JUDITH JARVIS THOMSON   Turning the Trolley

I

The trolley problem is by now thoroughly familiar, but it pays to begin with a description of its origins.

In "The Problem of Abortion and the Doctrine of the Double Effect," Philippa Foot described a variety of hypothetical cases, in some of which we regard it as permissible for the agent to act, in others of which we regard it as impermissible for the agent to act, and she asked the good question what explains the differences among our verdicts about them.[1] Her aim was to assess whether the Doctrine of Double Effect provides a plausible answer. She concluded that it doesn't, and went on to offer an answer of her own. It is her own answer that will interest us.

Here are two of her hypothetical cases. In the first, which I will call Judge's Two Options, a crime has been committed, and some rioters have taken five innocent people hostage; they will kill the five unless the judge arranges for the trial, followed by the execution, of the culprit. The real culprit is unknown, however. So the judge has only two options:

Judge's Two Options: he can
   (i)   let the rioters kill the five hostages, or
 (ii)   frame an innocent person for the crime, and have him executed.

1. Foot's article was first published in the *Oxford Review* in 1967; it is reprinted in her collection *Virtues and Vices* (Oxford: Basil Blackwell, 1978).

Most people would say that the judge must not choose option (ii).

In the second case, which I will call Driver's Two Options, the driver of a runaway tram has only the following two options:

Driver's Two Options: he can
    (i)   continue onto the track ahead, on which five men are working, thereby killing the five, or
   (ii)   steer onto a spur of track off to the right on which only one man is working, thereby killing the one.

Most people would say that the driver may choose option (ii).

What explains the difference between our verdicts about what the agents may do in these two cases? After all, in both cases, the agents must choose between five deaths and one death.

Foot suggested that the difference is explainable by appeal to two principles. "There is worked into our moral system a distinction between what we owe to people in the form of aid and what we owe to them in the way of non-interference." She suggested that we call what we owe to people in the form of aid our positive duties, and what we owe to people in the way of non-interference our negative duties. She then invited us to accept that negative duties are weightier than positive duties. Markedly weightier. So much so that, as I will express her first principle:

Letting Five Die Vs. Killing One Principle: A must let five die if saving them requires killing B.

That explains why the judge must not choose option (ii).

Things are otherwise in Driver's Two Options. The driver doesn't face a choice between letting five die and killing one, so the first principle is irrelevant to his case. So Foot appealed to a second principle, namely

Killing Five Vs. Killing One Principle: A must not kill five if he can instead kill one.

Given this principle, the driver must choose option (ii), and, a fortiori, he may choose it. So that explains why the driver may choose option (ii). So we now have an explanation of the difference between our verdicts in the two cases.

Her proposal is very attractive. The ideas about negative and positive duties expressed in the two principles are not new, but they are intuitively very plausible, and Foot shows that given those two principles, we

have a satisfying explanation of the differences among our verdicts in all of the cases she drew attention to. I should perhaps add that the principles she appeals to are intended merely as *ceteris paribus* principles, since further information about the six potential victims might make a difference in our views about what the agent may do. For example, finding out that one or more of the six potential victims is at fault for the coming about of the situation they now face might well make such a difference. What she had in mind is just that *other things being equal* the agent must or must not choose such and such an option. It will perhaps be useful, however, if I make explicit the assumption I make throughout that no one of the six in any of the cases we consider is at fault.

But are those two principles true? A doubt might be raised about the second principle.[2] I will ignore it, however. I will assume that the second principle is satisfactory, and focus instead on a doubt that was in fact raised about the first.

In an article provoked by Foot's, I suggested that we should take our eyes off the driver; we should eliminate him.[3] (Make him have dropped dead of a heart attack.) Then let us imagine the situation to be as in the case I will call Bystander's Two Options. A bystander happens to be standing by the track, next to a switch that can be used to turn the tram off the straight track, on which five men are working, onto a spur of track to the right on which only one man is working. The bystander therefore has only two options:

> Bystander's Two Options: he can
>   (i)   do nothing, letting five die, or
>   (ii)   throw the switch to the right, killing one.

Most people say that he may choose option (ii).

