Impermissible Altruism: A Response to Pummer*

Abstract: I raise three problems for Pummer's account of impermissible yet praiseworthy actions. First, Pummer does not offer a plausible account of the motives of his agents, and this presents a major obstacle to evaluating those agents' conduct. Second, the success of his argument hinges on a problematic way of individuating actions. Finally, there is a better candidate for a type of action that is both praiseworthy and impermissible, but this candidate presents novel challenges for Pummer's account.

In a stimulating paper published recently in the pages of this journal, Theron Pummer argues that actions can be impermissible yet overall praiseworthy. I shall argue that in making his case, Pummer relies on a problematic way of individuating actions. On a more plausible way of individuating actions, the argument does not go through. In addition, the actions in the types of cases Pummer offers are difficult to evaluate because it is not clear what motives could lead agents to act in the ways described. At the end of this commentary, I will suggest that there *may* be cases in which an action is both impermissible and praiseworthy, but those are cases of a fundamentally different type, a type that creates a different problem for Pummer's account.

1. Pummer's Argument

Pummer's candidates for praiseworthy impermissible actions involve, generally speaking, cases in which a person does something altruistic and possibly (very) costly but where a morally better alternative that isn't costlier is available. In some of the cases Pummer offers for consideration, the altruistic alternatives involve conflicts of interest between beneficiaries (e.g., you can save either A or (B and C)), while in other cases, the best alternative is Pareto-superior (e.g., you can save either A or (A and B)). The claim that a costly altruistic choice may nonetheless be impermissible seems most plausible in cases with a Pareto-superior alternative, particularly an alternative that's significantly better than the next best altruistic option. Consider:

> **Costly No-Conflict:** Two strangers face a deadly threat. You can do nothing, sacrifice your legs in a way that would save one stranger, or sacrifice your legs in another way that would save both.[1]

Given the cost, an agent in this case can permissibly do nothing, and choosing to save a stranger is praiseworthy. But given that she has the option of saving the second innocent stranger without increasing the cost, a choice to save only one seems impermissible. Pummer does not insist on this type of example, however, since he thinks that it is unclear that choosing the suboptimal alternative here is *overall* praiseworthy, as opposed to simply *pro tanto* praiseworthy. One may think that in picking the suboptimal alternative, an agent is allowing an innocent person to die completely gratuitously, which is blameworthy. If the praiseworthiness of saving one does not suffice to cancel out the blameworthiness of letting another die for no reason, then the action is not overall praiseworthy. Better examples, Pummer tells us, involve either cases in which there are conflicts of interest between beneficiaries or cases in which there is a Pareto-superior alternative, but it involves a relatively small amount of additional benefit, in proportion to the benefit the agent brings about. Consider an example of the latter type:

> **Red/Green No-Conflict**: You can do nothing, press a red button thereby saving the lives of ten strangers, or press a green button thereby saving the lives of these same ten strangers and saving another stranger's finger. Pressing either button will also cause you to drop into a fiery pit and die.[2]

And an example of the former type:

> **Red/Green Conflict**: You can do nothing; press a red button thereby saving the lives of ten strangers and saving the life of stranger X; or press a green button thereby saving the lives of these same ten

strangers and saving the lives of strangers Y and Z. Pressing either button will also cause you to drop into a fiery pit and die.[3]

In these two cases, it seems perfectly permissible to do nothing and avoid death. Doing nothing would be neither praiseworthy nor blameworthy. If, on the other hand, you choose to save several strangers at the price of your life, you would be doing something (highly) praiseworthy. But if you opt to save ten strangers without saving the eleventh stranger's finger, or you choose to save ten plus stranger X rather than saving ten plus Y and Z, the praiseworthy thing, on Pummer's reckoning, would be at the same time impermissible. While there is sufficient justification to save no one and your choice to save someone is praiseworthy, there is no sufficient justification to save or otherwise help the smaller number rather than the larger. Importantly, neither the failure to save an eleventh stranger's finger nor the failure to save Y and Z rather than only X would be sufficiently blameworthy to counterbalance the praiseworthiness of saving ten people: the loss of a finger is a relatively minor bad compared to the good of saving several lives (at the price of your own), and saving Y and Z, while better from a numbers point of view, would still involve sacrificing the life of X, so opting to save X instead of Y and Z is not so blameworthy as to outweigh the praiseworthiness and make the action overall blameworthy. This is what distinguishes cases *Red/Green No-Conflict* and *Red/Green Conflict* from *Costly No-Conflict*, mentioned earlier. Pummer concludes therefore that choosing the second-best altruistic alternative in *Red/Green Conflict* and *Red/Green No-Conflict* is both overall praiseworthy and impermissible.

I find Pummer's cases intriguing, but I am not convinced. In what follows, I will suggest that first, the types of cases discussed by Pummer involve a feature that makes the actions described difficult to evaluate (Section 2). Second, even if we bracket the issue with this feature, the argument does not go

through (Section 3). Finally, I hypothesize that there may be actions that are both praiseworthy and impermissible, but they are of a type such that their possible existence points to a problem in Pummer's notion of praiseworthiness (Section 4).

