately

moral wrongs. We cannot hold that wrongs are real, unless we can believe that real actions caused them.

But moral theory of the non-consequentialist sort has also typically insisted on realism about entities beyond visible causes and effects. Intentions have been a focus for many Kantians and others who believe that what an agent intended to do is relevant to the moral evaluation of what he or she in fact did. Common law and case law agree with this modification of realism in that the unintentional killing of innocent people is treated as much less serious than intentional killing. Inclusion of intentions in a realist ontology seems reasonable, then, in a non-consequentialist moral realism. But admitting intentions also brings up particular difficulties for the analysis of online communication. These difficulties are apparent in the quite common example of the e-mail that is taken to be insulting by the recipient, though it is intended by the sender to be funny. If intentions are opaque in everyday experience, they are more so in online behavior.

On the most basic level, the actions which we wish to focus on in the LambdaMOO and other online communities are forms of communication. The type of communication in these communities can be understood, following the view of language set out by J.L. Austin, as examples of speech acts.³ Speech acts can have several purposes. The purpose of communication is seldom merely the conveyance of information, as in a normal declarative sentence expressed by a *constative* utterance such as 'It is raining in Chicago right now'. Each constative utterance expresses a declarative sentence, and is thus an act of saying something potentially true or false. On the other hand, with *performative* kinds of utterances, such as "Take a hike!", the communication is an act which does not potentially express a true or false claim. The agents of the virtual community act in this performative way much as people do in any social realm when, by means of language, they flirt, cajole, honor, promise, chastize, and so on. These performative utterances include exhortations, condemnations, and many other kinds of socially significant expressions. Austin called what is intended in these speech acts the *illocutionary* force of the performative act. The *perlocutionary* force is the social effect, i.e., what happens in the world as a result of someone (or character) promising, condemning, flirting, and so on. According to Austin, a performative act will have its intended effect only when certain social conventions are followed, and only when the proper "felicity" or success conditions obtain.⁴ Illocu-

tion and perlocution might match up – this much is what the speaker wants – but there is no guarantee that they will. So speech act theory itself has a connection to social practices, one which we will exploit later. It also is compatible with a kind of realism. As Smith argues, speech act theory gives us a world which "contains promissings, obligatings, claims, commands, and relations of authority, just as it contains instances of biological and logical species such as *lion* and *tiger* or *judging* and *inferring*."⁵

An ontology which contains such linguistic performances is certainly broader than the traditional sparse realism of properties and substances, all of which are allegedly independent of human minds. The combination of speech act theory and realism about intentions suggests a more inclusive realism since it grants as real such entities as performative acts (and intentions so to act), even when they are put into effect by the mediation of computer programs. In simpler terms, this view holds that what a person intends to do and achieves by acting and uttering, are really part of the world. Because what an agent says, intends, and achieves is real, it is the subject matter for moral judgment, even when his or her agency is mediated by computers. The utterances and intentions of agents in cyberspace are partly constitutive of the socio-technical system that we call the online world. That world, I contend, is genuinely open to normative as well as information-theoretic analysis.

Close cases of social practice

To apply these lessons from our consideration of moral realism and speech act theory to the LambdaMOO case, we should begin by acknowledging that there were real actions by members of this virtual community, and not merely virtual actions by the characters on a screen. The controller of Bungle initiated programs by typing commands on his keyboard. Further, he acted in typing text-based messages that appeared to everyone logged on to the MOO that night. He continued to act throughout the evening of the virtual rape. No one would doubt that his entering commands and typing messages were real and intentional events. Were they, though, the causes of what happened to the two characters, legba and Starsinger? I think it is clear that they were. His finger movements on the keyboard had behind them some intentions – some illocutionary force. Through the mediation of his computer, the server, and the various programs of the MOO, his actions also had a perlocutionary force, pp. 29–61, 1990, also at <http://ontology.buffalo.edu/smith/articles/speechact.html>

⁵ Smith, 1990.

³ Cicognani, 1997.

