ROB LAWLOR

TAUREK, NUMBERS AND PROBABILITIES

ABSTRACT. In his paper, “Should the Numbers Count?” John Taurek imagines that we are in a position such that we can either save a group of five people, or we can save one individual, David. We cannot save David *and* the five. This is because they each require a life-saving drug. However, David needs all of the drug if he is to survive, while the other five need only a fifth each.

Typically, people have argued as if there was a choice to be made: either numbers matter, in which case we should save the greater number, or numbers don't matter, but rather there is moral value in giving each person an equal chance of survival, and therefore we should toss a coin. My claim is that we do not have to make a choice in this way. Rather, numbers do matter, but it doesn't follow that we should *always* save the greater number. And likewise, there *is* moral value in giving each person an equal chance of survival, but it doesn't follow that we should *always* toss a coin.

In addition, I argue that a similar approach can be applied to situations in which we can save one person or another, but the chances of success are different.

KEY WORDS: Aggregation, numbers, probabilities, saving lives, Taurek, weighing goods.

If we can save the life of a single stranger or the lives of five strangers, but we cannot save all six, what should we do? Should we save the five, or should we toss a coin? In recent years, there has been a resurgence of interest in this question. Although John Taurek's contribution is always acknowledged, his paper "Should the Numbers Count?" is often only mentioned in passing. If we do consider Taurek's arguments in detail, however, we will see that Taurek's arguments against aggregation are entirely without force. Nevertheless, I will also argue that Taurek's paper is not entirely without merit, and I will argue in favour of Taurek's claim that there is moral value in giving each person some chance of survival.

My paper will be in four parts. In the first I will consider Taurek's paper in detail, to show that his arguments against aggregation are flawed and unconvincing. In the second part, I will argue in favour of the claim that there is moral value in giving each person a chance of survival, but will argue that this must be weighed against the value of saving the greatest number. In the third part, I will consider two possible objections to the position I have defended. Finally, in the fourth section, I will show that much of what we conclude about numbers can also be applied to probabilities. That is, the arguments considered in response to Taurek and his critics can also be appealed to in order to explain what we ought to do when faced with a situation in which we can *try* to save one person or another, but the chances of success are different in each case.

Part one: Taurek and aggregation

John Taurek imagines that we are in a position such that we can either save a group of five people, or we can save one individual, David. We cannot save David and the five. This is because they each require a life-saving drug. However, David needs all of the drug if he is to survive, while the other five need only a fifth. Most people's natural intuition is that we should save the five. Taurek, however, argues that the way to show equal concern and respect for each person is to give each individual the same chance of survival. The way to do this, Taurek says, is to toss a coin. If we give the drug to the five, we give them a 100% chance of survival, but David has no chance. If we toss a coin, however, all have a 50% chance of survival.

In this paper, I will concede that there does seem to be some moral value in tossing a coin in such a situation, and will even concede that it may be that, if the difference in size between the two groups is small enough, we should toss a coin. I deny, however, that we should *always* toss a coin and, to deny this, I must deny the claim that the numbers should not count. In the cases where we should toss a coin, I claim that this is not because the numbers do not count, but because the moral value of tossing a coin may *sometimes* outweigh the moral value of saving the greater number (consider a group of 1,000,000 and a group of 1,000,001).

In addition to claiming that the best way to show equal concern and respect to each individual is to give each individual an equal chance of survival, Taurek also objects to the idea of aggregating harms. It may be natural for us to think five people dying is worse than one dying, but Taurek (1977, p. 303-4) asks, worse for whom? He writes:

> For each of these six persons it is no doubt a terrible thing to die. Each faces the loss of something among the things he values most. His loss means something to me only, or chiefly, because of what it means to him. It is the loss to the individual that matters to me, not the loss of the individual. But should any of these five lose his life, his loss is no greater a loss to him because, as it happens, four others (or forty-nine others) lose theirs as well. And neither he nor anyone else loses anything of greater value than does David, should David lose his life. Five individuals losing his life does not add up to anyone's experiencing a loss five times greater than the loss suffered by any one of the five. (Taurek, 1977, p. 307.)

And regarding aggregation, Taurek writes

> The claim that one ought to save the many instead of the few was made to rest on the claim that, other things being equal, it is a worse thing that these five persons should die than that this one should. It is this evaluative judgement that I cannot accept. I do not wish to say in this situation that it is a worse thing were these five persons to die and David to live than it is or would be were David to die and these five to continue living. I do not wish to say this unless I am prepared to qualify it by explaining to whom or for whom or relative to what purpose it is or would be a worse thing. (Taurek, 1977, p. 303-4.)

Taurek concedes that some will be "impatient with all this… They will insist that I say what would be a worse (or a better) thing, period." But

Taurek insists: "I cannot give a satisfactory account of the meaning of judgments of this kind."[1] (1977, p. 304.)

