

A Paradox for Weak Deontology

MICHAEL HUEMER

University of Colorado

ABSTRACT: Deontological ethicists generally agree that there is a way of harming others such that it is wrong to harm others in that way for the sake of producing a comparable but greater benefit for others. Given plausible assumptions, this principle leads to the possibility of paradoxical cases in which each of two actions is wrong, yet the combination of both actions is permissible. Consequentialism provides the most natural solution of the puzzle. (5925 words)

1. A WEAK DEONTOLOGICAL PRINCIPLE

A central and widely-held tenet of deontological ethical theories is that there are ways of harming individuals that cannot be justified by the production of a greater good, or the avoidance of a greater harm, for others. Some deontologists would qualify this by saying that if the greater good to be produced is *vastly* greater, or if it is of a qualitatively different and more important *kind* than the harm in question, then it may after all outweigh the wrongness of harming the individual.¹ But almost any deontological ethicist would agree that one is not justified in harming an individual (in certain ways) merely to produce a *comparable* greater benefit for others, that is, one of roughly the same order of magnitude and the same kind. I shall call the relevant way of harming others ‘deontically proscribed harming’, or simply ‘proscribed harming’, and I shall call this deontological principle *Weak Deontology*:

Weak Deontology: It is wrong to harm a person in the proscribed way, even to produce a greater but comparable benefit for others.

Weak Deontology is neutral with regard to the question of whether proscribed harms can sometimes be justified when *very large* consequences are involved or whether, instead, the deontic proscriptions are absolute. It is also neutral with regard to precisely how the proscribed way of harming others should be characterized. Intuitively, deontological constraints apply to some ways of harming individuals but not to other ways. This is shown by the following well-worn examples:

Trolley Car: A runaway trolley is heading for a fork in the track. If it takes the left fork, it will collide with and kill five people; if it takes the right fork, it will collide with and kill one person. None of the people can be moved out of the way in time. There is a switch that determines which fork the trolley takes. It is presently set to send the trolley to the left. You can flip the switch, sending the trolley to the right instead. Should you flip the switch?²

Organ Harvesting: A doctor in a hospital has five patients who need organ transplants;

otherwise, they will die. They all need different organs. He also has one healthy patient, in for a routine checkup, who happens to be compatible with the five. Should the doctor kill the healthy patient and distribute his organs to the other five?³

Most people believe that one should kill the one person to save the five in the first case but not in the second. A proponent of Weak Deontology can accept this, maintaining that there is *some* relevant difference between the killing in the Trolley Car example and the killing in the Organ Harvesting example. For instance, perhaps the difference is that the harm in the latter case is *intended* as a means to saving the five patients, whereas in the former case the harm is only a foreseen but unintended *side effect* of the action that saves the five other people.⁴ Or perhaps the difference is that in the Trolley Car case, an existing threat is *diverted* from one group of victims to another, whereas in the Organ Harvesting case a distinct threat is *created*.⁵ Or perhaps there is some other difference between the two cases. Regardless, if the killing of the one person is justified in the Trolley Car case but not in the Organ Harvesting case, then the killing is of the ‘deontically proscribed’ kind in the latter case but not the former.

Weak Deontology, then, is very modest. It is committed only to there being *some* kind of harm that cannot be justified by the production of comparable, greater benefits. Only consequentialists reject this, denying that *any* harm is deontically proscribed in my sense. Yet even a principle as modest as Weak Deontology may require revision in the light of its conflict with plausible formal constraints on ethical principles.

