

DOES 'OUGHT' STILL IMPLY 'CAN'?

Clayton Littlejohn

Forthcoming in *Philosophia*.

Graham (2011) offers a novel argument against the thesis that 'ought' implies 'can':

OIC: Necessarily, if you are obligated to ϕ , you can ϕ .

While his argument has considerable intuitive force, it ought to be rejected.

According to Graham, it is possible for an agent to have an obligation to ϕ even if she does not have the opportunity and ability to ϕ . While this example poses no threat to OIC, it is an important example for the purposes of our discussion:

Transplant: A surgeon has ten patients, each of whom will die of organ failure if he does not receive an organ transplant. The surgeon wants to save her patients and is convinced by philosophical arguments to the effect that it would be morally permissible to kill two people in order to save them. She notices that in another room of the hospital there are two innocent and unconscious tonsillectomy patients who are perfect organ matches for her patients. The only means by which the hospital janitor, who is aware of the situation, can stop the surgeon from chopping up the two and redistributing their organs among the ten is by shooting her with his pistol. He does so and thereby kills her (2011: 345).

Graham says two things concerning **Transplant**:

1. It is morally permissible for the janitor to kill the surgeon.
2. If the janitor had not killed the surgeon, the surgeon would have acted impermissibly in killing the two patients.

Both of these claims strike me as being intuitively correct. Moreover, I think Graham might be right that (1) is true in **Transplant** because (2) is.

Consider a variant of **Transplant**:

Compulsion: Everything is as it is in **Transplant** except that the surgeon cannot refrain from killing the two because the ten are her grandchildren, and she is as compelled to save them as is the most severe kleptomaniac to steal (2011: 346).

Concerning **COMPULSION**, Graham says this:

As in **Transplant**, in ... [**Compulsion**] (1) is true. But not only does the addition of the surgeon's compulsion not change this fact, it also seems irrelevant to it. In other words, whatever explains the truth of (1) in **Transplant**, it seems, must also explain the truth of (1) in **Compulsion**. But if this is right, then OIC must be false. For if, as I shall argue, what explains the truth of (1) in **Transplant** is (2), and what makes (1) true in **Compulsion** is the same as that which makes (1) true in **Transplant**, then it must be the case that (2) is true in **Compulsion** (2011: 346).

In light of these remarks, Graham offers us this argument against OIC:

- P1. We can explain why (1) is true in **Transplant** in terms of (2).
- P2. In **Compulsion**, (1) is also true.
- P3. Intuitively, the salient difference between **Transplant** and **Compulsion** (i.e., that the physician cannot refrain from killing the two patients) is irrelevant to the explanation of (1) in **Compulsion**.

C1. Thus, we can explain (1) in the same way in **Compulsion** and in **Transplant**, by citing (2).

P4. In **Compulsion**, the surgeon cannot refrain from killing the patients.

C2. Thus, it is possible for someone to act impermissibly even if she cannot refrain from acting in that way.

If this argument is sound, OIC is sunk. The case that Graham builds against OIC is impressive. To defend OIC, someone might argue against (P3) and try to show that there is some important difference between **Compulsion** and **Transplant** that would allow us to say that the permissibility of shooting the surgeon to save the patients has nothing to do with whether or not the surgeon would act impermissibly if the surgeon chopped up the patients and redistributed their organs. To defend OIC, we have to find an explanation as to why (1) is true in **Compulsion** that does not presuppose (2). This is what I shall provide here.

Graham's argument against OIC will not persuade those mad-dog consequentialists who say that the janitor acted impermissibly by preventing the surgeon from redistributing the organs to save the greatest number of patients. They would reject (P1) and (P2) because they would accept (1) and reject (2). The argument might not persuade those mild-mannered consequentialists who say that there is a consequentialist rationale for using force to prevent surgeons from chopping up patients without following hospital protocol. The rationale for using force might not depend upon whether those who would chop up the patients would be acting impermissibly. So, these mild-mannered consequentialists could reject (P1) and (P2) on the grounds that they do not think (1) tells us much about (2). Thus, the mad-dog and mild-mannered consequentialists among us would reject (P1) and (P2) either because they reject (1) or reject the suggestion that (2) explains (1). The success of Graham's argument seems to depend upon some sort of non-consequentialist principle. Should non-consequentialists be persuaded by this argument? I do not think his argument succeeds if we assume some broadly Kantian principles.