Once we understand that Taurek's concern is to show equal respect for each person, and that he therefore wants to give each person an equal chance of survival, his conclusion may not be quite as counter-intuitive as it first seemed.

Nevertheless, there are competing views which too are intuitively appealing – and don't lead to the counter-intuitive conclusion that even if it is a choice between David and five million, we should toss a coin. Parfit, Kamm and Scanlon have all argued that showing equal concern and respect for all requires that more should count as more.[2]

Parfit (1978, p. 301), for example, writes, "Why do we save the larger number? Because we *do* give equal weight to saving each. Each counts for one. That is why more count for more." Similarly, Kamm (2000, p. 221) writes, "If the presence of each individual person would make no difference, this seems to deny equal significance to each person."

---

[1] Taurek's position becomes complicated here, because – immediately after claiming that he cannot make sense of judgements of this kind – he does in fact go on to make sense of judgements of this kind. At this stage, I will assume Taurek meant what he wrote when he claimed not to understand statements of this kind. At the end of this section, however, I will reconsider Taurek's arguments in the light of the arguments that follow this claim.

[2] At least in cases when we are dealing with equal harms or losses on each side. I will not consider the question of whether a smaller harm to a number of people can outweigh a bigger harm to one person. I confess to not knowing what to say about this. On this issue, see Parfit (1978); Scanlon (1998, p. 229-241); Norcross (1997 & 1998); and Ridge (1998).

Kamm also presents the following argument. It would be worse if B and C die, than if B alone dies. Again, it is worse still if B, C and D die. This judgement is "made from a point of view outside of any person" (2000, p. 220), and she writes, "Nonconsequentialists as well as consequentialists, can evaluate states of affairs from an impartial point of view" (2000, p. 220). Kamm then argues that, given that B and C dying is worse than B alone dying, we can "substitute A for B on one side of the moral equation… and get that it is worse if B and C die than if A dies" (2000, p. 220).

At this point, however, I suggest that we have a stalemate. Taurek will simply deny the premise on which Kamm's argument is based. That is, he will deny that B and C dying is worse than B alone dying.

People often find it difficult to believe that Taurek would deny this premise. For many it is natural to think that if death is bad for those that die then presumably it is worse if 5 die than if only one does.[3] But remember that, as Kamm says, this judgement is "made from a point of view outside of any person", and, as we have seen, Taurek claims that he cannot make sense of these judgements. According to Taurek, judgements of better or worse only make sense relative to a particular person's point of view (or a particular goal).

Thus, if it is a choice between B alone dying or B and C both dying, then Taurek will agree that the latter outcome is worse for C, but he will deny that it is worse, period. And, of course, Taurek will agree that we should save C, even if we can't save B. But – for Taurek – this is not because two dying is worse than one dying. Rather, we should save C

[3] I thank Daniel Bristow and an anonymous referee for *Ethical Theory and Moral Practice* for this point.

simply because it is better *for C* if C lives. Thus, Taurek would agree with Kamm regarding what we should do in this situation, but he can agree with this without accepting Kamm's premise that it would be worse if B and C die than if B alone dies. And because he doesn't accept this premise, Taurek will not be convinced by Kamm's argument.

On Taurek's account, A dying and B and C living is the worst outcome for A, but the best for B and C. Likewise, A living and B and C dying is the best outcome for A, but the worst outcome for B and C. But – on Taurek's account – neither outcome is better or worse than the other, period, because such statements simply don't make sense. Presumably, when Kamm says that it would be worse if B and C die than if A alone dies, Taurek must claim that she is simply failing to say anything that he can make sense of.

This does not show that Kamm is wrong. It shows only that her argument will not convince anyone who denies her premise. And Taurek *does* deny it. Similarly, Taurek will simply deny Parfit and Kamm's claims that, by not counting more as more, he does not show equal respect and concern for each person.[4] Thus, there is a stalemate.

If we are to break this stalemate we need arguments that do not rely on premises that will be accepted only by those on one side of the debate. We need arguments that have some hope of winning converts from one side of the debate to the other, by showing that the position defended follows from premises that others accept, or by showing that the opposing view conflicts with moral intuitions we are not willing to reject.

---

[4] See Otsuka (2000, p. 291) for a defence against this objection, on Taurek's behalf.

To his credit, this is exactly what Taurek attempts in "Should the Numbers Count?" I will argue, however, that his arguments fail.[5] Taurek (1977, p. 294) claims that, "at least some of those who accept [that we should save the bigger group] fail to appreciate the difficulty of reconciling their thinking here with other convictions they are inclined to hold with even greater tenacity." Thus, Taurek presents examples in which we have clear intuitions about what is permissible, and then attempts to argue that, if we share his convictions on this issue, we will not be able to reconcile this with our claim that we should save the larger group when faced with the dilemma we are concerned with here.

