

# The Duplicity of Online Behavior\*

Joseph Ulatowski  
University of Texas at El Paso  
El Paso, TX 79968 · jwulatowski@utep.edu

May 30, 2014

## Abstract

People commonly believe that any form of deception, no matter how innocuous it is and no matter whether the deceiving person intended it otherwise, is always morally wrong. In this paper, I will argue that deceiving in real-time is morally distinguishable from deceiving online because online actions aren't as fine-grained as actions occurring in real-time. Our failure to detect the fine-grained characteristics of another avatar leads us to believe that that avatar intended to do a moral harm. Openly deceiving someone on Facebook or Twitter is not a way to build wholesome virtual friendships but to destroy them. This paper will show how the traditional understanding of the doing / allowing distinction fails to apply in cyberspace.

This is a DRAFT -- Do not cite without permission.

Comments, criticisms, and questions are welcome.

© 2014 Joseph Ulatowski

People commonly believe that any form of deception, no matter how innocuous it is and no matter whether the deceiving person intended it otherwise, is always morally wrong. In this paper, I will argue that

---

\*I am grateful to Dr. Jane Kaworski and Jim Nico for inviting me on their internet radio show, *The Social Network Show*, where we discussed issues that led me to thinking about the problems associated with online deception and drafting this paper.

deceiving in real-time is morally distinguishable from deceiving online because online actions are not as morally innocuous as actions occurring in real-time. Our failure to detect the fine-grained characteristics of another virtual avatar leads us to believe that the person intended to do a moral harm. Openly deceiving someone on Facebook or Twitter is not a way to build wholesome virtual friendships but to destroy them. This paper will show how the traditional understanding of the doing / allowing distinction fails to apply in cyberspace.

## 1 Innocuous passive deception

In this section and the next, I will provide some examples of deception we commonly employ in our lives that do not yield a moral harm. If the consequence of these forms of deception are not morally wrong, then we will have to entertain the possibility that online actions of virtual avatars are not morally harmful. What I will go on to show is that online deception lacks some of the characteristics of real-time deception. Missing these characteristics show that online deception cannot help but be morally harmful.

Suppose Alma and Bonita are deeply in love with one another and have been together for a very long time. They have such a tight emotional bond with one another that each of them knows that the other would never do anything to harm the other or the relationship they have cultivated over the years. One morning Alma and Bonita decide that they should treat themselves to a night of dinner and dancing. Both agree to meet for dinner at 6:00pm. Alma arrives at the restaurant around 5:50pm only to discover that Bonita has not arrived yet. The maitre d' seats Alma, even though Bonita has not arrived. After she has been seated, Alma sends a text message to Bonita asking her whereabouts. Bonita, though comfortably seated in a business meeting with colleagues that will not likely end for another ten minutes, responds with the following message: "OTW."<sup>1</sup> Although, in fact, Bonita is not on the way to meet Alma, her text's content expresses to Alma that she is on her way. This implies that Bonita has deceived Alma.

Under the circumstances, We ought not believe that Bonita harmed

---

<sup>1</sup>OTW is txtspeak shorthand for "On the way."

Alma, even though Bonita has clearly deceived Alma into believing that she is on her way to the restaurant. Conversationally, in txtspeak anyhow, “OTW” means not only that a person is physically moving toward a particular destination but also that one will be on the way shortly, one is on the way but ran into someone on the way out the door, or—perhaps as is more frequently the case in larger metropolitan areas—one is stuck in pedestrian or vehicle traffic and should arrive shortly.

Despite that there are these other alternative meanings to “OTW,” what Bonita has texted to Alma is what one might call a bald-face lie. Bonita is *in the meeting*. She’s not walking out the door. She’s not stuck in traffic. She’s not even intending to leave the meeting early to meet up with her long-time friend, partner, and lover for the dinner they planned earlier in the day. So, the alternative conversational implicatures of “OTW” do not apply in this case.