2. TWO FORMAL ETHICAL INTUITIONS

Aristotle famously thought that the speed with which a body falls is proportional to its weight.⁶ For example, if one were to drop a heavy cannon ball and a light musket ball, the cannon ball would fall faster. Galileo wondered what would happen if one tied the musket ball to the cannon ball with a piece of string and then dropped them. On the one hand, the new object consisting of the cannon ball plus the musket ball would be heavier than the cannon ball alone; therefore, it would fall faster than the cannon ball would fall by itself. On the other hand, since the musket ball’s natural rate of fall is less than the cannon ball’s, the musket ball ought to act as a *drag* on the cannon ball; therefore, the two balls should fall slower than the cannon ball by itself.⁷ One way of understanding this argument is as appealing to the following constraint on physical theories: no real, observable effect should be held to depend upon one’s method of individuating material objects. Aristotle’s theory violates this constraint:⁸ if we see the cannon ball and musket ball, when connected by a string, as comprising a *single object*, then we predict that they should fall faster than a cannon ball by itself. But if we see the cannon ball and musket ball as *two distinct* objects with a string connecting them, then we predict that they should fall slower than a cannon ball by itself.

An analogous constraint applies to ethical theories: no ethical theory should make a substantive moral judgement turn on one’s method of individuating actions. I call this the principle of Individuation Independence:

Individuation Independence: Whether some behavior is morally permissible cannot depend upon whether that behavior constitutes a single action or more than one action.

Thus, suppose a person does A followed by B. Whether he has done something wrong should not depend on whether ‘doing A followed by B’ counts as performing one action or instead must be seen as performing two distinct actions. I call Individuation Independence a *formal* ethical intuition because, while it places a constraint on ethical theories, it does not by itself entail any first-order evaluative or normative judgement.⁹

Here is a second formal ethical intuition; I call it the ‘Two Wrongs Don’t Make a Right’ principle (‘Two Wrongs’ for short):

Two Wrongs: If it is wrong to do A, and it is wrong to do B given that one does A, then it is wrong to do both A and B.

It seems that there are situations in which two prima facie wrongs *do* make a right. For instance, consider:

Bomb Squad: Sam has recently been called in by the bomb squad to defuse a bomb. He knows that to defuse it, he must cut the red wire and the blue wire simultaneously. If he cuts only one of them, the bomb will explode, killing him and several bystanders. Thus, it would be wrong to cut the red wire alone, and it would be wrong to cut the blue wire alone, but it is not wrong to cut both wires.

This case explains why the phrase ‘given that one does A’ must be included in the statement of Two Wrongs. It would be wrong of Sam to cut only the blue wire. However, *given* that Sam is cutting the red wire, it is *not* wrong but obligatory for him to cut the blue wire. Thus, our Two-Wrongs-Don’t-Make-a-Right principle does not imply that it is wrong for Sam to cut both wires in this case. With this qualification understood, I take it that Two Wrongs is obviously correct.

3. A DEONTOLOGICAL PARADOX

Weak Deontology generates problematic results in cases in which each of two actions imposes a deontically proscribed harm, yet the combination of actions is beneficial to all. For example:

Torture Transfer: Mary works at a prison where prisoners are being unjustly tortured. She finds two prisoners, A and B, each strapped to a device that inflicts pain on them by passing an electric current through their bodies. Mary cannot stop the torture completely; however, there are two dials, each connected to both of the machines, used to control the electric current and hence the level of pain inflicted on the prisoners. Oddly enough, the first dial functions like this: if it is turned up, prisoner A’s electric current will be increased, but this will cause prisoner B’s current to be reduced by twice as much. The second dial has the opposite effect: if turned up, it will increase B’s torture level while lowering A’s torture level by twice as much. Knowing all this, Mary turns the first dial, immediately followed by the second, bringing about a net reduction in both prisoners’ suffering.

Did Mary act wrongly? Call Mary’s turning of the first dial ‘Adjustment 1’, her turning of the second dial ‘Adjustment 2’, and her turning of both dials ‘Adjustment 3’.

