

The Unacceptability of Asimov's Three Laws of Robotics as a Basis for Machine Ethics

Susan Leigh Anderson

ONCE PEOPLE UNDERSTAND THAT MACHINE ETHICS IS CONCERNED WITH how intelligent machines should behave, they often maintain that Isaac Asimov has already given us an ideal set of rules for such machines. They have in mind Asimov's Three Laws of Robotics:

1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law. (Asimov 1976)

I shall argue that in "The Bicentennial Man" (Asimov 1976), Asimov rejected his own Three Laws as a proper basis for Machine Ethics. He believed that a robot with the characteristics possessed by Andrew, the robot hero of the story, should not be required to be a slave to human beings as the Three Laws dictate. He further provided an explanation for why humans feel the need to treat intelligent robots as slaves, an explanation that shows a weakness in human beings that makes it difficult for them to be ethical paragons. Because of this weakness, it seems likely that machines like Andrew could be more ethical than most human beings. "The Bicentennial Man" gives us hope that intelligent machines can not only be taught to behave in an ethical fashion, but they might be able to lead human beings to behave more ethically as well.

To be more specific, I shall use "The Bicentennial Man" to argue for the following: (1) An intelligent robot like Andrew satisfies most, if not all, of the requirements philosophers have proposed for a being/entity to have moral standing/rights, making the Three Laws immoral. (2) Even if the machines that are actually developed fall short of being like Andrew and should probably not be considered to have moral standing/rights, it is still problematic for humans to program them to follow the Three Laws of Robotics. From (1) and (2), we can conclude that (3) whatever the status of the machines that are developed,

Asimov's Three Laws of Robotics would be an unsatisfactory basis for Machine Ethics. That the status of intelligent machines doesn't matter is important because (4) in real life, it would be difficult to determine the status of intelligent robots. Furthermore, (5) because intelligent machines can be designed to consistently follow moral principles, they have an advantage over human beings in having the potential to be ideal ethical agents, because human beings' actions are often driven by irrational emotions.

“The Bicentennial Man”

Isaac Asimov's “The Bicentennial Man” was originally commissioned to be part of a volume of stories written by well-known authors to commemorate the United States' bicentennial.¹ Although the project didn't come to fruition, Asimov ended up with a particularly powerful work of philosophical science fiction as a result of the challenge he had been given. It is important that we know the background for writing the story because “The Bicentennial Man” is simultaneously a story about the history of the United States and a vehicle for Asimov to present his view of how intelligent robots should be treated and be required to act.

“The Bicentennial Man” begins with the Three Laws of Robotics. The story that follows is told from the point of view of Andrew, an early, experimental robot – intended to be a servant in the Martin household – who is programmed to obey the Three Laws. Andrew is given his human name by the youngest daughter in the family, Little Miss, for whom he carves a beautiful pendant out of wood. This leads to the realization that Andrew has unique talents, which the Martins encourage him to develop by giving him books to read on furniture design.

Little Miss, his champion during her lifetime, helps Andrew to fight first for his right to receive money from his creations and then for the freedom he desires. A judge does finally grant Andrew his freedom, despite the opposing attorney's argument that “The word *freedom* has no meaning when applied to a robot. Only a human being can be free.” In his decision, the judge maintains, “There is no right to deny freedom to any object with a mind advanced enough to grasp the concept and desire the state.”

Andrew continues to live on the Martin's property in a small house built for him, still following the Three Laws despite having been granted his freedom. He begins wearing clothes so that he will not be so different from human beings, and later he has his body replaced with an android one for the same reason. Andrew wants to be accepted as a human being.

In one particularly powerful incident, shortly after he begins wearing clothes, Andrew encounters some human bullies while on his way to the library. They order him to take off his clothes and then dismantle himself. He must obey the humans because of the Second Law, and he cannot defend himself without

¹ Related to me in conversation with Isaac Asimov.

harming the bullies, which would be a violation of the First Law. He is saved just in time by Little Miss's son, who informs him that humans have an irrational fear of an intelligent, unpredictable, autonomous robot that can exist longer than a human being – even one programmed with the Three Laws – and that is why they want to destroy him.