Further evidence seems to rule against a charitable reading of Bonita’s text message. There is no doubt that Bonita will be on her way as soon as the meeting ends ten minutes hence. The trouble is that the form of the text message suggests she’s on her way to meeting Alma *right now*. Texting Alma that she’s “OTW” suggests to Alma that Bonita is “walking” or “driving” towards the restaurant. The use of an acronym in place of the phrase, “On the way,” suggests the text was sent in haste with some urgency. Moreover, it triggers in Alma’s mind that she should not follow up that message with another because it could distract Bonita, whether she’s driving or walking to the restaurant.<sup>2</sup>

Suppose the story were to continue in the following way: Bonita arrives thirty minutes late to the restaurant and pleads with great supplication for Alma to forgive her for being so tardy. Bonita’s telling Alma where she was when she txt-ed Alma does not improve her thinning patience. Perhaps surprisingly Alma is not quite angry with Bonita for making her wait at the restaurant. Bonita could have called Alma before the meeting began warning her of a belated arrival. That might not even cross Alma’s mind; instead, she’s happy to see Bonita, have some dinner and fun with her very close and loving friend.

Under other circumstances, Alma might not be so accommodating.

---

<sup>2</sup>One might argue that since Bonita is in a business meeting her use of “OTW” is justifiably urgent because she does not want to be disturbed by Alma again. Even if one were to argue that her being in a business meeting justifies Bonita’s use of acronyms, it hardly saves her from the charge that she’s lied to her close friend, Alma.

For instance, if Alma discovered that Bonita had been unfaithful or that Bonita had drained her bank accounts and moved to Acapulco, Alma probably would not have been as forgiving as she was this evening. Why is this the case? In all of the examples, Bonita lied to and deceived Alma. In some cases we seem to overlook deception, while in others it's difficult to turn a blind eye to it.

The case of Alma and Bonita is just one way in which we deceive our closest friends, spouses, or siblings using social networking platforms, such as txt-ing.<sup>3</sup> The deception did not result in harming either Alma or Bonita, but what Bonita did was morally wrong if we believe that we ought not lie to others, especially if it is as self-serving as the lie Bonita told Alma.

Perhaps one would not want to call what Bonita did an act of lying. She might have overlooked the fact that her meeting might continue for thirty minutes. On one commonly held definition of lying, it is the act of making a false statement with the intention to deceive (Kagan, 1998). Since Bonita may not have intended to deceive Alma, she did not lie to her. We might even say that Bonita did not deceive Alma. Regardless, overlooking a particular fact seemingly frees Bonita from any moral wrongdoing because she did not *intend* to deceive her lover and long-time partner. Given that no moral harm resulted from Bonita's txt, we judge her passive act of deception to be harmless and innocuous.

## 2 Innocuous active deception

In the previous section, we were introduced to a real-time example, i.e., one commonly occurring in the actual world amid practical affairs, where a person's permitting someone else to be deceived did not result in any kind of moral harm. So, we concluded that the act involved no moral wrongdoing. We might call such cases 'innocuous passive deception'. A critic might contend that such innocuous actions that involve no moral

---

<sup>3</sup>For purposes of this paper, I consider all forms of technology that provide us with a greater reach to others a "social network." Not only will the usual suspects be discussed, like Facebook, Twitter, SnapChat, Pinterest, LinkedIn, etc., but also the less obvious examples of social networking: smartphones, the internet, chatrooms, discussion boards, blogs, etc. I realize that this is fairly wide definition of social networking with which some might take issue.

wrongdoing are not forms of deception. The critic maintains that all forms of deception yield harmful or painful results, so those sorts of actions that fail to fit into such a category must not be a form of deception.

In this section, I will argue that there are active forms of deception that fail to be a case of moral wrongdoing. Despite the critic's conception of deception always involving a moral harm and moral wrongdoing, these examples of active deception are just as innocuous and morally inert as the example of Alma and Bonita. Once I have shown that innocuous forms of active deception involve no moral wrongdoing, I will be able to discuss cases of online deception that do not divide easily between innocuous and egregious forms of active and passive deception. If this is true, then the kind of nuance our real-time behaviors permit do not translate into permissible online behaviors. We ought, then, refrain from performing any action online that could be understood as an egregious and morally harmful action.