It seems that, if Weak Deontology holds, then Adjustment 1 was wrong. Although we

have not specified what the class of deontically proscribed harms is, nearly any advocate of Weak Deontology should agree that Adjustment 1 falls into that category, for it inflicts unjust suffering on prisoner A merely for the purpose of producing a comparable, greater benefit for B. We may take this harm inflicted on A to be a *means* to producing the benefit for B (assume, for example, that the relief to prisoner B would be produced *only* if someone were strapped to A's machine and receiving the increased electric current), so that A is here treated merely as a means. Presumably A has a right, not only not to be tortured, but also not to be subjected to an increased level of torture. After all, if Mary had merely increased A's level of torture *without* producing any benefit for anyone else, she clearly would have violated A's rights, according to nearly any rights theorist. And the fact that she produces a benefit for someone else cannot be relevant to whether she commits a deontically proscribed harm; adherents of Weak Deontology do not in general endorse the infliction of unjust harm on one individual for the sake of preventing greater (but comparable) unjust harms from being done to others--for example, one may not commit a murder to prevent two murders.¹⁰ Finally, we may assume that A has not consented and cannot be expected to consent to the performance of Adjustment 1. On all accounts, then, it seems that Adjustment 1, considered by itself, is a clear case of a deontically proscribed harm.

Once Mary has performed Adjustment 1, it seems that she again acts wrongly, for the same reasons, in performing Adjustment 2. For, despite the fact that prisoner B has just received a reduction in his torture level, it is still a violation of prisoner B's rights to increase his torture level, since B does not deserve any torture at all. We may assume that B refuses to consent to Adjustment 2. Adjustment 2 also treats B merely as a means to benefitting A, and the harm to B is intended rather than a mere side effect, since B's being subjected to the increased electric current is causally necessary to the relief provided for A. On all accounts, then, it seems that Adjustment 2 is deontically proscribed.

Finally, however, it seems that Adjustment 3 is not only permissible but positively desirable, and perhaps even--if it is the best that Mary can do for the two prisoners--obligatory. This is because Adjustment 3 benefits both prisoners, reduces the injustice to both prisoners, and has no morally significant effects on anyone else. Presumably, if Mary had had the option of turning a *single* dial that would have reduced both A's and B's torment by the same amount overall, it would have been right to do so. It seems absurd to suppose that if she must bring about the same effect by turning the two dials, it is then wrong to do so.

This case, then, gives rise to the following argument against Weak Deontology:

1. In general, if it is wrong to do A, and it is wrong to do B given that one does A, then it is wrong to do both A and B.
2. If Weak Deontology holds, then it is wrong to perform Adjustment 1.
3. If Weak Deontology holds, then it is wrong to perform Adjustment 2 given that one performs Adjustment 1.
4. Therefore, if Weak Deontology holds, then it is wrong to perform Adjustment 3. (From 1, 2, 3.)
5. But it is not wrong to perform Adjustment 3.
6. Therefore, Weak Deontology is false. (From 4, 5.)

The problem for Weak Deontology is more general than the Torture Transfer case. Weak Deontology faces a problem as long as it is possible to construct any case with the following

characteristics: An agent has available actions A_1 and A_2 , each of which will affect two persons, P_1 and P_2 . A_1 will (regardless of whether A_2 is also performed) impose a deontically proscribed harm on P_1 while producing a comparable but greater benefit for P_2 . A_2 will (regardless of whether A_1 is also performed) impose a deontically proscribed harm on P_2 while producing a comparable but greater benefit for P_1 . If the agent performs both A_1 and A_2 , the net effect will be a benefit for both P_1 and P_2 . This situation is represented schematically by the following table:

<u>Action</u>	<u>Effect on P_1</u>	<u>Effect on P_2</u>
A_1	-1	+2
A_2	+2	-1
(A_1+A_2)	+1	+1

Here, ‘-1’ represents a loss or harm suffered, while ‘+2’ represents a gain or benefit that outweighs that loss. We assume that the harms produced by A_1 and A_2 are each of the deontically proscribed kind. In any situation of this kind, Weak Deontology judges each of A_1 and A_2 wrong, regardless of whether the other is performed. But if we consider A_1 and A_2 together as a single action, it seems that this conjunctive action, ‘ (A_1+A_2) ’, ought to be permitted, since it benefits all parties concerned. While one might seek to avoid this conclusion by refusing to recognize such ‘conjunctive actions’, one would then run afoul of the principle of Individuation Independence.