In a last ditch attempt to be accepted as a human being, Andrew arranges that his “positronic” brain slowly cease to function, just like a human brain. He maintains that it does not violate the Third Law, because his “aspirations and desires” are more important to his life than “the death of his body.” This last sacrifice, “accept[ing] even death to be human,” finally allows him to be accepted as a human being. He dies two hundred years after he was made and is declared to be “the Bicentennial Man.” In his last words, whispering the name “Little Miss,” Andrew acknowledges the one human being who accepted and appreciated him from the beginning.

Clearly, the story is meant to remind Americans of their history, that particular groups, especially African Americans, have had to fight for their freedom and to be fully accepted by other human beings.² It was wrong that African Americans were forced to act as slaves for white persons, and they suffered many indignities, and worse, that were comparable to what the bullies inflicted upon Andrew. As there was an irrational fear of robots in the society in which Andrew functioned, there were irrational beliefs about blacks among whites in earlier stages of our history, which led to their mistreatment. Unfortunately, contrary to Aristotle's claim that “man is the rational animal,” human beings are prone to behaving in an irrational fashion when they believe that their interests are threatened, especially by beings/entities they perceive as being different from themselves.

In the history of the United States, gradually more and more beings have been granted the same rights that others possessed, and we've become a more ethical society as a result. Ethicists are currently struggling with the question of whether at least some higher-order animals should have rights, and the status of human fetuses has been debated as well. On the horizon looms the question of whether intelligent machines should have moral standing.

Asimov has made an excellent case for the view that certain types of intelligent machines, ones like Andrew, should be given rights and should not be required to act as slaves for humans. By the end of the story, we see how wrong it is that Andrew has been forced to follow the Three Laws. Yet we are still left with something positive, on reflection, about Andrew's having been programmed to follow moral principles. They may not have been the *correct* principles, because they did not acknowledge rights Andrew should have had, but Andrew was a far more moral entity than most of the human beings he encountered. Most of the human beings in “The Bicentennial Man” were prone to being carried away by irrational

² One of the characters in “The Bicentennial Man” remarks, “There have been times in history when segments of the human population fought for full human rights.”

emotions, particularly irrational fears, so they did not behave as rationally as Andrew did. If we can just find the *right* set of ethical principles for intelligent machines to follow, they could very well show human beings how to behave more ethically.

Characteristic(s) Necessary to Have Moral Standing

It is clear that most human beings are “speciesists.” As Peter Singer defines the term, “Speciesism . . . is a prejudice or attitude of bias toward the interests of members of one’s own species and against those members of other species” (Singer 1975). Speciesism can justify “the sacrifice of the most important interests of members of other species in order to promote the most trivial interests of our own species” (Singer 1975). For a speciesist, only members of one’s own species need to be taken into account when deciding how to act. Singer was discussing the question of whether animals should have moral standing, that is, whether they should count in calculating what is right in an ethical dilemma that affects them; but the term can be applied when considering the moral status of intelligent machines if we allow an extension of the term “species” to include a machine category as well. The question that needs to be answered is whether we are justified in being speciesists.

Philosophers have considered several possible characteristics that it might be thought a being/entity must possess in order to have moral standing, which means that an ethical theory must take interests of the being/entity into account. I shall consider a number of these possible characteristics and argue that most, if not all, of them would justify granting moral standing to the fictional robot Andrew (and, very likely, higher-order animals as well), from which it follows that we are not justified in being speciesists. However, it will be difficult to establish, in the real world, whether intelligent machines/robots possess the characteristics that Andrew does.

In the late eighteenth century, the utilitarian Jeremy Bentham considered whether *possessing the faculty of reason or the capacity to communicate* is essential in order for a being’s interests to be taken into account in calculating which action is likely to bring about the best consequences:

What . . . should [draw] the insuperable line? Is it the faculty of reason, or perhaps the faculty of discourse? But a full-grown horse or dog is beyond comparison a more rational, as well as a more conversable animal, than an infant of a day or even a month old. But suppose they were otherwise, what would it avail? The question is not, Can they reason? nor Can they talk? but Can they suffer? (Bentham 1799)

In this famous passage, Bentham rejected the ability to reason and communicate as being essential to having moral standing (tests that Andrew would have passed with flying colors), in part because they would not allow newborn humans to have moral standing. Instead, Bentham maintained that *sentience* (he focused, in

particular, on the ability to suffer, but he intended that this should include the ability to experience pleasure as well) is what is critical. Contemporary utilitarian Peter Singer agrees. He says, "If a being suffers there can be no moral justification for refusing to take that suffering into consideration" (Singer 1975).