I will argue that, in each case, Taurek's arguments suffer from one or more of the following flaws. Either he fails to isolate the issue we are concerned with, or he moves from the question of what it would be best to do to questions of permissibility or obligation.

Where he fails to isolate the issue, the most common problem is that his examples often involve implicit appeals to entitlements, which then render his "analogies" disanalogous. His second mistake is that he fails to see that there are often disparities between the level of sacrifice required in different cases (or he fails to see the significance of these disparities). Again, as a result, his "analogies" are not analogous. As a result, I suggest that there is no problem with reconciling our belief that numbers matter with the intuitions we have about the situations described.

---

[5] It will not be possible to go through each and every argument one by one in a short essay such as this. I hope, however, that once I have highlighted the errors in Taurek's approach the reader will be able to return to the Taurek paper and see that these errors do recur throughout.

In the case of his first flaw, moving from questions of what would be best, to questions of what is permissible, Taurek argues for the claim that it is permissible to save David, and in doing so takes himself to be arguing against the claim that numbers count (for example, see 1977, p. 301). But neither the claim that we don't have an obligation to save the five, nor the claim that it is permissible to save the one, entails the conclusion that numbers do not matter. We could argue that it is permissible to save David, but insist that it would be *better* to save the five – because numbers *do* count. As such, the mere fact that it is permissible to save David – even if we accept this as a fact – does not support Taurek's claim that numbers should not count. What Kamm and Parfit claim, and Taurek must deny, is that it would be *better* to save the five. If Taurek wants to oppose Parfit and Kamm here, it is not enough to claim that saving David is *permissible*. Rather, he must claim that saving David is just as good as saving the five, or that tossing a coin is *the best* option. In short, the question we need to ask is not, what is permissible, but what is the morally *best* action: to save the five, the one, or toss a coin? Questions of obligations and permissibility are further and separate issues.

Taurek argues that it doesn't make sense to adopt a point of view outside of an individual, and to ask if a state of affairs is better or worse than another, period. Therefore, we cannot aggregate harms: "should any of these five lose his life, his loss is no greater a loss to him because, as it happens, four others… lose theirs as well." (Taurek, 1977, p. 307.) And because we cannot aggregate, and therefore five deaths is not worse than one, then the way to show equal respect for each individual is to give each individual an equal chance of survival. Taurek uses the following

example in an attempt to show that the idea of aggregating harms is outrageous.

Again, a drug can be used to save David or to save the five. It cannot save them all. This time though, Taurek supposes that David owns the drug. Taurek asks, would you try to persuade David to give the drug away? Taurek (1977, p. 299) argues that this "utilitarian reasoning would be comical if it were not so outrageous."

It is important, however, that we distinguish two issues. First there is the issue of the demandingness of utilitarianism, particularly in the light of David's rights of ownership: "It is his drug" (Taurek, 1977, p. 299). Second, there is the issue of aggregation. In the context of this paper, Taurek clearly takes himself to be arguing against the idea of aggregation. The intuitive appeal of his argument, however, comes from the fact that we would be making an unreasonable demand on David if we asked him to give up the drug for the sake of the five, especially considering his entitlement to the drug.

In contrast, consider a case in which David has no such entitlement. Perhaps the six of them are ill, in the wilderness. By chance, they find a first aid kit, including some of the medicine they require. No one has any more claim to the drug than anyone else.

It is far less absurd now to think that David might be moved by utilitarian reasoning. David may well think that it would be selfish of him to demand that his life should be given as much weight as the lives of the other five combined, all of whom will die without a fifth of the drug. The utilitarian reasoning here is neither comical nor outrageous.

Thus, Taurek's parody of utilitarianism is successful only to the extent that it shows that utilitarianism doesn't recognise entitlements

and is too demanding. Neither of these claims offers the slightest support to the claim that harms cannot be aggregated. If Taurek's argument shows anything, it shows only that, even if numbers do matter, they don't matter as much as rights.

Furthermore, even when Taurek does recognise that there is a question of demandingness, he fails to recognise the significance of this issue. Taurek (1977, p. 301) argues that *if* it is permissible for David to save himself, it must also be permissible for me to choose to save David over the five. That is, it must be permissible for me to take David's perspective, and to do what is best for David – even if he is a stranger to me.

First, I have already argued that if Taurek wants to show that numbers do not count, and that we cannot aggregate harms, he cannot do this merely by showing that it is *permissible* to save David. Rather, he would need to show that this option was *just as good* as saving the five, or that tossing a coin was the *best* option. Second, even if we do stick with permissibility, I am not convinced that Taurek reaches the right conclusion. Not only has he asked the wrong question, it is likely that he has also given the wrong answer. The fact that it is permissible for David to save himself does *not* entail that it is permissible for *me* to save David.