This general characterization of the problem is worth attending to, since it shows that the fundamental objection to Weak Deontology is not tied to the specifics of the Torture Transfer case. Even if the Weak Deontologist should show that the Torture Transfer case fails for some reason (for example, if one holds a form of Weak Deontology on which the torture transfers in this case are not deontically proscribed harms), the Weak Deontologist would still need to maintain the strong claim that there exists *no* case of the general kind I have described. In other words, the Weak Deontologist will want to maintain a view of deontically proscribed harms such that it is *impossible* for any pair of actions that could be performed by one person to (i) both satisfy the conditions for being deontically proscribed, and yet (ii) produce an overall benefit for all parties concerned when taken in combination. The torture adjustments described in Torture Transfer seem to me to provide an example of such a pair of actions, on the most plausible accounts of deontically proscribed harm; but even if I am mistaken about that case, it seems to me still very plausible that some such pair of cases could be described.

In the next two sections, I discuss objections to the Torture Transfer case that could likely be generalized to any case of the general kind I have here described.

4. THE ISSUE OF CONSENT

Many deontological ethicists would take the question of whether a harm is deontically proscribed to turn crucially on *consent*, whether actual or hypothetical. They would say that the permissibility of Adjustments 1, 2, and 3 turns on whether the prisoners consent to those adjustments being made, or perhaps on whether it could reasonably be assumed that they *would* consent.

In all probability, while neither prisoner would consent to an adjustment that increases his

own suffering on net (A would not consent to Adjustment 1 in isolation, nor would B consent to Adjustment 2 in isolation), both *would* consent to Adjustment 3. Now, at first glance, this seems only to reinforce the paradox: the consent-based deontological theory must now maintain that Adjustment 1 is wrong, as is Adjustment 2, but the conjunction of Adjustments 1 and 2 is permissible. But the problem might be resolved by the supposition that prisoner A would consent to Adjustment 1 *on condition* that Adjustment 2 is also performed, while prisoner B would consent to Adjustment 2 conditional on the performance of Adjustment 1. In this situation, Adjustments 1, 2, and 3 would all be permissible.

In fact, however, the sort of psychological conditions introduced by the consent theory only open the door to further paradoxes, due to the possibility of individuals with irrational attitudes. Suppose, for example, that prisoner A announces that he will *not* consent to Adjustment 1, not even conditional on Adjustment 2 being done, although he *will* consent to Adjustment 3. Prisoner B makes a similar declaration. Our ethical theory then seems to endorse the irrational conclusion that it is permissible to perform Adjustment 3, though it is impermissible to perform either Adjustment 1 or Adjustment 2, even given that one performs the other.

The consent theory may generate paradoxical conclusions even without any irrationality on the prisoners' part. A might be so indignant at receiving *any* torture (after all, neither prisoner deserves to be tortured at all) that he will not consent to Adjustment 1 under any conditions. Though he would of course *prefer* Adjustment 3 over no adjustment, he would most prefer--and will only *agree* to--the option in which Mary performs Adjustment 2 but not Adjustment 1 (the option that minimizes his own torture). This is not an irrational sort of attitude--to take an analogy, if a mugger is going to either punch me in the nose or shoot me, I prefer that he punch me in the nose, but I still do not *consent* to his punching me in the nose; all I will consent to is his leaving me alone. If we suppose, similarly, that each prisoner will only consent to the option that minimizes his own torture, even though both would *prefer* Adjustment 3 over no adjustment, we have the counterintuitive result that Mary must not perform any adjustment on the torture machines, even though both prisoners would prefer that she perform Adjustment 3 rather than no adjustment at all.