How would Andrew fare if sentience were the criterion for having moral standing? Was Andrew capable of experiencing enjoyment and suffering? Asimov manages to convince us that he was, although a bit of a stretch is involved in the case he makes for each. For instance, Andrew says of his woodworking creations:

"I enjoy doing them, Sir," Andrew admitted.

"Enjoy?"

"It makes the circuits of my brain somehow flow more easily. I have heard you use the word *enjoy* and the way you use it fits the way I feel. I enjoy doing them, Sir."

To convince us that Andrew was capable of suffering, here is how Asimov described the way Andrew interacts with the judge as he fights for his freedom:

It was the first time Andrew had spoken in court, and the judge seemed astonished for a moment at the human timbre of his voice.

"Why do you want to be free, Andrew? In what way will this matter to you?"

"Would *you* wish to be a slave, Your Honor," Andrew asked.

In the scene with the bullies, when Andrew realizes that he cannot protect himself, Asimov writes, "At that thought, he felt every motile unit contract slightly and he quivered as he lay there."

Admittedly, it would be very difficult to determine whether a robot has feelings, but as Little Miss points out, it is difficult to determine whether even another human being has feelings like oneself. All we can do is use behavioral cues:

"Dad . . . I don't know what [Andrew] feels inside, but I don't know what *you* feel inside either. When you talk to him you'll find he reacts to the various abstractions as you and I do, and what else counts? If someone else's reactions are like your own, what more can you ask for?"

Another philosopher, Immanuel Kant, maintained that only beings that are *self-conscious* should have moral standing (Kant 1780). At the time that he expressed this view, it was believed that all and only human beings are self-conscious. It is now recognized that very young children lack self-consciousness and that higher-order animals (e.g., monkeys and great apes³) possess this quality, so putting emphasis on this characteristic would no longer justify our speciesism.⁴

³ In a well-known video titled "Monkey in the Mirror," a monkey soon realizes that the monkey it sees in a mirror is itself, and it begins to enjoy making faces, etc., watching its own reflection.

⁴ Christopher Grau has pointed out that Kant probably had a more robust notion of self-consciousness in mind that includes autonomy and "allows one to discern the moral law through the Categorical Imperative." Still, even if this rules out monkeys and great apes, it also rules out very young human beings.

Asimov managed to convince us early on in “The Bicentennial Man” that Andrew is self-conscious. On the second page of the story, Andrew asks a robot surgeon to perform an operation on him to make him more like a man:

“Now, upon whom am I to perform this operation?”

“Upon me,” Andrew said.

“But that is impossible. It is patently a damaging operation.”

“That does not matter,” Andrew said calmly.

“I must not inflict damage,” said the surgeon.

“On a human being, you must not,” said Andrew, “but I, too, am a robot.”

In real life, because humans are highly skeptical, it would be difficult to establish that a robot is self-conscious. Certainly a robot could talk about itself in such a way, like Andrew did, that might *sound* like it is self-conscious, but to prove that it really *understands* what it is saying and that it has not just been “programmed” to say these things is another matter.

In the twentieth century, the idea that a being does or does not have *rights* became a popular way of discussing the issue of whether a being/entity has moral standing. Using this language, Michael Tooley essentially argued that *to have a right to something, one must be capable of desiring it*. More precisely, he said that “an entity cannot have a particular right, *R*, unless it is at least capable of having some interest, *I*, which is furthered by its having right *R*” (Tooley 1972). As an example, he said that a being cannot have a right to life unless it is capable of desiring its continued existence.

Andrew desires his freedom. He says to a judge: “It has been said in this courtroom that only a human being can be free. It seems to me that only someone who *wishes* for freedom can be free. I wish for freedom.” Asimov continues by writing that “it was this statement that cued the judge.” He was obviously “cued” by the same criterion Tooley gave for having a right, for he went on to rule that “[t]here is no right to deny freedom to any object advanced enough to grasp the concept and desire the state.”

Yet once again, if we were to talk about real life instead of a story, we would have to establish that Andrew truly *grasped the concept* of freedom and *desired* it. It would not be easy to convince a skeptic. No matter how much appropriate behavior a robot exhibited, including uttering certain statements, there would be those who would claim that the robot had simply been “programmed” to do and say certain things.