If the bystander may choose option (ii) in Bystander's Two Options, however, then Foot's first principle won't do. For if the *Letting Five Die*

---

2. Modeled on John Taurek's doubt about the moral relevance of the numbers in cases where what is in question is distributing a benefit.

3. "Killing, Letting Die, and the Trolley Problem," *The Monist* 59 (1976): 204–17. In that article, the driver was eliminated in favor of a passenger. The case I will call Bystander's Two Options—in which the driver is eliminated in favor of a bystander—comes from my second article provoked by Foot's, namely "The Trolley Problem," *The Yale Law Journal* 94 (1985): 1395–415. Those two articles, along with Foot's and several others on the topic, were helpfully reprinted and discussed, *Ethics: Problems and Principles*, ed. John Martin Fischer and Mark Ravizza (Orlando, Fla.: Harcourt Brace Jovanovich, 1992).

Vs. Killing One Principle is true, then the bystander must not choose option (ii)—for if he chooses option (ii), he kills one, whereas if he chooses option (i), he merely lets five die.

But if the *Letting Five Die Vs. Killing One Principle* is not true, then it cannot be appealed to to explain why the judge must not choose option (ii) in Judge's Two Options.

Perhaps there is some other answer to the question why the judge must not choose option (ii) in Judge's Two Options? An answer resting on the role of a judge in a legal system? No doubt there is. So let us bypass that case. Consider a case I called Fat Man. In this case a fat man and I happen to be on a footbridge over the track. I have two options:

Fat Man: I can
   (i)   do nothing, letting five die, or
   (ii) shove the fat man off the footbridge down onto the track, thereby killing him, but also, since he's very big, stopping the tram and saving the five.

Most people would say that I must not choose option (ii)—just as they would say that the judge must not choose option (ii) in Judge's Two Options. Yet I am not a judge, and no facts about the role of a judge in a legal system could be appealed to to explain why I must not choose option (ii).

Indeed, Foot might have presented us at the outset, not with Judge's Two Options and Driver's Two Options, but with Fat Man and Driver's Two Options, and asked why I must not choose option (ii) in Fat Man, whereas the driver may choose option (ii) in Driver's Two Options. Let us call that Philippa Foot's problem. We might well have been tempted to answer, as Foot would have answered, that the *Letting Five Die Vs. Killing One Principle* explains why I must not choose option (ii) in Fat Man, whereas the *Killing Five Vs. Killing One Principle* explains why the driver may choose option (ii) in Driver's Two Options.

But of course that answer won't do if the *Letting Five Die Vs. Killing One Principle* is false. So Philippa Foot's problem remains with us.

What is of interest is that we also have a second, different, problem before us. As I said, most people say that the bystander may choose option (ii) in Bystander's Two Options. In both Fat Man and Bystander's Two Options, the agent can choose option (i), letting the five die; in both, the agent who chooses option (ii) kills one. Why is it impermissible for

the agent in Fat Man to choose option (ii), but permissible for the agent in Bystander's Two Options to choose option (ii)? Nothing we have in hand even begins to explain this second difference.

Moreover, it is not in the least easy to see what might explain it. Since trams are trolleys on this side of the Atlantic, I called this "the trolley problem." (Besides, that is more euphonious than "the tram problem.")

It spawned a substantial literature. Unfortunately, nobody produced a solution that anyone else thought satisfactory, and the trolley problem therefore also remains with us.

II

A few years ago, an MIT graduate student, Alexander Friedman, devoted a chapter of his thesis to a discussion of the most interesting solutions to the trolley problem on offer in the literature.[4] He did a very good job: he showed clearly that none of them worked. What was especially interesting, though, was what he concluded. He said: the reason why no adequate solution has been found is that something went wrong at the outset. He said: it just isn't true that the bystander may choose option (ii) in Bystander's Two Options!