One might believe, as the critic appears to believe, that as opposed to the moral inertness of innocuous passive deception it must be the case that all forms of active deception are morally harmful and, therefore, a good example of moral wrongdoing. There are cases of active deception which do not clearly show the action to be a moral wrongdoing. The singular feature that stands out in these active forms of deception is that the agent acted intentionally. The agent intended to lie or to deceive others. Because of this, we might believe that it is morally wrong.

Some deceptive acts are undertaken (and intentionally so) not to harm others but to protect others (and oneself) from harm. Soldiers, for example, wear camouflage uniforms to hide from and to deceive enemy combatants. Some camouflage is more effective than other forms of camouflage because it enables soldiers to blend in with the elements of the environment. Deceiving enemy soldiers using camouflage preserves enemy combatant soldiers from harm, and not wearing camouflage, especially in light of strategies adopted in modern warfare (e.g., guerrilla tactics), would mean certain death.<sup>4</sup>

---

<sup>4</sup>Conventional war-time strategy where soldiers meet on the battlefield in regular and patterned formation had to be abandoned shortly following the invention of high-velocity automatic weaponry, e.g., the machine gun. Too many casualties were suffered because regular formation made the soldier vulnerable to any unskilled marksman. According to legend, for instance, at the Battle of the Somme (1916), an estimated 20,000 British soldiers were killed in the first twenty minutes of the battle by Axis machine

Wearing camouflage is certainly a form of deception. In fact, it is an intentional form of deception where an agent intends to deceive enemy combatants. Donning camouflage battle-dress protects the soldier wearing it by not permitting enemy combatants to detect his presence. If enemy combatants cannot detect a soldier's presence, then they will not know that their enemy is close. If the enemy cannot detect the proximity of enemy soldiers, then they will be less likely to open fire and bring about a fire-fight or, worse yet, inflict harm upon the soldiers or nearby innocent civilians. I believe it goes without saying that wearing camouflage is not morally wrong; in fact, one might go so far as to contend that not wearing camouflage in battle is not only uncouth but reckless and, therefore, possibly morally forbidden, even as an activity in an unjust war.

Besides the use of camouflage in war time, there are other forms of innocuous active deception. Think here of the boss who lays off an employee because she hates that the employee always submits work after institution-wide deadlines. When the boss "fires" the employee (after filing the appropriate paperwork, of course), she might say "the company has terminated the position you occupy" or "the company's moving in a different direction." Eliminating one's position in the company is not identical to "firing" the employee because the employee might believe that the company can no longer afford to pay for the skills he possesses. In a sense, the employee does not feel that the company or administrators have "fired" him but that the company can no longer afford the services the employee provides. Of course, the boss has bamboozled the (now ex-) employee because it is the employee's failure to submit material on-time that has brought about the "separation."

The boss' firing is an active form of deception because she intends to deceive the employee. The boss' reason for action is to ensure separation of the employee from the company because the employee has failed to uphold terms of the contract which include following the guidelines set in the original job description. Undoubtedly, a part of that job description will be to follow company policies and procedures. One of those policies will likely involve an explanation of submitting work by the prescribed

---

gunners who took advantage of the soldiers "going over the top" and "running across the desolate no-man's-land." With nowhere to run or to hide, German gunners very easily cut down, quite literally, the advancing British.

deadline administrators and managers mandate. Since the boss tells the employee a bogus reason for removing the employee from the position in the company, the boss has lied and intentionally deceived the employee.

Despite the outright deception perpetrated by the employer in this case, our reaction to it might not be to ascribe some sense of moral wrongdoing to the employer; instead, our intuition tells us that the employer seemingly has preserved and protected the employee's psyché. No moral harm comes to the employee. Thus, there seems little reason to think that what the employer has done is morally wrong.