⁴ See Barry Smith. Towards a History of Speech Act Theory. In A. Burkhardt, editor, *Speech Acts, Meanings and Intentions. Critical Approaches to the Philosophy of John R. Searle*,

and (we assume) it was what he had intended. So if the thesis of CPR above is correct, then the event description 'what happened to legba and Starsinger' picks out a real event, and in this case a real effect of the acts initiated by Bungle's controller. The "total" performance of the controller of Bungle, including intentions and finger movements on his keyboard, issued in effects to other characters and other controllers.

Still, the attribution of moral realist claims to this virtual behavior may seem to go too fast. We need to test this preliminary application of the moral realist view of speech acts against cases in which the virtual events are both real and unreal. It will be helpful to compare some cases in which computers are not involved, since the online world is not the only virtual realm. Consider the following examples:

1. Mike sees his acquaintance Joan one day, and surmises that she is depressed. He knows that her mother recently passed away, and feels like he should mention this fact and express his condolences. Unfortunately, in the conversation that ensues, Mike asks whether Joan and her mother were emotionally close, and this issue is precisely what has been troubling Joan in recent days. His question causes her to worry even more over the quality of her relationship with her mother. After Mike leaves, Joan is much more agitated and unhappy than before their conversation.
2. Roberto has thought up an elaborate fantasy world, complete with characters that he identifies with as well as ones with which he does not. This fantasy world has the semblance of a work of literary fiction. Several times, he has revised the plot to better fit his sense of narrative coherence. One day Roberto decides to tell Alice about his fantasy world. Believing that Alice will sympathize, Roberto tells her this "story" in great detail. She understands that Roberto is invested psychologically in the fantasy world, and learns its details very well. After a week, Alice sees Roberto again, recites the details as he had given them the week before, and extends the story with several gruesome and unlikely twists of the plot. Roberto takes great offense, asserts ownership over the story, and disavows her extension of the fantasy world. In doing so, he shows that he is harmed emotionally by Alice's co-opting of the story.
3. A soldier is captured behind enemy lines, and is subject to torture, apparently without a purpose. Her captors do not wish to gain information from her, they merely wish to terrorize her and send her back to her comrades. She withstands some physical torture, but her captors seem to be more adept at psychological torture. They frequently tell her that her comrades have mostly deserted the war effort, and even print up fake newspapers that report of the desertions and the impending collapse of her country. She shows signs of distress in hearing these accounts, though it is unclear whether she fully believes what her captors are saying.
4. A tribe of bushmen in a remote area of the world continue to practice the oral histories and mythologies of their ancestors. They mix observed fact, remembrance, mystical explanation, and tacit fantasy in their oral tradition. One day, a young man extends their received oral history in a way that is not to the liking of the rest of the group, primarily because it seems to denigrate some of their ancestors. He asserts narrative control over their shared history and mythology in order to effect his dominance. This person has learned the tradition well, so he cannot simply be dismissed. Nonetheless, there is a feeling amongst the members of the group that the extension of the story is illegitimate and offensive.

Do any of these cases present real moral wrongs in their seemingly "virtual" effects? In the first case, we can assume for the sake of argument that Joan is *harmed* by Mike's question about her mother. Further, there is no doubt that Mike's utterance was a performative speech act, that is, a real event that effected a change on its object. In this case, the object was Joan (or perhaps her psyche), and Mike's act is therefore the cause of some of Joan's distress. But that Joan is harmed, and that Mike caused the harm, is not sufficient for concluding that Mike's act was a real wrong. In this case, Mike unwittingly set off a reaction in Joan to which she was predisposed. He did not, however, intend to worsen Joan's emotional state, nor to cause her harm. And he was not negligent, *ex hypothesis*, in failing to anticipate Joan's predisposition. So while there is a causal agent, and an effect that counts as a harm, there are also some intervening causes and enabling factors: Joan's strained relationship with her mother, and Joan's preoccupation with their relationship. Here there is harm, but the weak causal link and lack of intention to harm lead us to believe that there is no real moral wrong.