Friedman didn't offer an independent argument to that effect. He drew his conclusion from two premises. First, there is the fact, which, as I say, he showed clearly, that none of the most interesting solutions on offer worked. We shouldn't take that fact lightly. It is, of course, consistent with there actually being a solution to the trolley problem that nobody has been clever enough to find it. But we should be troubled by the fact that so many people have tried, for so many years—well over a quarter of a century by now—and come up wanting. Friedman's second premise was that it just is intuitively plausible that negative duties really are weightier than positive duties. Thus, in particular, that the *Letting Five Die Vs. Killing One Principle* is true. And if it is true, then of course the bystander must not, after all, choose option (ii) in Bystander's Two Options.

4. A. W. Friedman, *Minimizing Harm: Three Problems in Moral Theory.* Unpublished doctoral dissertation, Department of Linguistics and Philosophy, Massachusetts Institute of Technology (2002).

Friedman therefore said that we should see the (so-called) trolley problem "for what it really is—a very intriguing, provocative, and eye-opening non-problem."

Well, there's an unsettling idea! But if you mull over Friedman's unsettling idea for a while, then perhaps it can come to seem worth taking very seriously. So let us mull over it.

III

Here is a case that I will call Bystander's Three Options. The switch available to this bystander can be thrown in two ways. If he throws it to the right, then the trolley will turn onto the spur of track to the right, thereby killing one workman. If he throws it to the left, then the trolley will turn onto the spur of track to the left. The bystander himself stands on that left-hand spur of track, and will himself be killed if the trolley turns onto it. Or, of course, he can do nothing, letting five workmen die. In sum,

   Bystander's Three Options: he can
           (i)   do nothing, letting five die, or
          (ii)   throw the switch to the right, killing one, or
         (iii)   throw the switch to the left, killing himself.

What is your reaction to the bystander's having the following thought? "Hmm. I want to save those five workmen. I can do that by choosing option (iii), that is by throwing the switch to the left, saving the five but killing myself. I'd prefer not dying today, however, even for the sake of saving five. So I'll choose option (ii), saving the five but killing the one on the right-hand track instead."

I hope you will agree that choosing (ii) would be unacceptable on the bystander's part. If he *can* throw the switch to the left and turn the trolley onto himself, how dare he throw the switch to the right and turn the trolley onto the one workman? The bystander doesn't feel like dying today, even for the sake of saving five, but we can assume, and so let us assume, that the one workman also doesn't feel like dying today, even if the bystander would thereby save five.

Let us get a little clearer about why this bystander must not choose option (ii). He wants to save the five on the straight track ahead. That would be good for them, and his saving them would be a good deed on

his part. But his doing that good deed would have a cost: his life or the life of the one workman on the right-hand track. What the bystander does if he turns the trolley onto the one workman is to make the one workman pay the cost of his good deed because he doesn't feel like paying it himself.

Compare the following possibility. I am asked for a donation to Oxfam. I want to send them some money. I am able to send money of my own, but I don't feel like it. So I steal some from someone else and send *that* money to Oxfam. That is pretty bad. But if the bystander proceeds to turn the trolley onto the one on the right-hand track in Bystander's Three Options, then what he does is markedly worse, because the cost in Bystander's Three Options isn't money, it is life.

In sum, if A wants to do a certain good deed, and can pay what doing it would cost, then—other things being equal—A may do that good deed only if A pays the cost himself. In particular, here is a third *ceteris paribus* principle:

> *Third Principle*: A must not kill B to save five if he can instead kill himself to save the five.

So the bystander in Bystander's Three Options must not kill the one workman on the right-hand track in furtherance of his good deed of saving the five since he can instead save the five by killing himself. Thus he must not choose option (ii).

On the other hand, morality doesn't require him to choose option (iii). If A wants to do a certain good deed, and discovers that the only permissible means he has of doing the good deed is killing himself, then he may refrain from doing the good deed. In particular, here is a fourth *ceteris paribus* principle:

> *Fourth Principle*: A may let five die if the only permissible means he has of saving them is killing himself.

So the bystander in Bystander's Three Options may choose option (i).

Let us now return to Bystander's Two Options. We may imagine that the bystander in this case can see the trolley headed for the five workmen, and wants to save them. He thinks: "Does this switch allow for me to choose option (iii), in which I turn the trolley onto myself? If it does, then I must not choose option (ii), in which I turn the trolley onto the one workman on the right-hand track, for as the *Third Principle* says,

I must prefer killing myself to killing him. But I don't want to kill myself, and if truth be told, I wouldn't if I could. So if the switch does allow for me to choose option (iii), then I have to forgo my good deed of saving the five: I have to choose option (i)—thus I have to let the five die. As, of course, the *Fourth Principle* says I may."