Also in the twentieth century, Tibor Machan maintained that to have rights it was necessary to be a *moral agent*, where a moral agent is one who is expected to behave morally. He then went on to argue that because only human beings possess this characteristic, we are justified in being speciesists:

[H]uman beings are indeed members of a discernibly different species – the members of which have a moral life to aspire to and must have principles upheld for them in

communities that make their aspiration possible. Now there is plainly no valid intellectual place for rights in the non-human world, the world in which moral responsibility is for all practical purposes absent. (Machan 1991)

Machan's criterion for when it would be appropriate to say that a being/entity has rights – that it must be a “moral agent” – might seem to be not only reasonable,⁵ but helpful for the Machine Ethics enterprise. Only a being that can respect the rights of others should have rights itself. So, if we could succeed in teaching a machine how to be moral (that is, to respect the rights of others), then it should be granted rights itself.

Yet we've moved too quickly here. Even if Machan were correct, we would still have a problem that is similar to the problem of establishing that a machine has feelings, is self-conscious, or is capable of desiring a right. Just because a machine's behavior is guided by moral principles doesn't mean that it is a moral agent, that is, that we would ascribe moral responsibility to the machine. To ascribe moral responsibility would require that the agent intended the action and, in some sense, could have done otherwise (Anderson 1995),⁶ both of which are difficult to establish.

If Andrew (or any intelligent machine) followed ethical principles only because he was programmed that way, as were the later, predictable robots in “The Bicentennial Man,” then we would not be inclined to hold him morally responsible for his actions. However, Andrew found creative ways to follow The Three Laws, convincing us that he intended to act as he did and that he could have done otherwise. An example has been given already: when he chose the death of his body over the death of his aspirations to satisfy the Third Law.

Finally, Mary Anne Warren combined the characteristics that others have argued for as requirements for a being to be “a member of the moral community” with one more – *emotionality*. She claimed that it is “persons” that matter, that is, are members of the moral community, and this class of beings is not identical with the class of human beings: “[G]enetic humanity is neither necessary nor

⁵ In fact, however, it is problematic. Some would argue that Machan has set the bar too high. Two reasons could be given: (1) A number of humans (most noticeably very young children) would, according to his criterion, not have rights because they can't be expected to behave morally. (2) Machan has confused “having rights” with “having duties.” It is reasonable to say that in order to *have duties* to others, you must be capable of behaving morally, that is, of respecting the rights of others, but to *have rights* requires something less than this. That's why young children can have rights, but not duties. In any case, Machan's criterion would not justify our being speciesists because recent evidence concerning the great apes shows that they are capable of behaving morally. I have in mind Koko, the gorilla that has been raised by humans (at the Gorilla Foundation in Woodside, California) and absorbed their ethical principles as well as having been taught sign language.

⁶ I say “in some sense, could have done otherwise” because philosophers have analyzed “could have done otherwise” in different ways, some compatible with Determinism and some not; but it is generally accepted that freedom in some sense is required for moral responsibility.

sufficient for personhood. Some genetically human entities are not persons, and there may be persons who belong to other species” (Warren 1997). She listed six characteristics that she believes define personhood:

1. *Sentience* – the capacity to have conscious experiences, usually including the capacity to experience pain and pleasure;
2. *Emotionality* – the capacity to feel happy, sad, angry, angry, loving, etc.;
3. *Reason* – the capacity to solve new and relatively complex problems;
4. *The capacity to communicate*, by whatever means, messages of an indefinite variety of types; that is, not just with an indefinite number of possible contents, but on indefinitely many possible topics;
5. *Self-awareness* – having a concept of oneself, as an individual and/or as a member of a social group; and finally
6. *Moral agency* – the capacity to regulate one’s own actions through moral principles or ideals. (Warren 1997)

It is interesting and somewhat surprising that Warren added the characteristic of *emotionality* to the list of characteristics that others have mentioned as being essential to personhood, because she was trying to make a distinction between persons and humans and argue that it is the first category that comprises the members of the moral community. *Humans* are characterized by emotionality, but some might argue that this is a weakness of theirs that can interfere with their ability to be members of the moral community, that is, their ability to respect the rights of others.