In the second case, there is a cause (Alice's report of her extension of the story) and a perlocutionary effect (Roberto's taking offense at the extension). There is also something that comes close to a virtual world: Roberto's fantasy world. Insofar as offense counts as a harm, we might be tempted to say that Alice's harming of Roberto counts as a wrong, but that would be too hasty. What is significant about this case

is that the "world" that Alice corrupts is not the one of Roberto's making. Though he is upset, he ought not be, for the thing that Alice changes is not the story of Roberto's making. In co-opting his story, she has made some version of it her own but has not deprived Roberto of *his* fantasy world. Perhaps Roberto cannot be faulted for becoming upset. Nonetheless, there is an ontological difference between the fantasy that is his and the one that is hers. His version of the fantasy remains unchanged.

In the third case, the soldier is clearly harmed and wronged by the physical torture she endures. The question is whether her psychological torture, effected by deliberate distortions concerning the fate of her comrades and her country, constitutes a further wrong. I think it does not, primarily because she has no reason to take the reports of her captors seriously. She cannot be faulted in the event that she comes to believe such distortions, given her predicament, but still the context of her interactions with them does not warrant expectations of honesty. Without those expectations, it is hard to see why she should take their claims seriously. To put it another way, she is not in community with them, so they cannot break faith with her. As this case shows, being in a community with others is one condition which generates legitimate moral expectations of behavior. But not every collection of people counts as a community.

Finally, I want to claim that the last case is a plausible example of a real moral wrong occurring in a virtual world. The young man who has extended the oral history of the group has done so intentionally. His modification of the story is causally efficacious on its own, given the background of the *shared* tradition, knowledge of which he has exploited. He has violated reasonable expectations and implicit rules of the community that are established by a social practice. In the first three cases, some one aspect or other excluded them from the realm of real moral wrongs. There were failures of intentionality (Case 1), identity (Case 2), or community (Case 3). The last case is one of intentional harm to a shared virtual world of a well-defined community.

Rawlsian practices and LambdaMOO

Following Rawls,⁶ we could describe the relationship between implicit rules, expectations, and practices in the following way. When individuals are engaged in a practice, the meanings and moral boundaries of behaviors, understood by them as expectations and

implicit rules, are constructed from within the practice. In that context, an action like a punishing or a promising "is a performative utterance which presupposes the stage-setting of the practice and the proprieties defined by it." The behavior of the young member of the tribe did not merely cause offense in other members of the group, in some way that would be explainable outside of the stage-setting of the practice. While some cases of offense may seem to rise to the level of a moral wrong, and some do not, this determination is impossible to make without the guidance given by the practice. Arguably, this offense, violates reasonable expectations of behavior, and trades on the level of trust, all of which are generated from within the communal practice. These kinds of wrongs are both made possible by the practice, and in time would be destructive of it.

We are now ready to look more closely at the LambdaMOO case. As has been widely discussed by Huff, et.al, and Dibbell, the "rape" perpetrated through Bungle and various programs had a traumatic effect on the controllers of legba and Starsinger, and created repercussions felt throughout the LambdaMOO. We have seen that the philosophical questions raised by this act form a nest of ontological and moral issues. Is a harm done to a virtual character a harm to its real controller? Is that harm also a wrong? To get clearer on these issues, we must focus on the relationships between the character and controller, and between the controllers themselves, within the practice. These relationships will vary over controllers and over virtual communities. Nonetheless, we have a sufficiently rich description of the LambdaMOO to be able to characterize the relationships in this case.

Controllers stand in relations to their characters that are very rare between people and their creations. The characters are in fact conduits of the meanings and illocutionary force of the controllers' acts; under speech act theory, they *deliver* utterances that are performative in that they honor, entice, denigrate, amuse, flirt with, and confound other "objects," i.e., other characters. Let us call this kind of performative utterance, a kind in which the object is other-regarding, a transitive performative speech act, or t-performative for short. In this respect, the characters allow the individuals, partly through t-performative acts, to construct the virtual community and its nexus of behaviors and expectations that define the practice. So far, though, the characters in the LambdaMOO are no different from technologies like fax machines and bullhorns. As mediating tools of t-performances, they help people construct and influence the social world around them.