As you can imagine, he therefore examines the switch *very* carefully. Lo, he discovers that the switch doesn't allow him to choose option (iii). "What luck," he thinks, "I can't turn the trolley onto myself. So it's perfectly all right for me to choose option (ii)!" His thought is that since he can't himself pay the cost of his good deed, it is perfectly all right for him to make the workman on the right-hand track pay it—despite the fact that he wouldn't himself pay it if he could.

I put it to you that that thought won't do. Since he wouldn't himself pay the cost of his good deed if he could pay it, there is no way in which he can decently regard himself as entitled to make someone else pay it.

Of how many of us is it true that if we could permissibly save five only by killing ourselves, then we would? Doing so would be altruism, for as the *Fourth Principle* says, nobody is required to do so, and doing so would therefore be altruism; moreover, doing so would be doing something for others at a major cost to oneself, and doing so would therefore be major altruism. Very few of us would. Then very few of us could decently regard ourselves as entitled to choose option (ii) if we were in the bystander's situation in Bystander's Two Options.

IV

Very well, suppose that the bystander in Bystander's Two Options is among the very few major altruists who would choose option (iii) if it were available to them. Should we agree that *he* anyway can decently regard himself as entitled to choose option (ii)?

I stop to mention my impression that altruism that rises to this level is not morally attractive. Quite to the contrary. A willingness to give up one's life *simply* on learning that five others will live if and only if one dies is a sign of a serious moral defect in a person. "They're my children," "They're my friends," "They stand for things that matter to me," "They're young, whereas I haven't much longer to live," "I've committed myself to doing what I can for them": these and their ilk would make sacrificing one's life to save five morally intelligible. Consider, by

contrast, the man who learns that five strangers will live if and only if they get the organs they need, and that his are the only ones that are available in time, and who therefore straightway volunteers. No reputable surgeon would perform the operation, and no hospital would allow it to be performed under its auspices. I would certainly not feel proud of my children if I learned that they value their own lives as little as that man values his.

Perhaps you disagree. I therefore do not rely on that idea. It remains the case that the altruistic bystander is not entitled to assume that the one workman is equally altruistic, and would therefore consent to the bystander's choosing option (ii). Altruism is by hypothesis not morally required of us.

Suppose, then, that the bystander knows that the one workman would not consent, and indeed is not morally required to consent, to his choosing option (ii). The bystander has a permissible alternative, namely choosing option (i)—that is, letting the five die. I think it very plausible therefore that there is no way in which he can justify to himself or to anyone else his choosing option (ii), and thus that he cannot decently regard himself as entitled to choose it.

v

If those arguments succeed, then Friedman was right:

> *Letting Five Die Vs. Killing One Principle*: A must let five die if saving them requires killing B

is safe against the objection I made to it in drawing attention to Bystander's Two Options, since the bystander may not in fact proceed in that case. And if so, two consequences follow. First, Bystander's Two Options is no threat to solving Philippa Foot's problem as she would have done. She *can* explain the difference between our verdicts about the agents in Fat Man and Driver's Two Options as she would have done—that is, she *can* say that the *Letting Five Die Vs. Killing One Principle* explains why I may not proceed in Fat Man, whereas the

> *Killing Five Vs. Killing One Principle*: A must not kill five if he can instead kill one

explains why the driver may proceed in Driver's Two Options. Second, as Friedman said, the (so-called) trolley problem is a nonproblem. The

bystander in Bystander's Two Options is no more free to turn the trolley than I am to shove the fat man off the footbridge into the path of the trolley; a fortiori, there is no difference between our verdict about the agent in Fat Man and our (now corrected) verdict about the agent in Bystander's Two Options to *be* explained.

But even if the (so-called) trolley problem is therefore in one way a nonproblem, it is therefore in another way a real problem, for if the bystander must not turn the trolley in Bystander's Two Options, then we need to ask why so many people who are presented with that case think it obvious that he may. I will make a suggestion about why they do.