Suppose Kamm is right, and B, C, D, E and F dying is (impartially) worse than A alone dying. The best thing to do would be to save the five, rather than the one. Thus, we ought to save the five. Here though, nothing has been said about what is demanded of you. We are assuming it is as easy to save five as to save one. The former does not require a greater sacrifice than the latter. Thus, if the best action available requires

no sacrifice on my part, we might plausibly claim that it *is* impermissible for me to do less.[6]

The case is clearly different, however, if saving five requires a great sacrifice, such as a sacrifice of my own life, my arm (Taurek, 1977, p. 302) or a friend (Taurek, 1977, p. 294-9).[7] Morality might be thought to be unreasonably demanding if it insisted that it was impermissible for me to save myself (or my arm or my friend). The different level of demandingness makes all the difference.[8]

Of course, if we take Taurek to mean that he cannot make sense of judgements made from the point of view outside of any person, then Taurek would deny the premise that saving the five is better than saving the one. However, this is irrelevant as Taurek's purpose at this point is to address those who do believe we can aggregate harms. His aim is to show that they cannot reconcile this belief with their intuitions on other issues. It is to this argument that I am responding, and therefore it is acceptable for me to appeal to Kamm's premise.

As I suggested earlier,[9] however, it is not clear what Taurek meant when he said "I cannot give a satisfactory account of the meaning of judgements of this kind" (Taurek, 1977, p. 304), given that he then goes on to give an account. One plausible interpretation would be to claim that he was not claiming that such judgements are meaningless. Rather,

---

[6] For a more detailed discussion of the relation between sacrifice and permissibility, see Mulgan (2001, p. 127-139).

[7] Also see Parfit (1978, p. 289-291).

[8] My argument here is influence by Parfit (1978, p. 287-292) and Mulgan (2001, p. 131, 137-8).

[9] See footnote 1 of this paper.

he was merely being modest, and suggesting that the account that he intended to give would not be perfect, but would suffice for further discussion.

I will argue that this interpretation seems to jar with other claims made by Taurek. Nevertheless, I will also consider the implications of reading Taurek in this way, and argue that we can still reject Taurek's arguments.

If we interpret this claim as being nothing more than modesty, it would be difficult to explain *why* Taurek denies the claim that if one person dying is bad for that person then five people dying is worse.

Also, we should consider the context in which Taurek makes this claim. Taurek writes:

> I grant that for each one of the five persons, it would be worse were David to survive and they to die than it would be if David were to die and the five to survive. But, of course, from David's perspective the matter is otherwise… From my perspective, I am supposing in this situation that it does not really matter who lives and who dies. My situation is not worsened or bettered by either outcome…
>
> Some will be impatient with this. They will say it is true, no doubt, but irrelevant. They will insist that I say what would be a worse (or a better thing), period. It seems obvious to them that from the moral point of view, since there is nothing special about any of these six persons, it is a worse thing that these five should die while this one continues to live than for this one to die while these five continue to live. It is a worse thing, not necessarily for anyone in particular, or relative to anyone's particular ends, but just a worse thing in itself. (1977, p. 304.)

And it is here, in response to this line of thought, that Taurek writes:

> I cannot give a satisfactory account of the meaning of judgements of this kind.

In this context, this does seem to be a rejection of the reader's insistence that Taurek tell them what would be a worse – or better – thing, *period*. He is refusing, precisely because he does not understand what they are demanding from him. Furthermore, this interpretation would seem to fit well with other claims that he makes throughout the paper. Consider the passages already quoted in this paper. In particular, remember that he states:

> I do not wish to say [it is or would be a wrong thing] unless I am prepared to qualify it by explaining to whom or for whom or relative to what purpose it is or would be a worse thing. (Taurek, 1977, p. 304.)

Also consider the following:

> Such reasoning seems appealing to many. I find it difficult to understand… (Taurek, 1977, p. 295.)

> Don't you think David might demur? Isn't he likely to ask: "Worse for whom?" (Taurek, 1977, p. 299.)

> This is, of course, not to say that he thinks he is more valuable, period, than any of them, or than all five of them taken together. (Whatever could such a remark mean?) (Taurek, 1977, p. 300.)

> It is not my way to think of them as each having a certain *objective* value... (Taurek, 1977, p. 307.)

> This reflects a refusal to take seriously in these situations any notion of the sum of two persons' separate losses. (Taurek, 1977, p. 308.)

> I want to stress that it does not seem natural in such a case to attempt to add up their separate pains. (Taurek, 1977, p. 309.)