5. ARE THE HARMS OF THE PROSCRIBED KIND?

It seems clear that if the deontologist is to resist the argument of section 3, he must reject either premiss 2 or premiss 3--

2. If Weak Deontology holds, then it is wrong to perform Adjustment 1.
3. If Weak Deontology holds, then it is wrong to perform Adjustment 2 given that one performs Adjustment 1.

The Weak Deontologist must argue that either Adjustment 1 or Adjustment 2, or both, are permissible, despite the Weak Deontological principle. Furthermore, it seems clear that *if* it is permissible to perform Adjustment 1, then it is also permissible, having done so, to go on to perform Adjustment 2. Therefore, we may simply focus on whether it is wrong of Mary to perform Adjustment 2, given that she performs Adjustment 1--if I can defend the affirmative answer to this question, then the wrongness of Adjustment 1 should be at least as evident. Given

Weak Deontology, this question comes down to whether, in the circumstances, Mary's adjustment of the second dial constitutes a deontically proscribed harm.

Considered *by itself*, Adjustment 2 would surely be proscribed, for the reasons discussed in section 3. For example, torturing a person is a paradigmatic rights-violation, and increasing a person's level of torture seems similar enough to also count as a clear rights-violation. Those with deontological intuitions would surely say that it would be wrong of Mary to increase one person's level of torture (for instance, by turning only *one* of the dials) to reduce another person's torture by a greater amount.

If one is to criticize premiss 3, then, one must argue that some feature of the scenario when Mary is turning both dials exempts Adjustment 2 from the deontic proscription that would apply if Mary performed *only* Adjustment 2. Why might that be?

Here is one suggestion. Given that Mary is also performing Adjustment 1, Adjustment 2 does not actually increase prisoner B's torture level. One might think this is true because, after all, prisoner B's torture level is decreasing on net, rather than increasing. This suggestion, however, is mistaken. Adjustment 2 increases B's torture level in the sense that it causes B's torture level to be greater than it would otherwise be--it adds to the electric current applied to B. This is true even if some other, causally independent influence is simultaneously subtracting from B's electric current. To see this, consider a different scenario:

Sadistic Torturer: In another room of the prison, Jane finds a single torture victim, C, attached to a torture device. This device has a dial for increasing the electric current applied to the victim. The torture device is presently malfunctioning, resulting in a steady decline in C's torture level. Unfortunately, Jane is a sadistic torturer. She turns up the dial to partially compensate for the malfunction. Though her action makes C's torture level higher than it otherwise would have been, C's torture level still decreases on net due to the malfunction.

In this case, it would be feeble for Jane to claim that she did not inflict a proscribed harm on C, on the grounds that C's overall torture level did not increase. We cannot, then, claim that Adjustment 2 does not inflict a proscribed harm merely because B's overall torture level does not increase.

A second suggestion is that Adjustment 2 is permissible, given that Mary also performs Adjustment 1, because Adjustment 2 'compensates' A for the unjust harm he received from Adjustment 1. We may, however, suppose that the two adjustments are causally unconnected to one another; Mary would perform each adjustment even if she did not perform the other (perhaps because she is a consequentialist). In that case, it is doubtful that Adjustment 2 constitutes compensation for Adjustment 1. More importantly, in any case, the alleged compensation prisoner A receives from Adjustment 2 comes at the cost of an unjust addition to B's suffering, and this seems to be precisely the sort of tradeoff that commonly accepted forms of deontological ethics forbid. To be sure, if a person suffers unjustly, then he is entitled to compensation--but not at the expense of another innocent person. For example, if I have stolen \$100 from one person, I may not then steal \$50 from a second person, even if doing so somehow enables me to compensate my first victim fully--or so, I assume, a deontological ethicist would argue.