There is a question that we need to answer first, however.

VI

What I have in mind is this. I have suggested that consideration of Bystander's Three Options brings out that there is trouble for the idea that the bystander may turn the trolley in Bystander's Two Options. What has to be asked is whether consideration of an analogous case, namely Driver's Three Options, brings out that there is analogous trouble for the idea that the driver may turn the trolley in Driver's Two Options. If it does, then something must be wrong with my arguments in the preceding sections. Let us see why.

By way of reminder, here is Driver's Two Options again.

> Driver's Two Options: he can
>> (i)   continue onto the track ahead, on which five men are working, thereby killing the five, or
>> (ii)   steer onto a spur of track off to the right on which only one man is working, thereby killing the one.

Most people would say that the driver may choose option (ii). What would make that true? According to the *Killing Five Vs. Killing One Principle*, the driver must choose option (ii). I said in Section I that a doubt might be raised about that principle, but I also said that I would ignore it—thus that I would assume the principle is true. Very well then, the driver must choose option (ii). It follows that he may.

Here, now, is Driver's Three Options:

> Driver's Three Options: he can

(i)   continue onto the track ahead, on which five men are working, thereby killing the five, or

(ii)   steer onto a spur of track off to the right on which only one man is working, thereby killing the one, or

(iii)   steer onto a spur of track off to the left, which ends in a stone wall, thereby killing himself.

If consideration of this case makes trouble for the idea that the driver in Driver's Two Options may choose option (ii), then, as I said, something must be wrong with my argument in the preceding sections.

Before attending to the question whether it does, we should take note of an objection that some people make when presented with these cases. If the bystanders in Bystander's Two Options and Bystander's Three Options choose option (i), then they on any view don't kill the five; they merely let the five die. But the people I refer to object that that is also true of the drivers in these cases. After all, choosing option (i) in these cases isn't *steering* onto the track ahead. The drivers merely continue onto the track ahead. We may assume that they don't turn the steering wheel. Indeed, we may assume that they take their hands entirely off the steering wheels—letting the trolleys continue onto the tracks ahead, where *they* (the trolleys) will kill the five.

It really won't do, however, if we can also assume that the drivers themselves started their trolleys, and were steering them up to the time at which the brakes failed—so let us assume it. Suppose Alfred takes out his car and drives toward a restaurant where he expects to meet his friends. He suddenly sees five people on the street ahead of him, but his brakes fail: he cannot stop his car, he can only continue onto the street ahead or steer to the right (killing one) or steer to the left (killing himself). If he doesn't steer to one or the other side, if he simply takes his hands off the wheel, he runs the five down and kills them. He cannot at all plausibly insist that he merely lets them die. So similarly for the trolley drivers. Let us therefore return to them.

In Section III, I argued that if the bystander in Bystander's Two Options is not a major altruist, and, in particular, would not choose option (iii) if he could, then he cannot decently regard himself as entitled to choose option (ii). Here is a reminder of that argument. The bystander in Bystander's Three Options can choose option (iii)—that is, he can kill himself—and

> *Third Principle*: A must not kill B to save five if he can instead kill himself to save the five

therefore yields that he must not choose option (ii). He isn't required to choose option (iii), for choosing option (iii) would be major altruism, and he therefore may instead choose option (i). But he must not choose option (ii). I then went on to say that if the bystander in Bystander's Two Options is not a major altruist, and in particular, would not choose option (iii) if it were available to him, then he cannot decently regard himself as entitled to choose option (ii).

Does that argument have an analogue for the driver in Driver's Two Options? Well, should we accept yet another *ceteris paribus* principle? Namely

> *Variant Third Principle*: A must not kill B to avoid killing five if he can instead kill himself to avoid killing the five.

If so, then since the driver in Driver's Three Options can choose option (iii)—he can kill himself—he must not choose option (ii). Are we to go on to say that if the driver in Driver's Two Options would not choose option (iii) if it were available to him, then he cannot decently regard himself as entitled to choose option (ii)? Can that be right—given that the *Killing Five Vs. Killing One Principle* yields that the driver in Driver's Two Options must choose option (ii)?