Nevertheless, it is true that, immediately after claiming that he cannot make sense of judgements made from the point of view outside of any person, Taurek does in fact go on to give an account of what it might mean to make judgements of this kind. He writes:

> When I judge of two possible outcomes that the one would be worse (or better) for this person or this group, I do not, typically, thereby express a preference between these outcomes. Typically, I do not feel constrained to admit that I or anyone *should* prefer the one outcome to the other. But when I evaluate outcomes from an impersonal perspective (perhaps we may say from a moral perspective), matters are importantly different. When I judge that it would be a worse thing, period, were this to happen than were that to happen, then I do, typically, thereby express a preference between these outcomes. Moreover, at the very least, I feel

> constrained to admit that I *should* have such a preference, even if I do not. It is a moral shortcoming not to prefer what is admittedly in itself a better thing to what is in itself a worse thing. (Taurek, 1977, p. 304-5.)

Thus, on this account, I take it that these preferences are not merely on a par with normal personal preferences. After all, if I fail to prefer mint chocolate chip ice cream to vanilla ice cream, this is not a moral failing. You can prefer vanilla ice cream if you like. But, on this account, if we evaluate something from a moral perspective, we *do* think it is a moral shortcoming not to have the corresponding preference. On this account, Kamm or Parfit could say that five dying is worse than one dying, and therefore it is a moral shortcoming to fail to prefer the outcome in which the five are saved rather than the one.

Thus, it seems that Taurek *can* make sense of judgements from a point of view outside of any person. If we interpret Taurek in this way, he can no longer resist the claim that five dying is worse than one dying by simply insisting that such statements don't make sense. Rather, he has to show *why*, on the account given, it is not true that five dying is worse than one dying. In defence of his position, he writes:

> I could not bring myself to say to this one person, "I give my drug to these five because, don't you see, it is a worse thing, a far worse thing, that they should die than that you should." (Taurek, 1977, p. 305.)

But it is not clear whom this is likely to convince. Presumably, Kamm and Parfit would be perfectly happy saying such things. Parfit would say

to him, “Each counts for one, and therefore, I’m afraid, more must count for more, and so I must save the five.” And Kamm would explain, “if the presence of each individual person would make no difference, this seems to deny equal significance to each person. And so, I am afraid I must save the five.”

Taurek continues:

> I do not expect that David, or anyone in his position, should think it a better thing were he to die and these five others to survive than it would be were he to survive and they to die. (Taurek, 1977, p. 305.)

But, of course, that is because David would, naturally, evaluate the situation from his own point of view. But, even if it is true that David would not prefer to die and for the five to survive, it doesn’t follow that he would object to the reasoning given above. He might say, “Of course, I would prefer to live, but I must concede that, if I was in your position, I would save the five rather than one, and I accept your decision.” And why would he say that? Because, although he thinks it would be better if he lived (viewing the situation from his own point of view), he can see that, from an impartial point of view, it would be better to save the five.

Commenting on the fact that David would prefer to live, Taurek writes:

> I do not think him morally deficient in any way because he prefers the outcome in which he survives and others die to the outcome in which they survive and he dies. (Taurek, 1977, p. 305.)

But, again, Taurek is ignoring the importance of the demands made on the agent. We can say that five dying is worse than one dying, but still go on to say that we would not consider David deficient for preferring the outcome in which he survives. We would justify this by arguing that morality would be too demanding if it insisted that David not only *accept* our reasoning, but actually *prefer* that outcome too.

However, if we consider the case from the point of view of the rescuer who can save the one or the five, and has no special attachments to any of them, the situation is very different. If *he* prefers the situation in which one lives and five die, it *does indeed* seem plausible to claim that the rescuer is morally deficient.

If the drug is David's, we don't consider him morally deficient if he keeps it to himself. To give the drug away would be to give up his life, and we argue that morality doesn't require agents to make such great sacrifices. But the impartial rescuer doesn't give up anything significant either way, so what reason do we have to deny that he is morally deficient if he saves the one, and lets five die.[10]

Also, reconsider Kamm's argument that if B and C dying is worse than B alone dying then we can conclude that B and C dying is worse than A dying. Originally, I argued that Taurek could resist this argument by denying the premise that B and C dying is worse than B alone dying. And, I argued, he could deny this premise because, on Taurek's account, this statement simply didn't make sense.

On this new interpretation of Taurek, however, this option is not available to him. As we have seen, Taurek does offer an account to make sense of judgements of this sort, but rather tries to argue that it is not

[10] Again, see Mulgan (2001, p.131 and 137-8).

worse for five to die than for one. But how can he resist Kamm's claim that B and C dying is worse than B alone dying. On Taurek's account, the question we have to ask is: is a person morally deficient if they prefer B and C to die than for B alone to die? Presumably the answer must be yes, and so, we must accept Kamm's premise. B and C dying *is* worse than B alone dying. And now it is less clear how Taurek can resist Kamm's conclusion that B and C dying is worse than A dying.