The business world is pregnant with examples of deception. Besides the one just outlined regarding the firing of an employee, there is the more common situation in which job-seekers tend to inflate their own self-worth, either in a resume or in a cover letter accompanying a job application. The job applicant might embellish the responsibilities of a previous job or bedeck line upon line of a resume with nearly false but not untrue statements about honors for which the job candidate had been "nominated." The statements are considered harmless because search committee members know that job candidates want to look better to them than they actually are because they want seek at least an interview if not landing a job with the company.<sup>5</sup>

Like the example of firing the employee is innocuous, so too is the job seeker who reports near falsehoods on a resume. The job candidate intends to deceive the reader into believing that the near falsehoods are accurate representations of the accomplishments. Since members of the business world have come to expect some exaggeration from job candidates and no moral harm arises from these near falsehoods, it seems inappropriate to call what the job seeker is doing is morally wrong. Even the most fastidious search committee member would overlook the candidate's taradiddle.

---

<sup>5</sup>In the academic world, it is an open secret that many letter writers, particularly those from well-established prestigious universities, knowingly deceive search committees with high praise in confidential letters of recommendation. So much so is this true in philosophy that letter writers have begun writing letters in a code only very few members of the philosophy profession understand. When educated search committee members review the letter, they know that the letter writer believes candidate-X is a top-notch researcher and should be placed at an R-1 research university or whether the candidate is best fit for a liberal arts university and college. Unfortunately, not everyone is aware of this code and remains deceived by what would otherwise be considered a *good* letter of recommendation.

Perhaps in reading the above examples one might conclude that innocuous active deception is mostly found in professions and circumstances where the agent justifies a certain means to reach a highly beneficial end. In war, soldiers want to survive. As a means of survival, they must wear camouflage. In business, job candidates and employers seek success. As a means to success, they must report half-truths. If there are other more intimate settings in which active deception arises, then we might be less dismissive of the other examples.

There are, in fact, intimate settings in which active deception seems to occur. Many of them focus on relationships in families. Suppose that a parent takes away a loud toy from a child because it is too loud, but the parent tells the child that the toy is broken. The toy is not broken, but the parent intentionally deceives the child so that they do not have to listen to the loud toy any longer. No moral harm comes to the child, even if the child really wants to play with the toy that has been taken away. In this example, the parent intends to deceive the child by telling her that the toy is broken and is no longer fit for play, even though the parent knows the toy is in perfect working order.

Similarly, think of the parent who tells a child that “Santa Claus will not be bringing her any toys this Christmas because she’s been a naughty little girl.” This is doubly deceiving. First, there is no Santa Claus to bring gifts to the little girl. Santa Claus is a fictitious entity who does not exist. The parent has perpetuated a lie about the existence of an entity that does not actually exist. Second, the deceit seems a way for parents to get what they want from the child—in this instance, they are treating the child merely as a means to an end. Acting nice, the child will get what she wants and the parents will get what they want—perhaps, a better behaved child. The parent’s telling the child Santa will not be delivering presents to her is doubly deceiving, but it is a common tactic parents use to get what they want from a child.

In the last two sections, I have introduced examples occurring in real-time that seem to show there are morally innocuous forms of active and passive deception. The question with which we will be concerned in the next two sections is whether active and passive deception occurring online and in virtual social worlds are equally morally innocuous. It is my contention that the conditions for active and passive deception do not easily translate from the real-time world to the virtual world. Netizens ought to exercise greater care in their deceptive behavior than they would



in real-time affairs because online behaviors have real-time consequences, some of which—unfortunately, as we have witnessed—can be catastrophic.

### 3 Online deception, innocuous or dangerous?

There is a widespread belief that the real-time world and the virtual world are completely distinct. We live life differently in each, and we believe that the things we say and do in the virtual world does not carryover to the real-time world. In the previous section, I set out to show that there are innocuous forms of active and passive deception. For this section, I will begin by briefly illustrating why we believe the online virtual world and the real-time world are distinguishable. Then, I will show that our belief in this distinction is wrong. An argument for this view will come in §4 in a discussion of doing something versus letting something happen. Finally, in the last section (§5), I will show how virtual world deception, whether active or passive, yields seriously harmful consequences. If my analysis is correct, then we should refrain from deceiving others online.