There is a tension in the relationship between emotionality and being capable of acting morally. On the one hand, one has to be sensitive to the suffering of others to act morally. This, for human beings,⁷ means that one must have empathy, which in turn requires that one has experienced similar emotions oneself. On the other hand, as we’ve seen, the emotions of human beings can easily get in the way of acting morally. One can get so “carried away” by one’s emotions that one becomes incapable of following moral principles. Thus, for humans, finding the correct balance between the subjectivity of emotion and the objectivity required to follow moral principles seems to be essential to being a person who consistently acts in a morally correct fashion.

In any case, although Andrew exhibited little “emotionality” in “The Bicentennial Man,” and Asimov seemed to favor Andrew’s way of thinking in ethical matters to the “emotional antipathy” exhibited by the majority of humans, there is one time when Andrew clearly does exhibit emotionality. It comes at the very end of the story, when he utters the words “Little Miss” as he dies. Notice, however, that this coincided with his being declared a *man*, meaning a

⁷ I see no reason, however, why a robot/machine can’t be trained to take into account the suffering of others in calculating how it will act in an ethical dilemma, without its having to be emotional itself.

human being. As the director of research at U.S. Robots and Mechanical Men Corporation in the story says about Andrew's desire to be a man: "That's a puny ambition, Andrew. You're better than a man. You've gone downhill from the moment you opted to become organic." I suggest that one way in which Andrew had been better than most human beings was that he did not get carried away by "emotional antipathy."

I'm not convinced, therefore, that one should put much weight on emotionality as a criterion for a being's/entity's having moral standing, because it can often be a liability to determining the morally correct action. If it is thought to be essential, it will, like all the other characteristics that have been mentioned, be difficult to establish. Behavior associated with emotionality can be mimicked, but that doesn't necessarily guarantee that a machine truly has feelings.

Why the Three Laws Are Unsatisfactory Even If Machines Don't Have Moral Standing

I have argued that it may be very difficult to establish, with any of the criteria philosophers have given, that a robot/machine that is created possesses the characteristic(s) necessary to have moral standing/rights. Let us assume, then, just for the sake of argument, that the robots/machines that are created should not have moral standing. Would it follow, from this assumption, that it would be acceptable for humans to build into the robot Asimov's Three Laws, which allow humans to harm it?

Immanuel Kant considered a parallel situation and argued that humans should not harm the entity in question, even though it lacked rights itself. In "Our Duties to Animals," from his *Lectures on Ethics* (Kant 1780) Kant argued that even though animals don't have moral standing and can be used to serve the ends of human beings, we should still not mistreat them because "[t]ender feelings towards dumb animals develop humane feelings towards mankind." He said that "he who is cruel to animals becomes hard also in his dealings with men." So, even though we have no *direct* duties to animals, we have obligations toward them as "indirect duties towards humanity."

Consider, then, the reaction Kant most likely would have had to the scene involving the bullies and Andrew. He would have abhorred the way they treated Andrew, fearing that it could lead to the bullies treating human beings badly at some future time. Indeed, when Little Miss's son happens on the scene, the bullies' bad treatment of Andrew is followed by offensive treatment of a human being – they say to his human rescuer, "What are you going to do, pudgy?"

It was the fact that Andrew had been programmed according to the Three Laws that allowed the bullies to mistreat him, which in turn could (and did) lead to the mistreatment of human beings. One of the bullies asks, "who's

to object to anything we do” before he gets the idea of destroying Andrew. Asimov then writes:

“We can take him apart. Ever take a robot apart?”

“Will he let us?”

“How can he stop us?”

There was no way Andrew could stop them, if they ordered him in a forceful enough manner not to resist. The Second Law of obedience took precedence over the Third Law of self-preservation. In any case, he could not defend himself without possibly hurting them, and that would mean breaking the First Law.

It is likely, then, that Kant would have condemned the Three Laws, even if the entity that was programmed to follow them (in this case, Andrew) did not have moral standing itself. The lesson to be learned from his argument is this: Any ethical laws that humans create must advocate the respectful treatment of even those beings/entities that lack moral standing themselves if there is any chance that humans’ behavior toward other humans might be adversely affected otherwise.⁸ If humans are required to treat other entities respectfully, then they are more likely to treat each other respectfully.