Ultimately, we have two interpretations of Taurek. On the first, he claims that judgements from an impartial point of view don't make sense, but this is undermined by the fact that he himself does make sense of the judgements. And even if he fails to give a "satisfactory" account, he doesn't give any good reason to think a satisfactory account is, in principle, impossible. And, on the second interpretation, according to which we can make sense of judgements from an impartial point of view, Taurek fails to show that five dying is not worse than one dying. Furthermore, the majority of the arguments against Taurek are effective on either interpretation. Thus, on either interpretation, Taurek's arguments lack force.

Part two: the moral value in giving each person an equal chance of survival

My arguments, however, do not effect the claim that, in the circumstances described, there is *some* moral value in tossing a coin, thereby giving everyone an equal chance of survival. I see no way of arguing against this claim, and furthermore see no reason why we

should want to. On the contrary, I intend to defend this part of Taurek's position.

Many people, when presented with Taurek's arguments for the first time, find it highly counter-intuitive, but, at the same time, they see that there is something positive in giving every person an equal chance of survival. When considering the case of one versus two, or even one versus four or five, people will often think this presents a real dilemma. As soon as we consider bigger disparities, one versus a million for example, all sense of there being a dilemma falls away. As Taurek presents his case, this is irrational. If one accepts his argument, we should see that we should toss a coin, regardless of how many lives are at stake. This looks implausible. Even if we were to defend the claim that we should toss a coin when it is one life against a million, it is implausible to think that this answer should be as easy to reach as when it is a case of one versus two. But, of course, this is the implication of Taurek's arguments. If we can save group A or group B, but not both, we should toss a coin. We do not even need to know the numbers involved – according to Taurek, the numbers do not matter.[11] If the groups are roughly the same size, or if one group is a billion times bigger than the other, it makes no difference. For Taurek it is clear what one should do. One should toss a coin. This doesn't look plausible.

It should be noticed, however, that those who simply think we should save the greatest number have a similar problem. If you think you should simply save the greatest number, it makes no difference whether it is one versus two or one versus a million. The answer should be just as

[11] Assuming that we are not considering the possibility of empty groups such that, for example, there is *no one* in group A.

clear in the first case as it is in the latter. But, again, this doesn't seem plausible. Note, I am not denying that we should save the two instead of the one. I am merely denying that it is obvious that this is what we should do in the way that it is obvious that we should save the million instead of the one.

Consider a new scenario. This time, we can save 1,000,000 or we can save 1,000,001. I agree with Parfit and Kamm that numbers do count, but it is not clear that the fact that we can save one extra life by saving the second group should be the deciding factor.[12] Rather, it seems more likely to me that the moral value of giving each person an equal chance of survival *can* outweigh the moral good of saving the biggest group in this particular case.[13]

Now consider this case in contrast with a case in which you can save one, or you can save two. Someone might argue that in both cases it is only one extra life that is at stake in either case, so why should we think that we should save the greater number in one case, but toss a coin in the other. The answer is simple. In the case where we can save one, or we can save two, the extra life we can save is weighed against the moral value of giving one person – the lone individual – some chance of survival. In the case where we can save 1,000,000 or we can save 1,000,001, however, the one extra life we can save is weighed against the moral value of giving *a million* people a chance of survival. This explains why many will have the intuition that we should toss a coin in one case, but not the other.

---

[12] Furthermore, I don't want to suggest that Parfit and Kamm would be committed to this position. It's not clear to me that they need be.

[13] See also Parfit (1978, p. 300-1n).

A summary of my conclusion (regarding numbers)

In short, if Taurek says that there is no value in saving the greatest number, but insists instead that we should simply toss a coin, giving everyone an equal chance of survival, he cannot make sense of the intuition that there is a real dilemma in some cases, but not others. Similarly, if Taurek's critics deny that there is any value in giving everyone a chance of survival, but claim instead that we should simply save the biggest group, it also looks like they won't be able to make sense of the intuition that there is a real dilemma in some cases but not others.[14] However, if we acknowledge that there is some value to saving the greatest number, but also acknowledge that there is some value to giving each person an equal chance of survival, it is *not* irrational to see a dilemma in some cases but not in others. Thus, the position defended here, in which we give some weight to both issues, has the advantage of being able to explain and justify people's intuitions on this matter.

Part three: objections

A weighted lottery?

[14] Unless, of course, they offer an alternative explanation for this intuition.

In response to these arguments, a number of people have suggested that we ought to have a weighted lottery? To some extent I am sympathetic to this view, as my main concern is to acknowledge that there is value in saving as many people as possible and also that there is value in giving each person a chance of survival, and a weighted lottery is indeed one way in which we can take both values into consideration. My concern, however, is that it doesn't do so in the right way.