There is a distance between us and virtual netizens that permit us to do and to say things that we would not normally do or say to people on the street. When the internet began to take shape and online interactions were more common toward the end of the 1990s, I can remember entering a chat room named “Philosophy” or “Nietzsche on Value.” I presumed that we would have a discussion about the chatroom’s title. Sometimes the discussion was fruitful, but most of the time it was not. Online chat rooms had the tendency to devolve into school playground arguments. Arguments would amount to “yah-huh” and “nun-ahh.” What should have passed for an enlightening chat became nothing interesting whatsoever.

There is at least one reason why school yard arguments take place in online forums. The discussants use anonymous handles. Anonymity provides them the opportunity to use baseless claims and offer unguarded opinions. The degeneration of the discussion might have been avoided if members of the chatroom had not anonymized their own profiles. Profiles, at least in the 1990s, were easy to anonymize. Discussants could easily cloak their own identity. For example, the only thing known about the person chatting was their handle, “fideo1991” or “veritasetvirtus.” Handles typically do not provide an adequate amount of information to

identify the handle with any person.

There was no other place one could be as anonymous as they were when they were online. This is true today, too. Look at any *New York Times* editorial generating interest and comments among its readers, and it is easy to find a few anonymous commentators posting ugly opinions surreptitiously.

Fast-forward to the current era of social networking sites like Facebook or LinkedIn. Although it is not recommended, one can create fictitious Facebook or Twitter accounts.<sup>6</sup> A person is capable of manipulating one's online identity in a way he is unable to do so in real-time. A person wandering around town in a mask would likely garner some attention from not only ordinary people but the local authorities as well. Given that manipulation of one's identity is so easy to do online, this seems to be one underlying reason why we believe the online world and the real-time world is distinct.

Clandestine interactions online have a far greater possibility of going undetected than secret meetings in real-time. No matter what it is a person does in real-time there is a chance that the encounter will be filmed by closed-circuit televisions, that the N.S.A. is watching, or the encounter is being surveilled by a private investigator. Even if the person wears a costume, mask, wig, or fake mustache, the existence of facial recognition software or other means of identifying a person will give away the person's true identity. No such luck online. It seems possible to cruise the internet undetected without any repercussions for the individual who covertly engages in nefarious and illicit affairs.

For these reasons, I believe culture has disentangled the world wide web from the web of humanity. No longer do we believe that our online behavior has any serious consequences, good or bad, in real-time. The trouble is that our online behavior does have consequences.

There have been tragic consequences resulting from the creation and maintenance of fake social networking profiles. Megan Meier, a 13-year-old girl suffering from depression and attention deficit disorder, committed suicide after a "fake" person she began corresponding with on MySpace ended their friendship. It was discovered later that the "fake" person's profile had been set up by neighborhood bullies who berated

---

<sup>6</sup>In fact, if I am not mistaken, I believe Facebook has a policy against the creation and maintenance of accounts of fictitious virtual people.

Megan in chats had on the once popular social networking site. In another incident, Brandon Wentzell consumed a lethal cocktail of vodka and Dilaudid after his long-distance Facebook girlfriend, Clarissa Chistiakov, canceled an in-person meeting and discontinued the relationship they had fostered. It turns out that Clarissa did not exist.

The consequences of any action can be either good or bad. Whenever we consider deceitful behavior, we think that the outcome will be bad for whatever agent is the victim of the deceit. Online behavior seems to make one more prone to being a victim of bad action because of the anonymity involved in online interactions. If it is possible for us to refrain from behaving badly online and to prevent bad things from happening to other members of the social network, then we have a moral obligation to prevent unnecessary suffering of our fellow internet surfers (cf. Singer 1971). Not doing anything seems tantamount to permitting very bad things to happen to others who are innocent cyber-bystanders.