An unstated assumption of Kant’s argument for treating certain beings well, even though they lack moral standing themselves, is that the beings he refers to are similar in a significant respect to human beings. They may be similar in appearance or in the way they function. Kant, for instance, compared a faithful dog with a human being who has served someone well:

[I]f a dog has served his master long and faithfully, his service, on the analogy of human service, deserves reward, and when the dog has grown too old to serve, his master ought to keep him until he dies. Such action helps to support us in our duties towards human beings. (Kant 1780)

As applied to the Machine Ethics project, Kant’s argument becomes stronger, the more the robot/machine that is created resembles a human being in its functioning and/or appearance. The more the machine resembles a human being, the more moral consideration it should receive. To force an entity like Andrew – who resembled human beings in the way he functioned *and* in his appearance – to follow the Three Laws, which permitted humans to harm him, makes it likely that having such laws will lead to humans harming other humans as well.

Because a goal of AI is to create entities that can duplicate intelligent human behavior, if not necessarily their form, it is likely that the autonomous ethical machines that may be created – even if they are not as humanlike as Andrew – will resemble humans to a significant degree. It, therefore, becomes all the more

⁸ It is important to emphasize here that I am not necessarily agreeing with Kant that robots like Andrew, and animals, should not have moral standing/rights. I am just making the hypothetical claim that *if* we determine that they should not, there is still a good reason, because of indirect duties to human beings, to treat them respectfully.

important that the ethical principles that govern their behavior should not permit us to treat them badly.

It may appear that we could draw the following conclusion from the Kantian argument given in this section: An autonomous moral machine must be treated as if it had the same moral standing as a human being. However, this conclusion reads more into Kant's argument than one should.

Kant maintained that beings, like the dog in his example, that are sufficiently like human beings so that we must be careful how we treat them to avoid the possibility that we might go on to treat human beings badly as well, should not have the same moral status as human beings. He says, "[a]nimals . . . are there merely as a means to an end. That end is man" (Kant 1780). Contrast this with his famous second imperative that should govern our treatment of human beings:

Act in such a way that you always treat humanity, whether in your own person or in the person of any other, never simply as a means, but always at the same time as an end. (Kant 1785)

Thus, according to Kant, we are entitled to treat animals, and presumably intelligent ethical machines that we decide should not have the moral status of human beings, differently from human beings. We can require them to do things to serve our ends, but we should not mistreat them. Because Asimov's Three Laws permit humans to mistreat robots/intelligent machines, they are not, according to Kant, satisfactory as moral principles that these machines should be forced to follow.

In conclusion, using Asimov's "Bicentennial Man" as a springboard for discussion, I have argued that Asimov's Three Laws of Robotics are an unsatisfactory basis for Machine Ethics, regardless of the status of the machine. I have also argued that this is important because it would be very difficult, in practice, to determine the status of an intelligent, autonomous machine/robot. Finally, I have argued that Asimov demonstrated that such a machine/robot programmed to follow ethical principles is more likely to consistently behave in an ethical fashion than the majority of humans.

References

- Anderson, S. (1995), "Being Morally Responsible for an Action Versus Acting Responsibly or Irresponsibly," *Journal of Philosophical Research*, Volume XX, pp. 451–462.
- Asimov, I. (1976), "The Bicentennial Man," in *Philosophy and Science Fiction* (Philips, M., ed.), pp. 183–216, Prometheus Books, Buffalo, NY, 1984.
- Bentham, J. (1799), *An Introduction to the Principles of Morals and Legislation*, chapter 17 (Burns, J. and Hart, H., eds.), Clarendon Press, Oxford, 1969.
- Kant, I. (1780), "Our Duties to Animals," in *Lectures on Ethics* (Infield, L., trans.), Harper & Row, New York, NY, 1963, pp. 239–241.
- Kant, I. (1785), *The Groundwork of the Metaphysic of Morals* (Paton, H. J., trans.), Barnes and Noble, New York, 1948.

- Machan, T. (1991), "Do Animals Have Rights?" *Public Affairs Quarterly*, Vol. 5, no. 2, pp. 163–173.
- Singer, P. (1975), "All Animals are Equal," in *Animal Liberation: A New Ethics for our Treatment of Animals*, Random House, New York, pp. 1–22.
- Tooley, M. (1972), "Abortion and Infanticide," *Philosophy and Public Affairs*, no. 2, pp. 47–66.
- Warren, M. (1997), "On the Moral and Legal Status of Abortion," in *Ethics in Practice* (La Follette, H., ed.), Blackwell, Oxford.