In addition to mediating transitive performances, the characters also facilitate constructions of new

⁶ See John Rawls. Two Concepts of Rules. *Philosophical Review*, 64(1): 3-32, 1955.

personae for their own controllers. That is, the characters perform acts "directed" at their controllers, in addition to those directed at other characters. These created personae are not always veridical; some controllers experiment with temperamental, sexual, political, racial, and species modifications. Whatever the outcomes of "character development" in these contexts, the relation to the controller remains a self-regarding performance. In building various aspects of the character, the controller is engaging in *reflexive performative* (r-performative) speech acts. The object which the controller hopes to affect is his or her virtual character. Of course, to some extent the controller is acting on him or herself in experimenting with various adornments, genders, moods, or expressions of ideology. The range of actions in r-performances may be similar to the t-performances, or may be different. A controller may honor all other characters in the text messages that he or she sends, while at the same time performing acts of self-effacement, or the controller may vary his or her speech acts. In no case, however, would we expect entirely new forms of interaction — ones never before experienced in real societies.⁷ The eventual content of the sum total of interactions within a virtual community, effected by nothing more than performative utterances disguised as information exchange, may be quite different from any "real" community that has existed. But its elements will be as familiar as the constitutive acts of real communities. In these more familiar communities, reflexive and transitive speech acts are commonplace.

Seen as technologies, the LambdaMOO characters, which are elements of a larger technology — the database — which is part of the (larger still) internet, have dual natures. These dual natures present the key to understanding the peculiarity of such technologies. The transitive performances of the characters build the community by establishing boundaries and expectations, i.e., they effect what Rawls calls stage-setting. No doubt, these constructions are themselves effects of indirect communication; except where one character says to another "You should not do that!," we should suppose that stage-setting is achieved by a confluence of t-performances such as insults, entreaties, and acts of praise. The reflexive performances, on the other hand, adorn one's character in acts of feigning, gender-switching, flirting, and the like. In these ways, r-performances connect the character to the controller to a degree unimaginable with other, merely t-performative technologies. One may be able to influ-

ence many people by use of a bull horn or fax machine, or other transitive technologies. But we would never say that normal agents could *identify with* technologies such as these. As Dibbell has explained, it is precisely this aspect of identification that prefigures in the harms felt by the victims of the cyber-rape. Thus, the strength of the character-controller identification and the moral boundaries set for the controllers by the practices of the virtual community lay out the basis for a moral judgment. In other words, the character-controller identification allows harm to a character to become a wrong to a controller. Granting that this judgment could only be made with certainty from within the LambdaMOO, there are strong grounds for thinking that those members were right in condemning, in a moral realist fashion, the acts of Bungle's controller.

Other candidates for wrongs in virtual communities

In the years since Dibbell's account of the LambdaMOO cyber-rape was published, many psychologists have begun to look closely at online behavior. In various online contexts, behaviors such as flaming, spamming, spoofing (using another's IP address), and "grief playing" (the deliberate obstruction of player-to-player online role-playing games) have been studied as forms of deviance.⁸ Much has already been written about the "disinhibiting" effect of online existence; hate speech and other kinds of anti-social behavior seem to be more prevalent in the virtual than in the embodied world. Are these kinds of behaviors only of interest to psychologists? Can the realist speech-act theory be applied to other kinds of online behavior in order to facilitate moral analysis?

Let us turn to cases of actions which appear to be the least harmful. Anyone who has spent much time online in discussion boards, chat rooms, and such cyberplaces has no doubt witnessed flaming and off-topic posts. For whatever reasons, some participants effectively burden the commons by misusing shared IT resources. It is doubtful, however, that these behaviors are anything more than breaches of netiquette. First, though they involve intentional speech acts, it is not always clear in context that they violate community norms established by a practice. Some discussion and chat forums are simply more anarchistic than others, and whatever rules there may be are up to the moderators to enforce. In joining an unmoderated

⁷ This view was put forth by Simmel in his theory of social forms. See "How Is Society Possible?" In Donald N. Levine, editor, *Georg Simmel: On Individuality and Social Forms*, 1971 (original work published in 1908).