Now that I have argued for the view that the real-time world is not easily distinguished from the online world, I can move on and discuss one way in which we might justify to ourselves why we let certain things happen on the internet we would not otherwise allow in real-time. We believe there is some spatial distance between us and other cyber-travelers. If no such distance justifies treating others differently online, then we must cautiously approach the distinction between doing and allowing online. Since the doing/allowing relationship resembles the connection between active/passive deception, we also have reason to give up such a distinction online.

## 4 Doing vs. Allowing

It seems intuitively obvious that doing some harmful act is morally worse than allowing some harmful act to occur. No matter how clear our intuitions may seem to be, controversy over the question, “Is doing harm morally worse than merely permitting harm to occur?” still remains. Suppose that a doctor may save six innocent people only if he takes the vital organs, such as the heart, lungs, liver, and kidneys, from a healthy prisoner. To do so the doctor will have to end the prisoner’s life. This seems like a morally reprehensible act.

For a second case, instead of the doctor having to terminate the pris-

oner's life, suppose that the prisoner is allowed to die. The doctor refrains from putting the prisoner on life support system. After the prisoner's brain ceases functioning he performs the surgery to save six people's lives. Our intuitions tell us that the second scenario seems to be less morally problematic than the first one.

Permitting a person to die in order to save six people seems less problematic than deliberately ending a person's life in order to save six people. There are numerous ways to account for the distinction between doing and allowing (cf Bennett, 1966, 1993; Dinello, 1971; Lichtenberg, 1982; Steinbock and Norcross, 1994), of which the most sophisticated analyses is Warren Quinn (1989).

Quinn has argued that the distinction between doing and allowing depends on whether an agent's most direct contribution is an action or an inaction.

Harmful positive agency is that in which an agent's most direct contribution to the harm is an action, whether his own or that of some object. Harmful negative agency is that in which the most direct contribution is an inaction, a failure to prevent the harm. (Quinn, 1989, 301f)

Quinn explains that an agent's most direct contribution to a:

harmful upshot of his agency is the contribution that most directly explains the harm. And one contribution explains harm more directly than another if the explanatory value of the second is exhausted in the way it explains the first. (Quinn, 1989, 301)

The primary difference is between cases where an agent produces the harm by an action and cases in which an agent produces the harm by inaction. Ending the prisoner's life is an action, and refraining from putting the prisoner on life support is an inaction. Since ending the prisoner's life is an action and is the contribution that most directly explains the harm, the agent, on Quinn's view, is positively relevant to a harmful upshot. Therefore, ending the prisoner's life is worse than refraining from putting the prisoner on life support.

Quinn offers a range of cases to show that the distinction between doing and allowing depends on the agent's most direct contribution being

an action. In Rescue I, an agent may choose to save five people in danger of drowning or the agent may choose to save one person. The agent cannot save all six. If the agent decides to save the five people at the expense of the one person, the agent's most direct contribution to the one person's death is an *inaction*. Since the agent has not acted in a way that contributes directly to the one person's death, the agent has not "done" anything to harm the person.

In Rescue II, however, an agent can only save five people by driving over and—presumably—killing one person. The agent's saving five people contributes directly to the death of the one person. So, the agent's most direct contribution is an action, and in this case a "doing."

Quinn contends that there are special cases where an agent may be positively relevant to a harmful upshot, even though the agent's most direct contribution may be characterized as inaction. His example, called "Rescue III," is:

We are off by special train to save five people who are in imminent danger of death. Every second counts. You have just taken over from the driver, who has left the locomotive to attend to something. Since the train is on automatic control you need do nothing to keep it going. But you can stop it by putting on the brakes. You suddenly see someone trapped ahead on the track. Unless you act he will be killed. But if you do stop, and then free the man, the rescue mission will be aborted. So you let the train continue. (Quinn, 1989, 298)

The person has not done anything to "let the train continue." Because on Quinn's account the train may act as the person's agent, the person killed the man on the tracks.