Essentially, the idea is that if it's a choice between saving one person or another, then we would toss a coin. On the other hand, if it's a choice between saving one or saving five, then we have a lottery that is sensitive to the number of lives that are at stake on either side. So, for example, we might throw a die. If the number 1 is thrown, then we save the one, but if any other number from 2 to 6 is thrown, then we save the five. And so on, for other numbers. This way, everyone is given some chance of survival, but the fact that there is value in saving the greatest number is also taken into consideration by weighing the odds in favour of the larger group (and doing so in proportion to the difference in size.)

So what is wrong with this approach? Consider again the situation in which you can save one, or you can save a million. Now suppose that we have a weighted lottery, and the result comes out in favour of saving the one. Of course, this result is not likely, but it is possible. In which case, we ought to save the one, and let the million die. (Otherwise, why did we bother with the lottery?). The procedure seems to have been fair. We ran the lottery, it was weighted according to the numbers, but as a matter of luck, the result was that a million people died. This seems to be the wrong result, and we could have avoided it by simply saving the million to begin with instead of running a lottery. And we could justify

this decision by saying simply that the moral value of saving a million lives (and making absolutely *sure* that we save them) outweighed the moral value of giving one person some chance of survival. Ultimately though, if your intuitions are not the same as mine, I will not have much more to say to try to convince you. My main point in this paper has just been to argue for the more modest claim that there are two values relevant to the scenarios Taurek discusses, not just one – and on this point we would be in agreement. Exactly how the two values should be incorporated into our moral judgements in these situations is a further question I do not intend to discuss further in this paper.

A levelling down objection

Another way in which some have responded to my arguments is to suggest that the *principle* of giving each person an equal chance of survival may invite a levelling down objection.[15]

Consider the following situation;

> Suppose that I can try and save either Smith or Jones from drowning. If I opt to save Smith there is a 10% chance that I'll succeed. If I opt to save Jones there is a 20% chance of succeeding. However, if I choose to try to save Jones, it is also possible to make

[15] These concerns were expressed by two anonymous referees for *Ethical Theory and Moral Practice*.

> the rescue attempt more difficult for myself by tying bricks to my feet, thereby reducing my chances of success to 10%.[16]

The worry is that if our principle is that we should give each person an equal chance of survival, it seems that we ought to toss a coin, and if it lands on heads, I should do my best to save Smith. If it lands on tales, however, I should tie a couple of bricks to my feet, and then try to save Jones. This would give Smith and Jones an equal chance of survival. Clearly, this is absurd.

My response to this objection is simply to stress that Taurek was not looking for, or trying to defend, a general principle comparable, for example, to the utilitarian principle that we should maximise happiness. Rather, his focus was much narrower than that. He was simply trying to answer a particular question about a particular situation. Taurek writes:

> The situation is that I have a supply of some life-saving drug. Six people will all certainly die if they are not treated with the drug. But one of the six requires all of the drug if he is to survive. Each of the other five only requires one-fifth of the drug. What ought I to do? (Taurek, 1977, p. 294.)

And it is only in relation to this particular question that Taurek claims that we ought to toss a coin, thereby giving everyone an equal chance of survival. This response is intended as an answer to the question: what ought I to do in *this* particular situation. It does not apply in other

---

[16] This scenario was suggested to me by one of the anonymous referees for *Ethical Theory and Moral Practice.*

situations, unless they are relevantly similar. Likewise, this paper began with the following question:

> If we can save the life of a single stranger or the lives of five strangers, but we cannot save all six, what should we do?

And it is in relation to *this particular question* that I agree with Taurek that there is value in giving each person an equal chance of survival, but also agree with Parfit and Kamm (among others) that there is also some value in saving the greater number.

By introducing a new situation in which the chances of survival are different, we have simply changed the subject and asked a different question. Equally, someone might say, what if one of the people is an evil mass-murderer who doesn't deserve to live? Or what if the person who needs all the drug is on the verge of finding a cure for cancer? And so on. But these are just different questions, and therefore we should not be at all surprised if they require different answers.

Nevertheless, interesting issues arise when we change the circumstances by introducing probabilities. Furthermore, there are interesting parallels between this new question and the original question we started with. So, suppose that I can try to save either Smith or Jones from drowning, but I definitely cannot save both. If I opt to try to save Smith there is a 10% chance that I'll succeed. If I opt to try to save Jones there is a 20% chance that I'll succeed. What ought I to do?

Part four: Taurek and probabilities

As we have seen, we clearly do not want to conclude that I should give each person an equal chance of survival. We could, however, still conclude that I should toss a coin to decide who we will *try* to save. If heads, I will try to save Smith, and if tails I will try to save Jones (without tying bricks to my feet).