I agree with Graham that (P1) is rather plausible. It seems perfectly plausible that (1) is true in **Transplant** because (2) is. I also agree that it is not implausible that there is a common explanation as to why the janitor would be justified in intervening that covers both cases. Even if there is a common explanation of (1) in **Transplant** and **Compulsion**, (C1) might be false. The case for intervention in **Transplant** might be overdetermined. It might be that there is a case for intervention in **Transplant** is not a case for intervention in **Compulsion** even if there is a second case for intervention that is common to both. Suppose, for example, that there is (i) a *pro tanto* reason to stop someone from engaging in wrongdoing if you can do so and (ii) a *pro tanto* reason to protect the patients waiting for their tonsillectomies. If we suppose that OIC is true and that nobody can engage in wrongdoing unless they can refrain, we might say that (i) and (ii) provide a case for intervention in **Transplant** that would be sufficient to justify intervention in the absence of reasons to refrain from intervening. We might say that (ii) provides a case for intervention in **Compulsion** that would be sufficient to justify intervention in the absence of reasons to refrain from intervening even if (i) does not provide a case for intervening in this case. While the salient difference between **Transplant** and **Compulsion** is irrelevant when it comes to explaining why there is some case or other for intervention in these cases, it is not irrelevant to some case for intervention that applies to one of these cases. To show that (C2) is true, Graham has to show that the differences between **Compulsion** and **Transplant** are irrelevant to the totality of the morally relevant features in these cases. I do not think that this is something that he has done.

I do not think that Graham's argument fails just because the case for intervention in **Transplant** is overdetermined. If we want a common explanation as to why (1) is true in both **Transplant** and **Compulsion**, we know that Graham would say that (1) is true in both cases

because (2) is true in both cases. If this is right, OIC is sunk. To save OIC, we should try to find an alternative explanation of (1) consistent that is consistent with (2) being false in **Compulsion**. I think that a broadly Kantian explanation provides the explanation we seek.

Let us first suppose that the janitor has a *pro tanto* duty to protect the patients from someone who would chop them up to take their organs without their consent. Suppose further that the janitor sees the surgeon advance on the patients with scalpel drawn and knew (never mind how) that the surgeon intended to chop up the patients. The janitor would have a *pro tanto* duty to prevent the surgeon from chopping up the patients and so would have a *pro tanto* reason to use the means available to stop the surgeon. From his perspective, it would seem that if he shot the surgeon with the intention of saving the patients, he would have an adequate justification for doing so. Suppose that the surgeon explains why she needs the organs. She explains that she needs the organs because she plans on redistributing the organs so as to save ten lives. If the janitor decides at this point to refrain from saving the two patients and decides not to try to meet her *pro tanto* obligation to protect, she would do so on the grounds that these patients' organs could be used elsewhere to save the greater number. The janitor would decide against saving upon seeing that these organs could be redistributed and upon seeing how these patients could be used as instruments for bringing about a very good outcome. It seems rather un-Kantian to decide against assisting those that one has a *pro tanto* duty to protect because one sees that they can be treated as mere means if one decides against assisting. To decide against intervention on these grounds is to treat the patients merely as means. It is to treat them as things to be saved were it not for the fact that they would be used as instruments for producing a good end. This is true in both **Transplant** and **Compulsion**.