Despite that the agent of Rescue III has done nothing, it is the agent's failure to act that makes us believe the person is morally culpable. Quinn argues:

In this case it seems to me that you make the wrong choice. You must stop the train. It might seem at first that this is because you occupy, if only temporarily, the role of the driver and have therefore assumed a driver's special responsibility to drive the train safely. But, upon reflection, it would not make

much moral difference whether you were actually driving the train or merely had access to its brake. Nor would it much matter whether you were in the train or had happened upon a trackside braking device. The important thing from the standpoint of your agency is that you *can* stop the train and thereby prevent it from killing the one. (Quinn, 1989, 299)

Because the agent of Rescue III failed to stop the train, though the agent clearly could have done so, Quinn believes that the person has acted wrongly. Not acting, at least in this instance, yields harm to innocent bystander and therefore ought to be seen as moral wrongdoing.

To get as clear as possible about the role inaction and action play in the distinction between doing and allowing, Quinn presents one final version of the Rescue scenario, "Rescue IV." He writes:

Suppose. . . you are on a train on which there has just been an explosion. You can stop the train, but that is a complicated business that would take time. So you set it on automatic forward and rush back to the five badly wounded passengers. While attending to them, you learn that a man is trapped far ahead on the track. You must decide whether to return to the cabin to save him or stay with the passengers and save them. (Quinn, 1989, 299)

In Rescue IV, as in Rescue III, your inaction will result in the deaths of those people trapped on the track. The difference between Rescue III and Rescue IV, however, is that in Rescue III the agent intends an action of the train which brings about the trapped person's death. In Rescue IV, no such intention is present. The agent intends no action of the train that ultimately leads to the trapped person's untimely demise.<sup>7</sup> Ultimately:

---

<sup>7</sup>In an interesting variant of Rescue IV relegated to a footnote, Quinn admits that if the agent is "the driver or his designated replacement" who has a special responsibility to see to it that the train not run over and kill people trapped on the tracks, then the agent intends for the action of the train to run over the trapped person on the tracks. Interestingly, Quinn rules out this option because "[i]t does not seem to derive from any supposition that, if you stay with the passengers, you will really be taking the train forward or will somehow be party to the fatal action of the train itself" (Quinn, 1989, 299fn24).

In Rescue III, but not in Rescue IV, the train kills the man *because* of your intention that it continue forward. This implicates you, I believe, in the fatal action of the train itself. If you had no control, but merely wished that the rescue would continue—or if, as in Rescue IV, you had control but no such wish—you would not be party to the action of the train. But the combination of control and intention in Rescue III makes for a certain kind of complicity. Your choice to let the train continue forward is strategic and deliberate. Since you clearly *would* have it continue for the sake of the five, there is a sense in which, by deliberately not stopping it, you *do* have it continue. For these reasons your agency counts as positive. (Quinn, 1989, 300)

According to Quinn, the moral implication of distinguishing doing and allowing in this way depends on the distinction between negative and positive rights, a view he inherits from Philippa Foot (1967). When an agent's action or inaction is positively relevant to some harmful upshot, the agent has violated someone's negative rights. Positive rights have been violated when an agent's action or inaction is negatively relevant to some harmful consequence. Quinn holds that negative rights must be more strictly enforced than positive rights. So, a violation of an agent's negative rights is worse than a violation of an agent's positive rights. A critical aspect of Quinn's view is that the man on the railroad track has a say about what harm may befall him. If he does not have a say, then an agent's negative rights, something we must uphold more stringently, has been abandoned.

Quinn's primary objective is "to find the formulation of the [doing / allowing] distinction that best fits our moral intuitions" (Quinn, 1989, 288). He believed the distinction is of the *greatest general importance* because "it enters as a strand into many real moral issues and because it stands in apparent opposition to the most general of all moral theories, consequentialism" (Quinn, 1989, 288).

Quinn's four "Rescue" cases seem to speak directly to the concerns raised against forms of online deception. Because netizens are unable to discern whether a person is real or fake and because we do take what others say to us and about us very seriously, we ought to err on the side of caution and believe that the harmful upshot of online actions or inactions

are morally wrong. Not alerting other users of a fake social networking profile is undertaken intentionally. The likelihood of harm befalling a fake member's "friends" is increased because the anonymity provided by the profile allows the user to say potentially very harmful comments to unsuspecting members.