2. A Question of Motives

I wish to begin with the following observation: The motive for choosing the suboptimal alternative in the sorts of cases Pummer describes appears inscrutable – it is not clear what the agent's rationale to opt for the equally costly but less good option can possibly be. This is because once the cost is fixed, it is simply rational to opt for the larger benefit even from a self-interested point of view (you get more "bang for the buck"). It is, therefore, odd to deliberately choose the smaller benefit. This oddity is a problem because it is difficult to evaluate actions whose motives seem inscrutable.

Of course, we can imagine perfectly intelligible motives such as hatred or malice toward the agent who doesn't get saved, but Pummer does not suggest this, because bad motives threaten to erode the action's praiseworthiness, which would undermine the argument.[4] In order to avoid the bad motives problem, Pummer has us imagine that the agents in question have the best possible motives available. What can the best motives be? The motives Pummer envisions include what he dubs "innumerate altruism" and some variations thereof. Innumerate altruism, as Pummer defines it, involves being moved to help someone but being indifferent to *how many* you help.[5] This may be imagined to be the motive behind choosing the second-best alternative in *Costly No-Conflict*. Variations include "imbalanced altruism" – being "moved by the plights of each and indeed moved more by the plights of more" but not "motivated by an appropriate balance of the plights on each side,"[6] which may be the motive behind the suboptimal choice in *Red/Green Conflict* – and "innumerate smaltruism" – being "moved more by more big plights than by fewer" but "not moved more by only a few more relatively

small plights than by fewer,"[7] which may be imagined to be the motive behind *Red/Green No-Conflict* (the motive may incline you to save lives but not fingers).

But it is difficult to see how the agents in question could be motivated in the ways suggested. Pummer attempts to make the motives appear plausible by drawing an analogy to charitable giving. He notes that some real-world charitable contributions appear motivated in something like the innumerate altruist's way – a person wants to give money to some good cause, and that's the scope of the intention. Many a good cause would do so far as this type of altruist is concerned. Perhaps, ordinary altruists do not care about numbers, or the balance of plights, or relatively small additional goods – they just want to do enough good of some sort or other.

The problem is that the cases discussed by Pummer are quite different from ordinary charitable giving. There are two main differences which make innumerate altruism and its cousins implausible here. First, in the ordinary case, the agent is not presented with a precise number of equally costly alternatives one of which is clearly morally better. This is because in the real world, it is not obvious either how many alternatives there are or how to rank them. Given the overwhelming number of options and the work involved in ranking them, a person may, quite understandably, opt to satisfice and donate to a cause that seems good enough.

Second, even if we set aside the costs associated with attempts to reduce ignorance and uncertainty – as I believe Pummer wants to do[8] – and we assume that a person is aware of research on effective altruism, say, so that for her, no significant cost would be involved in improving on innumerate altruism or some variation thereof, a real-world altruist is generally promoting *some* important value. The choice to support the local symphony, for example, even when one is aware of the fact that the

money can be sent to a charity that saves lives, is neither irrational nor inexplicable. The ordinary altruist, thus, has an answer to the question, "Why did you do that?" which makes her behavior intelligible. What of the Pummerian agent? Why does *she* act as she does?

The answer we get from Pummer is that she just wants to help someone or other, and this desire would be satisfied by picking the suboptimal choice. But how precisely does not caring about numbers and proportionality lead a person to choose the suboptimal alternative in the sorts of cases Pummer envisions?

Let us grant that Pummer's satisficing altruist does not actively prefer to save the smaller number rather than the larger, since if she did, she would be an altruist who cares about the numbers but opts to save the fewer. How else must we construe the agent's motive? Perhaps the agent deliberately chooses at random between the smaller and the bigger number.

When one of the alternatives is Pareto-superior, there seems to be no reason whatsoever to pick at random, and the choice is extremely puzzling. What of scenarios without a Pareto-superior alternative, that is, cases of conflict such that you can save either a larger or a smaller number of people at the same cost, but they are not the same people?

As Pummer himself notes, some philosophers, such as John Taurek, have argued that one *should* choose randomly in such cases, and that the numbers shouldn't count.[9] Maybe Pummer's altruist is persuaded by Taurek-style arguments. But if she is (and we assume the arguments are reasonable), then it is not at all clear that she would be acting impermissibly.[10] Not surprisingly, Pummer assumes that choosing the smaller number is impermissible[11], and he rules out the possibility that the agent is

deliberately randomizing, writing:

> You do not deliberately employ a randomizing procedure, but you are just as likely to save the lesser number in order to help someone as you are to save the greater number in order to help someone, and on this occasion, you happen to go for the lesser number.[12]

So the idea is that your altruism is such that it does not incline you to save the bigger number any more than it inclines you to save the smaller. We are now back where we started and have not gotten a satisfactory answer to the motivation question. Once again, while the motivation Pummer ascribes to his agents may be plausible in the case of the garden variety altruist who is simply not paying attention to all the possible options and is unmoved by the options she isn't thinking about, or else is aware of the fact that her donation can be used to combat famine, but she chooses to send money to the symphony, Pummer's altruist picks an alternative that is inferior by her own lights. This is puzzling on the part of someone who has just demonstrated willingness to make the sacrifice that would be involved in choosing the morally best option. The choice therefore seems irrational and inexplicable.