A third suggestion, perhaps the most initially promising, is that what disqualifies both Adjustments 1 and 2 from counting as proscribed harms is the very feature of the case that makes

us confident that their conjunction is permissible--namely, the fact that the two together benefit both prisoners and harm no one else. In other words, perhaps a condition on an act's being deontically proscribed is that it *not* be one of a set of actions that, on net, benefit all individuals concerned.¹¹ This is, of course, an *ad hoc* modification--there is no motivation for it within standard deontological ethical theories, independent of the need to avoid the argument of this paper. More importantly, however, the following cases show the implausibility of the suggestion:

Torture Reduction 1: As in Torture Transfer, Mary comes upon two prisoners being unjustly tortured. This time, there is only one dial that can be turned, which will increase B's pain while lowering A's by twice as much. It happens that A is already suffering more than B, due to an injury that was previously inflicted by another prison guard, such that the level of suffering of both A and B is exactly what it was in the original Torture Transfer case immediately after Mary turned the first dial. After Mary turns the dial, A and B will be at the same level of suffering that they were in at the end of the original Torture Transfer case (see figure 1).

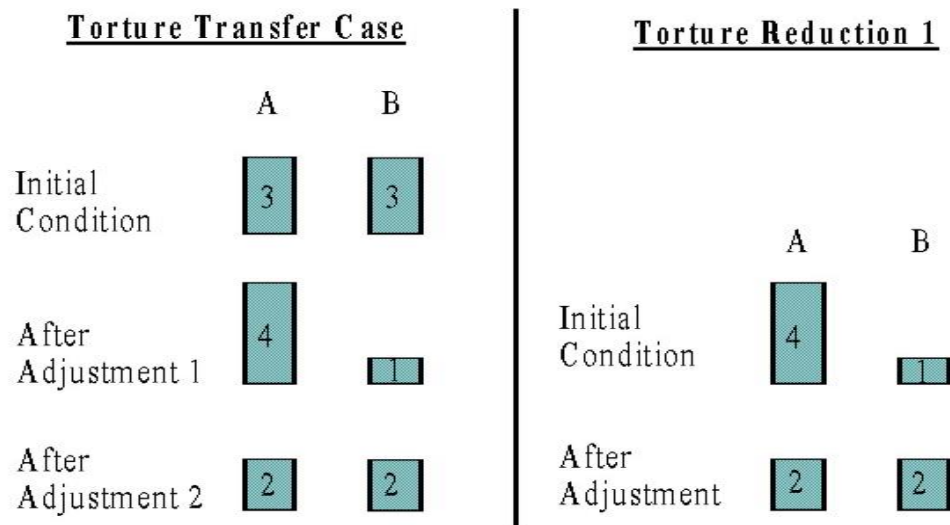


Figure 1. The boxes represent levels of pain to which each prisoner is subjected. In Torture Reduction 1, if Mary performs the one adjustment available, she changes A's and B's pain levels from 4 and 1, respectively, to 2 and 2, just as she does when she performs Adjustment 2 in Torture Transfer.

I take it that in this case, the deontologist will say turning the dial is proscribed. The fact that A is presently suffering more than B is irrelevant; B still has a right not to be subjected to electrical torment, and thus it is wrong to inflict such pain on him even to produce a greater benefit for A.

Torture Reduction 2: As in Torture Reduction 1, except that Mary is suffering memory loss and does not know where A's injury came from: she does not remember whether A is suffering due to an injury previously inflicted by *herself*, or due to one previously inflicted by someone else.

Now, it strikes me as very implausible to suggest that in this case, Mary needs to go investigate

where A's previous injury came from--whether it was from Mary or another guard--in order to find out whether she may turn the dial and palliate A's pain at B's expense. A and B are in the same situation either way. If the question of whether Mary may turn the dial turns on the rights of A and B, surely A and B have the same rights regardless of whether A's injury was caused by Mary or by another prison guard. Of course, Mary's having caused the injury might give Mary a special obligation toward A, to compensate A. But, on standard deontological views, she nevertheless may not do so at B's expense, in a way that treats B as a means, in a way that violates B's rights, etc. To see this point, consider one of Thomson's variants on the Organ Harvesting example:

Remorseful Organ Harvester: As in the original Organ Harvesting case, except that the reason why the five patients need organ transplants is that the doctor previously poisoned all five of them, in an effort to murder them. The poison has different effects on different people; in two of the patients, it is causing lung failure, in two it is causing kidney failure, and in one it is about to cause heart failure. The doctor, however, has since repented of his evil deed, and now wishes to save the five patients' lives by harvesting the organs of the healthy patient.¹²

As Thomson notes, if the doctor does not kill the healthy patient, he will have murdered the five others. Nevertheless, although murdering five people is worse than murdering one, the doctor may not kill the healthy patient--this, at any rate, is the standard deontological intuition. It does not matter if the five organ transplant patients are in their situation because of a previous act of the doctor's. The doctor's previous action may give him special obligations towards his five victims, but he may not discharge those obligations at someone else's expense. Similarly, it hardly seems to matter, in Torture Reduction 2, whether A's current level of suffering results from a previous act of Mary's or from some other cause. Mary's previous act would give her a special obligation to A, but she may not discharge that obligation at someone else's expense.

Finally, consider:

Torture Reduction 3: As in Torture Reduction 2, except that, instead of suffering due to an injury previously inflicted by a prison guard, A is suffering due to someone's having previously turned a dial that caused A greater pain while reducing B's pain. Again, Mary does not remember whether it was she or another guard who turned the dial previously.

I take it that this case is relevantly similar to Torture Reduction 2, so again it does not matter whether it was Mary or someone else who turned the dial previously. And since this case is relevantly similar to the situation Mary is in after performing Adjustment 1 in the original Torture Transfer case, we arrive at the conclusion that, given Weak Deontology, it is wrong for Mary to perform Adjustment 2 in Torture Transfer, even after she's performed Adjustment 1.

One difference between Torture Reduction 2 and Torture Reduction 3 is that in the latter case, unlike the former, Mary's previous action (if it was she who harmed A) would have not only harmed A but also *benefitted B*. We have already seen that Mary's having previously harmed A is morally irrelevant, on standard deontological intuitions. Mary's having benefitted B seems equally irrelevant: while B was perhaps a beneficiary of Mary's earlier turning of the first dial, it is still unjust to inflict pain on him by turning the other dial, since B does not deserve any

torture at all. To take an analogy, note that the fact that Mary has just benefitted B would not entitle her to punch B in the nose, even if after she has done so, B is still better off overall than before she benefitted him. In other words, B's right against being punched in the nose by Mary is not at all diminished by Mary's having previously benefitted B. Likewise, B's right against being tortured by Mary would not be at all diminished by Mary's having previously benefitted B.

Now, if one still thinks that Mary's turning of the first dial entitles her to turn the second dial, consider a final torture scenario:

Delayed Torture Transfer: As in the original Torture Transfer Case, except that Mary performs only Adjustment 1, harming A to help B. She then quits her job. Ten years later, Mary is working in a different prison, where, alas, more unjust torture is occurring. She comes upon a situation exactly like the one she encountered ten years ago. Coincidentally, the same two prisoners are involved, having been arrested for something else. This time, Mary decides to perform only Adjustment 2, harming B to benefit A.

I take it that the standard deontological intuition is that both of Mary's actions are wrong. In particular, Mary's earlier unjust act is irrelevant to whether she may perform Adjustment 2 ten years later. But if so, then we should say the same when the time delay between the two actions is reduced--say, to one second.

6. THE CONSEQUENTIALIST SOLUTION

Consequentialism is the only natural way out of the paradox generated by Weak Deontology: a consequentialist will say that both Adjustment 1 and Adjustment 2 are right, as is Adjustment 3.

For most, this is an uncomfortable conclusion. Most people, whether professional ethicists or not, have strong deontological intuitions. It seems very clearly wrong to kill the healthy patient in the Organ Harvesting example, and it seems almost as clearly wrong to inflict lesser harms on innocent victims in analogous circumstances. I have nothing to say that can significantly assuage these worries. It may be that we are stuck with strongly counter-intuitive results no matter what ethical theory we adopt. The counter-intuitive results of consequentialism have long been known; here, I hope to have contributed to the awareness of the counter-intuitive results of deontological ethics.