Someone could object that the Kantian principles do not apply to omissions, but this does not seem terribly plausible. If you see that someone is stepping into the street and is about to be flattened by a bus and then see that they are wearing a hat that reads 'Organ Donor' and thereby decide not to pull them to safety as you initially intended, it does seem that you are treating this person as a mere means. You are clearly treating them differently than you would treat people who wear shirts that read 'I'll take my organs with me to the grave'. If you see that a drone is about to launch that has been programmed to fly over a hospital and bomb it to smithereens and think about pushing the stop button but then think better of it because you see that it is destroying the hospital because the programmer has worked out that this will undermine the enemy and further an otherwise just cause, you have shown a willingness to make decisions about whether to save on the basis of whether saving prevents some from being treated as mere means. Surely such omissions show an indifference to the kind of moral considerations that Kantians care about.

Someone might object instead that the Kantians place the poor janitor in a terrible position.¹ If the janitor does not intervene, the suggestion is that the janitor would fail to meet one of his *pro tanto* obligations for reasons that could not justify his failure to do so—he would fail to meet one of his *pro tanto* obligations because he sees that by doing so he can allow a process to unfold in which people are treated as mere means to some further good end. If, however, the janitor does intervene to save the two patients awaiting tonsillectomies, the janitor would be treating the surgeon as a mere means to the end of saving two.

It would be horrible if the Kantian approach suggests that the janitor is neither permitted to intervene or to refrain from intervening. I think that the Kantians have an answer, which is to deny that the janitor is guilty of treating the surgeon as a mere means by denying that the janitor is treating the surgeon as a means to the intended end of saving the patients. It is important to notice that in cases like **Transplant**, **Compulsion**, and **Control**, the janitor could perform acts that the janitor could foresee would result in the death of the surgeon

¹ An anonymous referee raised this objection.

without intending the surgeon's death as either a means or an end. Typically people assume that principles like the doctrine of double effect would allow the janitor to take arms to protect the patients precisely because these are cases where we can describe the janitor's intended end as that of protecting the patient and the means as involving no persons as instruments to the production of that end. While my response does not assume that the doctrine of double effect is correct, notice that the applicability of the principle to cases like these assume that the janitor could perform actions known to result in the death of the aggressor without intending the death as a means or an end. If the janitor can defend the patients foreseeing the death of the surgeon without intending the death of the surgeon and without intending to use the surgeon as a means to protect the patients, I think that the janitor can act without treating the surgeon as a mere means to the end of protecting the patients. This assumes something that I think is quite plausible, which is that if someone can ϕ without intending that some effect E is produced, the agent can ϕ without treating E as a means to her end. In the typical case in which someone treats something as a means to their end, they treat something as an available instrument for bringing about some end and intend that instrument's involvement in their pursuit of their end. From the perspective of the janitor, the surgeon is not an instrument to use so as to save the patients. So, it does not follow from the mere fact that the surgeon is a threat to be dealt with that the surgeon becomes a means to any of the janitor's ends. And, if the surgeon is not necessarily a means to the janitor's ends, there is no obvious reason why the janitor would be treating the surgeon merely as a means when the janitor pursued his end of saving the patients. Unless we assume that threats that the agent intends to deal with automatically become instruments for dealing with these threats, I see no reason to say that the janitor must treat the surgeon as a means to some end if the janitor were to intervene.

I think that the Kantian explanation of (1) in **Transplant** and **Compulsion** is preferable to the explanation Graham offered. Consider a further variant on **Transplant**:

Control: Everything is as it is in **Transplant** with two exceptions. First, the surgeon cannot refrain from killing the two because the surgeon is controlled by remote by a neurosurgeon who has wired our surgeon's brain and taken control of the surgeon's limbs, forcing the surgeon to behave as the surgeon did in **Transplant**. Second, this surgeon would not have killed the two if the neurosurgeon had not wired up the surgeon's brain.