However, many will object to this response. They will suggest that I am more likely to succeed if I try to save Jones than if I try to save Smith and, for this reason, I ought not to toss a coin. Instead, I ought to save Jones. I suspect that many would consider this to be the common-sense answer. Furthermore, it does seem to be the parallel of the common-sense answer regarding numbers. In the original situation, the common-sense answer was that we should save the greatest number. In this new scenario, the common-sense answer is that we should try to save the person we are most likely to be able to save. Taurek rejected the common-sense answer in the first situation, but could he align himself with common-sense in this new situation?

In this section I will demonstrate that Taurek's reasoning can be applied in this new situation as well as to the original, and I will argue that, if he is to be consistent, Taurek ought *not* to accept the common-sense answer, but should again insist that I should toss a coin.

In this situation, the numbers are the same on either side, but the chances of success are different. So how would Taurek's reasoning apply to this new scenario? What is better: a situation in which I am swimming out to try to help Smith, or a situation in which I am swimming out to try to help Jones? Presumably, Taurek will claim again that this question does not make sense to him. Rather, he will claim that it is better *for*

Jones (but worse for Smith) if I try to save Jones, but it is better *for Smith* (but worse for Jones) if I try to save Smith. For Taurek, it doesn't make sense to say that one is better than the other, period.[17]

As such, it looks like Taurek should argue as follows:

It would be better for Jones if I try to save Jones
It would be better for Smith if I try to save Smith
But we cannot say that one option is better than another, period
Therefore, I ought to toss a coin to decide who I ought to save

Again, however, I think that Taurek's answer and the common-sense answer are both wrong, because they each focus on only one of the relevant factors. And again, we can give a number of different scenarios and stress that neither Taurek nor those who adopt the common-sense approach can make sense of the fact that we might see a dilemma in one case but not another.

For example, consider two more situations like the one just described, only this time the probabilities are as follows:

a) 19.9% chance of saving Smith, 20% chance of saving Jones.
b) 0.01% chance of saving Smith, 99.99% chance of saving Jones.

[17] Or, alternatively, he will concede that it does make sense, but will deny that trying to save one is better than trying to save the other. He will say "I could not bring myself to say to Smith, 'I will try to save Jones because, can't you see, he has a 20% chance of survival, so it is better that I try to save him.' I do not think that Smith would agree that this is better. He would prefer that I try to save him, and I would not think him morally deficient…" See Taurek, 1977, p. 305. And my objection to this line of argument will be the same as before.

Should I simply try to save Jones, or should I toss a coin? Should we give the same answer in both cases? Taurek – it seems – should be committed to tossing a coin in both cases. The common-sense answer is committed to saving Jones in both cases. If you think I should toss a coin in scenario a) but ought to save Jones in scenario b) then – unless you have some other explanation – I suggest that this is because you agree that (in this situation) there is value in my trying to save the person I am most likely to save successfully, but you also see that there is some value in tossing a coin such that there is at least some chance that I will try to save Smith. In scenario b), the latter outweighs the former, and in scenario a) the former outweighs the latter.

ACKNOWLEDGEMENTS

I am grateful to Brad Hooker for his generous comments on previous drafts of this paper, and I thank those who attended the Pavia postgraduate conference in political philosophy 2004 for their discussion of a version of this paper, which I presented there. I am also grateful to two anonymous referees from *Ethical theory and Moral Practice* for comments that led to a number of significant revisions and improvements.

REFERENCES

Kamm, F., Nonconsequentualism in Hugh LaFollette ed. *The Blackwell Guide to Ethical Theory*, Blackwell Publishers, 2000, pp. 205-227.

Mulgan, T., *The Demands of Consequentialism*, Oxford University Press, 2001.

Norcross, A., Comparing Harms: Headaches and Human Lives, *Philosophy and Public Affairs* 26, no. 2 (1997), 135-67.

Norcross, A., Speed Limits, Human Lives, and Convenience. A Reply to Ridge, *Philosophy and Public Affairs* 27, no. 1 (1998), pp. 59-64.

Otsuka, M., Scanlon and the claims of the many versus the one, *Analysis* 60 (2000), pp. 288-93.

Parfit, D., Innumerate Ethics, *Philosophy and Public Affairs* 7, no. 4 (1978) pp. 285-301.

Ridge, M., How To Avoid Being Driven to Consequentialism: A Comment on Norcross, *Philosophy and Public Affairs* 27, no. 1 (1998) pp. 50-58.

Scanlon, T., *What We Owe to Each Other*, Harvard University Press, 1998.

Taurek, J., Should the numbers count? *Philosophy and Public Affairs* 6 (1977), pp. 293-316.

Inter-Disciplinary Ethics Applied

A Centre for Excellence in Teaching and Learning

University of Leeds

8-12 Fenton Street

(Off Woodhouse Lane)

Leeds

LS2 9JT

Email: r.s.lawlor@leeds.ac.uk