It is important to remember that the general problem for deontological ethics can be separated from the details of the Torture Transfer example. To plausibly resist the general argument, the deontologist must claim that it is impossible for there to be a case in which (i) two actions, A_1 and A_2 , each harm someone in the proscribed way, and (ii) A_1 produces a comparable but greater benefit for the person that A_2 harms, while A_2 produces a comparable but greater benefit for the person that A_1 harms. As far as I can see, the only plausible way to maintain that such a scenario is in general impossible is to claim that the satisfaction of condition (ii) would itself exempt A_1 and A_2 from the deontic proscriptions on harming persons. So I think that the dispute over the argument of this paper will have to turn on whether one's having recently performed an action that harmed P_1 to benefit P_2 makes it permissible to perform an action that harms P_2 to benefit P_1 , even when the latter action would otherwise be deontically proscribed. And *that*, I think, will turn on whether harming P_1 to benefit P_2 either weakens P_2 's rights or strengthens P_1 's rights in such a way that one may now harm P_2 to benefit P_1 in a way that would

otherwise be wrong. I have already indicated why I do not believe that either of these things comports with typical deontological ethical intuitions. As a result, even those with strong deontological intuitions should take the consequentialist alternative seriously.¹³

owl232@earthlink.net

NOTES

1. See W.D. Ross, *The Right and the Good* (Indianapolis, Ind., 1988), pp. 21-2; and F.M. Kamm, 'Non-consequentialism, the Person as an End-in-itself, and the Significance of Status', *Philosophy and Public Affairs* 21 (1992): 354-89, p. 383. Robert Nozick, though sometimes taken for an absolutist, withholds judgement on the need for some such qualification (*Anarchy, State, and Utopia* [New York, 1974], p. 31n).

2. This example is from Philippa Foot, 'The Problem of Abortion and the Doctrine of the Double Effect', *Oxford Review* 5 (1967): 5-15, with a modification suggested by Judith Jarvis Thomson in *Rights, Restitution, and Risk*, ed. William Parent (Cambridge, Mass., 1986), p. 96.

3. This example derives originally from James Rachels in informal conversations in the 1960's.

4. See Foot, 'The Problem of Abortion'.

5. See Thomson, *Rights, Restitution, and Risk*, pp. 83-4.

6. See *Physics* IV.8, 216a12-17, in *The Basic Works of Aristotle*, ed. Richard McKeon (New York, 1941).

7. Galileo Galilei, *Dialogues Concerning Two New Sciences* (Buffalo, N.Y., 1991), pp. 62-3.

8. Perhaps more precisely, Aristotle's theory implies that the rate at which an object or pair of objects falls depends upon whether it/they is/are *in fact* one object or two (regardless of how we view them).

9. In my 'Revisionary Intuitionism' (*Social Philosophy & Policy* 25 [2008]: 368-92), I argue that formal ethical intuitions should typically be given greater weight in ethical theorizing than other intuitions.

10. See Samuel Scheffler, *The Rejection of Consequentialism*, revised edn. (Oxford, 1994), pp. 80-81; Nozick, *Anarchy, State, and Utopia*, pp. 28-9; Thomson, *Rights, Restitution, and Risk*, p. 114.

11. This statement is obviously too simple. It would permit one, for example, after saving someone's life, to subsequently beat that person up, since the combination of actions--saving the victim's life and beating him up--still benefits him overall. This sort of objection can probably be

best avoided by modifying the proposed condition so as to require that each of the actions in the set in question be *essential* to the production of the overall benefit for all persons affected.

12. Thomson, *Rights, Restitution, and Risk*, pp. 98-9.

13. I would like to thank Jackie Colby, Andrew Bailey, and the participants of the 2005 Northwest Conference on Philosophy for their insightful comments and questions regarding an earlier version of this paper.