Intuitively, if there are sufficient grounds for acting in **Compulsion**, there are sufficient grounds for acting in **Control**. Intuitively, they seem to be the very same grounds. We cannot explain why (1) is true in **Control** by appeal to (2). The *surgeon* doesn't act in **Control**. The *neurosurgeon* acts and the neurosurgeon's actions cause the surgeon's body to move as if the surgeon is acting. If the surgeon does not act, the surgeon does not act impermissibly. The explanation that Graham offers of (1) in **Compulsion** cannot apply to this case, but the explanation I offered of (1) does apply. And, to the extent that we ought to insist on a common explanation of **Transplant** and **Compulsion**, I think we ought to insist again on a common explanation of **Control** and **Compulsion**. So, by Graham's logic, we should reject (P3). If, however, (2) does not feature in explaining why the janitor is permitted to intervene here, I do not see why we should think that (2) must feature in an explanation of (1) in **Compulsion**. So, again, we would have reason to reject (P3).

Graham anticipates this kind of objection. He offers this challenge in response. Consider another case:

Rifle and Fork: A bystander can redirect an out-of-control train away from ten trapped track workers and toward two other trapped track inspectors. If the bystander does nothing, the train

will kill the ten, and if she redirects the train, it will kill the two. From a distance, a hunter sees that the bystander is about to redirect the train and realizes that he can prevent her from doing so only by shooting her with his rifle. He does so and thereby kills her (2011: 355).

Graham says that it is not permissible for the hunter to do this, and I agree. Graham asks what the difference is between this case where the hunter acts impermissibly and my case where the janitor kills a “passive threat” (2011: 358)? As he sees it, the hunter with the rifle in **Rifle and Fork** is in a similar moral situation as my janitor in **Control**. Both are positioned to use force to kill someone who I claim is not acting impermissibly. Yet, intuitively, it would be wrong for the hunter to shoot in **Rifle and Fork**.

There are a few points to say in response to this. First, he and I might simply disagree about the morality of defending others against passive threats. He seems to think that some account is needed that explains why it would be permissible for an agent to use violence to protect a potential victim from a passive threat. He notes that on the view he considers, it is permissible for an agent to use violence to protect someone from an innocent aggressor if that aggressor would otherwise violate the victim’s rights (Thomson 1991). In **Control**, my surgeon would not *violate* the patient’s rights because my surgeon is not acting. (At least, that is how I see the case.) What I do not see is why the violation of rights is required for the justification of intervention. Why is infringement not enough? As Thomson (1971: 54) notes, we do not have to assume that the fetus is *violating* the pregnant woman’s right to bodily autonomy to justify abortion in the case where the pregnant woman’s life is in danger and while Thomson grants for the sake of her discussion that the fetus is a person with rights, she does not pretend that the fetus is a tiny, morally responsible agent capable of violating a woman’s rights. Similarly, if an unconscious violinist is plugged into your kidneys without knowing that this is happening, the case for the permissibility of disconnection does not seem to depend upon whether the *violinist* is violating your rights. (Not if violations require an act, that is.) So, part of my response is simply that it is permissible to use force to intervene in cases like **Control** where it is clear that the permission to intervene does not depend upon whether the person who you harm to protect another was acting permissibly.

There are further differences between Graham’s **Rifle and Fork** and cases like **Compulsion** and **Control**. In **Compulsion** and in **Control**, the decision to do nothing would involve letting some die so that their organs can be used to save a greater number of patients. This would violate the Kantian prohibition against deciding to act in ways that involve treating others as mere means. (The only case for non-interference seems to be that by doing nothing, ten can be saved using the organs of the two and anyone who chooses not to intervene for these reasons violates the Kantian prohibition against treating others as mere means.) This is not a feature of **Rifle and Fork** and so **Rifle and Fork** is not analogous to **Control**. The Kantian reasoning that helps us understand why it is that the janitor acts permissibly seems not to apply in **Rifle and Fork** to help us understand why someone in the hunter’s position might shoot.

There is yet a further difference between **Rifle and Fork** and the other three examples that might matter as well. In **Rifle and Fork**, we have an agent who is using violence to prevent someone from acting justifiably. This is not a feature of **Control**, **Compulsion**, or **Transplant**. While there might be no absolute prohibition against using violence to prevent someone from performing an action they would be justified in performing, it is plausible that there is a presumption against that sort of thing. Intuitions about **Rifle and Fork** give us little reason to think that intuitions about **Control** cause trouble for Graham’s argument.