Should we accept the *Variant Third Principle*? One thing that might incline us to accept it should be set aside. We might be moved to accept it because we are moved by the fact that it is trolley drivers, and track workmen, whom we are concerned with here. Perhaps we think of a trolley driver as charged, as part of his duties, with seeing to the safety of the men who are working on the tracks. If we do, then perhaps we will think it true that since the driver in Driver's Three Options can choose option (iii)—he can kill himself—he must not choose option (ii). And the quite general *Variant Third Principle* may strike us as true for that reason.

But we should prescind from the possibility that the agents in the cases we are considering have special duties towards the other parties— special in that they are duties beyond those that any (private) human beings have towards any other (private) human beings. So let us return to Alfred. I invited you to imagine that he took out his car to drive toward a restaurant where he expects to meet his friends. He suddenly sees five

people on the street ahead of him, but his brakes fail: he cannot stop his car, he can only continue onto the street ahead, or steer to the right (killing one), or steer to the left (killing himself). I said that if he doesn't steer to one or the other side, if he simply takes his hands off the wheel, then he kills the five. If the *Variant Third Principle* is true, then he must steer to the left, killing himself. Should we agree?

There is reason to believe that we should. Alfred is driving the car; *he* is the threat to people. He will kill five if he does nothing, and it is not morally optional for him to do nothing: he must not kill the five. But since he will kill the five if he does nothing, *he* must be the one to pay the cost of his avoiding killing them. And if the cost is a life, then so be it: he is the one who must pay it.

(This rationale for saying that Alfred must pay the cost in this case is clearly very different from the rationale for saying that the bystander in Bystander's Three Options must pay the cost of saving the five if he is to save them. The difference lies in the fact that it is morally optional for the bystander, but not Alfred, to do nothing: the bystander may decline to do the good deed he would like to do, whereas Alfred must not kill five.)

On the other hand, there is the fact that Alfred is not at fault for the situation in which he now finds himself. Admittedly, it is not morally optional for him to do nothing, but why does he have to kill himself? Somebody has to pay the cost of his avoiding killing the five, but why him? Wouldn't it be fair in him to flip a coin? Is it mere high-mindedness that lies behind the thought that he had better kill himself?

(By contrast, the bystander in Bystander's Three Options must not flip a coin. It is not fair in him to impose a 0.5 risk of death on a person in order to do what it is morally optional for him to not do.)

On balance, I am more moved by the former consideration than by the latter, and thus prefer the idea that though Alfred is without fault, he must kill himself. (After all, the one on the right is also without fault.) And in sum, that we should accept the *Variant Third Principle*.

But I leave it open.[5] What matters, anyway, is what conclusion should be drawn about the driver in Driver's Two Options, whether or not the *Variant Third Principle* is true. Or, to avoid possible interference due to the thought that trolley drivers have special duties to track workmen,

---

5. The view one holds on this matter has a bearing on the views one can consistently hold on other issues in moral theory, on the moral limits to self-defense in particular.

what conclusion should be drawn about Alfred, if it turns out that he cannot after all steer his car to the left, killing himself. Thus if it turns out that he has only the following two options: continue straight (killing five), or steer to the right (killing one). Suppose now that it is true of him that if he had the third alternative of steering to the left (killing himself), he wouldn't choose it. Suppose he wouldn't even flip a coin. Can he decently regard himself as entitled to steer to the right?

It is unjust in him that he would not only not steer to the left, but not flip a coin, if the option of steering to the left were available to him. That remains the case whatever he does. But his not steering to the right would itself be unjust, for his only alternative to steering to the right is killing five. If he knows that those are his only two options, then he cannot decently regard himself as entitled to *not* steer to the right.

Both of those facts mark Alfred off from the non-altruistic bystander in Bystander's Two Options. It is not the case that it is unjust in the bystander that he would not only not throw the switch to the left, killing himself, but not flip a coin as to whether or not to do so, if the option of throwing the switch to the left were available to him. And it is not the case that his not throwing the switch to the right would itself be unjust. For it is morally open to him to do nothing.