⁸ See J.R. Suler and W. Phillips. The Bad Boys of Cyberspace: Deviant Behavior in Multimedia Chat Communities. *CyberPsychology and Behavior*, 1: 275-294, 1998.

discussion or chat, one assumes the risk. Second, participants have the relatively easy option of "voting with their feet." When a particular online forum becomes crowded with participants who flame, spam, and otherwise disrupt, the more committed members of the group will often contact one another off the list and agree to move the discussion somewhere else. As long as this "move" has no serious financial implications, the committed members have no serious moral claims to make. Finally, individual members who suffer textual attacks by others are often defended by the members of the group. While the keyboard may be mightier than the sword, it may have unintended effects on those who use text to abuse others online. Reputation does seem to matter online, as in the embodied world.

The more difficult cases arise in long-running, multiplayer role-playing games where participants are deeply committed to the characters or "avatars" they create, much like the controllers in the MOO's are committed to their characters. These avatars and the online world they inhabit are graphical as opposed to text-based. Unlike the MOO's, these cyberworlds give controllers possibilities for interacting (and insulting) which go beyond speech acts. Nonetheless, the graphical behavior is symbolic and can transmit meaning. Many of these games began as hacker enclaves and have expanded to huge commercial enterprises. Suler and Phillips cite many types of deviant behaviors in such contexts, and recommend certain "interventions" for moderators or "wizards" in dealing with them. In online role playing games, the deviant behaviors can range from low-end (being a clueless newbie, parodying other players) to high-end (hate speech, verbal exhibitionism, speech indicating pedophilia, or even crashing the server on which the game runs). Their analysis is thoroughly psychological and would seem to indicate that such SNERT's (snot-nosed eros ridden teenagers) are socially troublesome in both online and everyday worlds. What should the well-behaved player think of the disruption of the game by miscreants?

The inclination to feel insulted, frustrated, or indignant by a snert's actions reflects the tendency to invest a lot of psychological energy in one's online world. Users take it personally and feel very emotional when it comes to their virtual community. To them, it's as real as the real world. Perhaps the best defense against snerts is to unravel that psychological investment a bit. You can always turn off the computer and walk away. The Greek philosopher Epictetus said that people are not disturbed by things that happen to them, but by the views they take of those things. In other words, sticks and stones can

break your bones, but the snerts of virtual reality can rarely hurt you . . . unless you let them.⁹

Moral judgment is not the last recourse. At least one option remains open to all online participants: opt out. If one is offended by something online, then the offline world awaits.

An important difference between role-playing games and the LambdaMOO is in the respective expectations of the two practices. The role-playing games fit the general libertarian ideology of the internet; participation is a free choice, and offense does not count as harm. Minimal rules are established, and "fair play" is anything that falls within the rules. There are consequences for violating the rules, including "pinning", "gagging", and being "killed" or even banned by the moderators. Also, the very point of some of these role-playing games seems to lie in the expression of deviance, as Suler and Phillips suggest. It is a reasonable expectation, upon signing up to play the game, that your avatar at some point will be abused, violated, dismembered and exterminated. On the other hand, the early versions of the MOO's did not have explicit rules about the behavior of characters. Implicit rules and expectations were learned through participation in the practice. Certainly, there must have been deviants in the MOO's before the famous case of Bungle. But the very lack of explicit rules placed a greater burden of self-restraint on the participants. Further, the degree to which Bungle's controller violated the implicit rules showed that he exploited a social situation of communal tolerance. The normal interactions within the MOO appeared to track much more closely with everyday social interactions. Conversely, the more recent role-playing games *feature* kinds of deviance. They are part of the very software that runs the game – software that allows avatars to urinate on and behead one another. When psychologists go looking for deviance in these kinds of online worlds, it is difficult to imagine where they could draw lines.