When netizens deceive other netizens, tragic results could follow. Even if the netizen "did not mean" for the comments to hasten the death of another person, the harm was allowed to be brought about by the netizen's actions. Just in virtue of that fact, we have to judge the netizen's actions harshly. Clearly, the person bringing about the harmful upshot has done something morally wrong.

This section has outlined a distinction between doing and allowing, according to one of the debates main contributors, Warren Quinn. Then, I argued that if netizens allow harmful acts to be brought about by their inaction, i.e., their not alerting others to the deception, we have to say that they have done something morally wrong.

## 5 Online deception is harmful

There are some forms of deception, at least in real-time, we cannot do without. The soldier's camouflage is one example. Deception often encountered in virtual space of social networking fails to protect others in any way. In fact, it seems that online deception is prone to harm others, regardless of whether the deceiving agent meant or intended to do so. In this closing section, I will argue against some potential counterarguments one might provide in response to my argument.

First, a critic might contend that creating fake social networking profiles is morally permissible because it allows the author of the fake profile to express herself openly. This kind of argument presumes that free speech is protected, and, if the person feels that her rights might be violated or she could lose her job by openly criticizing her employer online, she should have the ability to do it pseudonymously. Free speech is protected, but free speech that harms others certainly is not. Just as we would not protect harmful speech acts like yelling "Fire!" in a crowded theater, we should not think harmful conversations bringing about a person's death is something that ought to be protected either.

Second, one might believe that since we can think of the virtual world



differently than we think of the real-time world, any event resulting from online discussions or comments is coincidental. This is a common excuse bullies use if one of their victims harms themselves. According to this interpretation, the words did not hurt or bring about the harm to the victim. After all, correlation is not causation. This argument has no merit. To believe that Megan Meier's suicide was not caused by her neighbors creating a false identity and berating her online is to overlook the fact that Megan and "Josh," her supposed online friend, had sown a friendship. "Josh's" terrible comments led Megan to commit suicide. Her death was no coincidence.

Finally, one could defend the use of online deception to protect oneself from real-time harm befalling oneself or others. If a person creates a fake Facebook profile and befriends his spouse to uncover the spouse's infidelity, then we might think that knowing the truth about the spouse outweighs the deception involved to discover it. Although this seems like a strong argument for the moral permissibility of online deception, the likelihood that it will lead to morally harmful actions is greater than the good that might be produced in learning the truth about one's spouse.

## 6 Conclusion

There are a variety of examples of real-time and online deception. No matter how innocuous online deception might seem and that the deceiving person intended it otherwise or did not mean for some consequence to follow, it is always morally wrong to deceive someone online. In this paper, I have argued that deceiving in real-time is morally distinguishable from deceiving online because online actions are not as morally innocuous as actions occurring in real-time. Our failure to detect the fine-grained characteristics of another virtual netizen leads us to believe that the person intended to do a moral harm. Openly deceiving someone on Facebook or Twitter is not a way to build virtual friendships but to destroy them and, at least sometimes, to end the life of innocent virtual bystanders.

## References

- Bennett, J. (1966). Whatever the consequences. *Analysis*, 26(23).
- Bennett, J. (1993). Negation and abstention: Two theories of allowing. *Ethics*, 104:75–96.
- Dinello, D. (1971). On killing and letting die. *Analysis*, 31:84–86.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5:5–15.
- Kagan, S. (1998). *Normative Ethics*. Westview Press, Boulder, CO.
- Lichtenberg, J. (1982). The moral equivalence of action and omission. *Canadian Journal of Philosophy*, 8:19–36.
- Quinn, W. (1989). Actions, intentions, and consequences: The doctrine of doing and allowing. *Philosophical Review*, 98:287–312.
- Steinbock, B. and Norcross, A., editors (1994). *Killing and Letting Die*. Fordham University Press, New York, second edition.