This difference between Alfred and the bystander is obviously due to the fact that whereas Alfred kills five if he does nothing, the bystander instead lets five die. Thus it is due to the very difference in weight between positive and negative duties that Foot said we should bring to bear on the cases she drew attention to, and that Friedman said was so plausible. I find myself strongly inclined to think they were right.

I add a proviso, though. I am sure it could go without saying, but it won't: it is one thing to say there is a difference in weight between positive and negative duties, and quite another to say what the source of that difference is. I know of no thoroughly convincing account of its source, and regard the need for one as among the most pressing in all of moral theory.

VII

We should return now to the question I set aside earlier, namely why so many people who are presented with Bystander's Two Options think it obvious that the bystander may turn the trolley.

Friedman suggested that they think this for two reasons, first because of "the subconscious pull of utilitarianism," and second because the bystander's turning the trolley would not be as strikingly abhorrent as the agent's acting in some of the other hypothetical cases described in the literature, such as Fat Man and another case that I will get to shortly.

Friedman's first reason is over-strong. As we know, a number of psychologists have recently been collecting data on people's reactions to cases of the kind we are looking at. Ninety-three percent of the seniors at South Regional High School in Dayton, Ohio, say that the bystander may turn the trolley in Bystander's Two Options! (Actually, I just invented that statistic, but it's in the right ballpark.) I doubt that those students were pulled, consciously or subconsciously, by anything as sophisticated as utilitarianism.

It is surely right, however, to think that the psychologists' informants are moved by the fact that more people will live if the bystander turns the trolley than if he doesn't. A utilitarian would of course be moved by that fact; but so also would many others.[6]

Friedman's second reason is not so much a reason as a restatement of what has to be explained. Consider Fat Man again. It would be strikingly abhorrent for me to shove the fat man off the footbridge down onto the track, thereby killing him, even though more people will live if I do than if I don't. Consider a case often called Transplant: I am a surgeon, and can save my five patients who are in need of organs only by cutting up one healthy bystander—a bystander who has not volunteered—and distributing his organs among the five. Here too it would be strikingly abhorrent for me to proceed. Even more so, in fact. By contrast, it does not strike people generally as abhorrent for the bystander to turn the trolley. However, that difference cannot be thought to explain why people think that the bystander may proceed whereas the agents in Fat Man and Transplant may not. Rather it is what has to *be* explained.

Perhaps the explanation is not deep but right up at the surface?

In those three hypothetical cases, more will live if the agent proceeds than if he or she does not. Yet it isn't open to any of the agents to arrange, by magic, as it were, that there be just that difference, namely that more live. The agents have to *do* something to bring that outcome about. By

---

6. My own account of how facts of that kind figure in support of a conclusion about what a person ought to do appears in *Normativity* (Chicago: Open Court Publishing, 2008).

what means are they to bring it about? Here are the only means by which I can bring it about in Fat Man: move the one into the path of the trolley currently headed toward the five. Here are the only means by which I can bring it about in Transplant: carve the one up and distribute his organs to the five who need them.

There is a wild efflorescence of hypothetical cases in this literature, and much strenuous theorizing about the differences among them. My impression is that when one backs off from all those cases—and one has to back off, lest one get bemused by the details, some of them thoroughly weird—what seems to vary is at heart this: how drastic an assault on the one the agent has to make in order to bring about, thereby, that the five live. The more drastic the means, the more strikingly abhorrent the agent's proceeding. That, I suspect, may be due to the fact that the more drastic the means, the more striking it is that the agent who proceeds infringes a negative duty to the one.

By contrast, if the bystander proceeds, then here are the means by which he brings about that more live: merely turn the trolley.

Some early attempts to explain why the bystander may proceed appealed to that fact about his means. Alas, they didn't succeed, since by turning the trolley, the bystander will kill the one, and thus will infringe a negative duty to the one; and there is no good reason to think that that fact about his means makes his infringing the negative duty count any the less heavily against his proceeding. No matter. We are not asking here why the bystander *may* turn the trolley. What we are asking is only why it seems to so many people that he may. The answer, then, may simply lie in our being overly impressed by the fact that if he proceeds, he will bring about that more live by merely turning a trolley.