So while there is a tendency to "explain away" much online behavior in psychological as opposed to moral terms, it is not clear that the psychological and moral analyses are mutually exclusive. Likewise, it is not clear that we ought to refrain from moral judgments about mass murderers, even as we make psychological judgments about them. In some contexts, online communities give rise to expectations of intentional behavior, where the expectations are backed up by the moral force of established practices, as opposed to the "legal" force of the moderators. It is impossible to designate *a priori* just what these online moral communities are; but it is a safe bet that

⁹ Ibid.

they will share many features of everyday communities which give rise to moral relations among participants. Discussion boards share some features with face-to-face discussion groups, but lack others. Role-playing games seem only to share features with a bizarre Hobbesian world, and hence must lack moral relations, just as did the state of nature on Hobbes' view.

Conclusion

As virtual communities become outlets for social impulses and remedies for individual isolation, it is to be expected that the technologies will become increasingly sophisticated in their ability to mimic "real life" communities. As these "real life" communities become more scarce, the question will arise as to which kinds of communities, virtual or real, support healthy and meaningful forms of sociability and individuality. The relatively long history of the LambdaMOO portends an increasing identity of real and virtual personae. We can only hope that participants of the MMORPG's do not increasingly identify with their avatars. These online communities also suggest analogs to the typical distinction between law and morality. As I have argued, the realm of morality in virtual communities is defined in the context of that community. Expectations, implicit rules, and the sense of propriety are defined in the "stage-setting" of the virtual practice. The explicit rules of the database administrators are akin to positive laws. Punishments can be moral or legal; characters, and thereby controllers, can be admonished or even banned, i.e., given the virtual death penalty. Finally, the closeness of virtual and real communities is expressed in the parallel between the irreplaceability of characters and the mortality of real people. The banning of one's character removes more than the name from the database. A controller can insert another character, and animate

and adorn it in ways similar to the first "deceased" character. But in the end it has a different identity and history, in the eyes of the particular virtual community, and will never be a replacement for the "real" thing.

References

- J.L. Austin. *How to Do Things With Words*. Oxford University Press, Oxford, 1962.
- A. Cicognani. On the Linguistic Nature of Cyberspace and Virtual Communities. *Virtual Reality: Research, Development and Application*, vol. 3, pp. 16–24. Springer-Verlag, London, 1998, also available at <http://www.arch.su.edu.au/~anna/papers/language.pdf>, at 22.01.03
- A. Cicognani and M.L. Maher. Design Speech Acts. 'How to Do Things With Words' in Virtual Communities. *Proceedings of CAAD Futures 1997*. Kluwer Academic Publishers, Dordrecht, 1997.
- Julian A. Dibbell. A Rape in Cyberspace: How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database in to a Society. *The Village Voice*: 36–42, December 23, 1993, also available at http://www.juliandibbell.com/texts/bungle_vv.html, at 24.01.03.
- Nelson Goodman. *Ways of Worldmaking*. Hackett Publishing, Indianapolis, 1978.
- C. Huff, K. Miller and D.G. Johnson. Virtual Harms and Real Responsibility: a Rape in Cyberspace. *IEEE Technology and Society Magazine*, 23(2), 2003.
- John Rawls. Two Concepts of Rules. *Philosophical Review*, 64(1): 3–32, 1955.
- Georg Simmel. How Is Society Possible? In Donald N. Levine, editor, *Georg Simmel: On Individuality and Social Forms*. University of Chicago Press, Chicago, 1971 (original work published in 1908).
- Barry Smith. Towards a History of Speech Act Theory. In A. Burkhardt, editor, *Speech Acts, Meanings and Intentions. Critical Approaches to the Philosophy of John R. Searle*. Walter de Gruyter, Berlin/New York, 1990.
- J.R. Suler and W. Phillips. The Bad Boys of Cyberspace: Deviant Behavior in Multimedia Chat Communities. *Cyber-Psychology and Behavior*, 1: 275–294, 1998.