My approach to these cases faces one final objection.² As an anonymous referee rightly points out, my approach to **Rifle and Fork** seems to leave me with no response to this loop variant of **Rifle and Fork**:

Rifle and Loop: A bystander can redirect an out-of-control train away from ten trapped track workers. These workers are trapped on a track that loops. If the train goes left, it would ordinarily travel the loop clockwise and crush them. If the bystander would flip the switch, it would send the trolley to the right. The train would go counter-clockwise and crush them. As it happens, there are two workers trapped on the right. If the bystander does nothing, the train travels left killing ten without harming the two. If the bystander flips the switch diverting the train to the right, it would travel right and be stopped by the two workers trapped on the right side without making it around to crush the ten. From a distance, a hunter sees that the bystander is about to redirect the train and realizes that he can prevent her from doing so only by shooting her with his rifle. He does so and thereby kills her.

The objector thinks that the hunter in this case is in the same position as the janitor in **Control**. If this is correct, my account suggests that the hunter in this case can fire on the bystander to prevent the bystander from redirecting the train. The objector thinks that this is deeply counterintuitive. The objector thinks that it is clearly *impermissible* for the hunter to intervene and that my Kantian account cannot explain this fact.

The objection seems to assume that **Rifle and Loop** is similar to **Rifle and Fork**. The objector is right that my account suggests that **Rifle and Loop** is more similar to **Control**. I think that this is the right thing to say, but I admit that this is not intuitively obvious. Indeed, I admit that this is not intuitive. If I were to rely just on my intuitive responses to these cases, I would admit that my intuitions do not suggest that it is wrong to flip the switch in **Rifle and Loop** and that my intuitions do not suggest that it is permissible for the hunter to fire on the bystander if the bystander were to flip the switch. That is *not* to say, however, that I think these responses are inconsistent with my intuitions about these cases. I have no immediate intuition either way. Upon reflection, it seems to me that there is no relevant moral difference between Rifle and Loop and the following:

Rifle and Drag: A bystander cannot redirect an out-of-control train away from ten trapped track workers because the train is running on a straight line with no forks. If the bystander does nothing, the train will kill the ten. If she drags two unconscious workers onto the tracks, she can use their bodies to prevent the train from travelling further and killing the ten trapped workers. From a distance, a hunter sees that the bystander is about to drag the bodies onto the tracks to stop the train and realizes that he can prevent her from doing so only by shooting her with his rifle. He does so and thereby kills her.

My intuitions about this case are quite clear. The bystander would be acting impermissibly if she dragged the bodies onto the tracks and the hunter would act permissibly if he fired his rifle at her to prevent her from using these two workers to stop the train. It seems to me that there is no moral difference between **Rifle and Drag** and **Rifle and Loop** since there is no moral difference between turning a train towards some bodies to use them to stop the train and dragging bodies into the path of the train to stop it. If this is so, then I think my account is right

² This objection is due to an anonymous referee.

about **Rifle and Loop** even if intuition does not vindicate this verdict. I think my account is right to treat **Rifle and Loop** as if it is roughly equivalent to **Control** and **Compulsion** because my account predicts quite plausibly that **Rifle and Drag** is equivalent to **Control**, **Compulsion**, and **Rifle and Loop**.

If the Kantian account that tells us to refrain from treating others as mere means can handle the sorts of objections discussed above, I think it can handle the cases that Graham uses to try to show that 'ought' does not imply 'can'. Since I think that the Kantian account does handle the cases discussed here and that cases like **Control** show that Graham's explanation as to why intervention is permissible in his cases is flawed, I think it is fair to say that OIC has withstood Graham's attack. If it was ever reasonable to think that 'ought' implies 'can', it still is.

References

- Graham, Peter. 2011. "Ought" and Ability. *Philosophical Review* 120: 337-82.
- Thomson, Judith. 1971. A Defense of Abortion. *Philosophy and Public Affairs* 1: 47-66.
- _____. 1991. Self-Defense. *Philosophy and Public Affairs* 20: 283–310.