Perhaps, we can, on behalf of Pummer, try to imagine an understandable motive that can help make the action intelligible.[13] Unfortunately, I do not think that this charitable move succeeds. Note first, that if the case involves a Pareto-superior alternative, then it is very difficult to imagine a morally acceptable motive for choosing the suboptimal alternative. Let us, however, focus on cases without a Pareto-superior alternative. It is not an accident that Pummer ascribes to his agents the motives that he ascribes. Rather, he is trying to avoid two pitfalls. On the one hand, there is the danger that if he were to ascribe a morally acceptable ordinary motive of the sort we can expect of agents in the types

of cases he envisions – for instance, loyalty to family and friends – the actions in question would no longer seem impermissible.

On the other hand, we can think of bad and otherwise morally problematic motives such as biases or animosity toward individuals or members of a given group that would help rationalize the actions in question but would undermine their praiseworthiness, and perhaps, render them blameworthy.[14] For instance, if Peter risked his life to save one non-Asian person, but he did so only because he would then have an excuse not to save two Asians, his action, despite resulting in a life saved, is not overall morally praiseworthy and may even be overall blameworthy.

We can, however, press further and attempt to imagine a case in which the motive falls in a goldilocks zone such that it makes an altruistic action of the type Pummer envisions intelligible while keeping it both overall praiseworthy and overall impermissible. Consider the following scenario:

> **Benighted Red/Green Conflict**: The agent chooses, at the cost of his own life, to save ten strangers + X rather than saving the ten strangers + Y and Z, where this choice is motivated by his (mistaken) thought that since women and children are often given priority in rescue operations, it also makes sense to give priority to saving X here, since X is a woman, over saving Y and Z, who are men.[15]

I suspect that readers' intuitions would differ when it comes to the question of whether or not the agent in Benighted is acting impermissibly overall, but perhaps many will intuit that he is. Let us assume for the sake of argument that what the agent does is impermissible. (Some readers who think that the agent is acting impermissibly may, on the other hand, think that the agent is sufficiently blameworthy to make the action overall unpraiseworthy, but let us set this issue aside as well.) The problem is that if the agent is acting impermissibly, the impermissibility is not of the sort Pummer

needs for his account as it derives from the questionable motive for choosing to save X rather than Y and Z. If, contrariwise, the motive is a permissible feeling of loyalty to one's group, then the action would no longer be impermissible. This is a problem because if the only intelligible altruistic actions that appear both praiseworthy and impermissible are impermissible on account of a problematic motive, Pummer's account loses its force. What Pummer needs is an argument for the claim that choosing a suboptimal alternative that's no costlier than the best alternative in cases of self-sacrifice is impermissible simply because you are saving a smaller number of people and not because you have false beliefs about the value of (some) people's lives. The idea is that conditional on a decision to bear a certain cost in the service of altruistic goals, an agent has a duty to pick the best alternative despite not having a duty to bear the cost to begin with. Cases such as Benighted do not speak to this claim. We cannot persuade the agent in Benighted to save the bigger number of people by giving him arguments to the effect that saving a smaller number of people is impermissible given a fixed cost. What we would have to do is persuade him that men's lives count for no less than women's lives.

One can argue, however, that Benighted gives us an example of an action that is both overall praiseworthy and overall impermissible. The example can be said to lend support to Pummer's claim that there are such actions though it may not support any general conclusions about the duties of altruists or real-world charitable giving of the sort Pummer favors. There are two points I wish to make in this connection. First, even this more limited claim depends on the assumption that the agent in Benighted is in fact acting impermissibly. As mentioned earlier, I granted this assumption for the sake of argument, but it can be disputed. Second, in the next section, I will give arguments which, if correct, will show that the claim depends on a misleading way of individuating actions.

Importantly, Pummer is aware of the fact the motives of the agents he describes may seem irrational. In a footnote, he suggests that even if this is so, it would not mean that the agents cannot be praiseworthy.[16] There are two points to make in response. First, the motivation of Pummer's agents appears *inexplicably* irrational. Importantly, many irrational actions are not inexplicable, since we understand the agent's reasons; it is just that those reasons fail to conform to rational standards. For instance, it may be irrational to try to aid a loved one's recovery through magical thinking, but we can morally evaluate such an action because we understand why it was done: we know that the agent intends to do good though she may choose ineffective (but harmless) means toward that end. In the case of Pummer's agents, by contrast, evaluation is difficult because we do not understand the motivation. Second, arguably, we can be more generous with praise and bestow it on the actions of agents who behave seemingly irrationally and perhaps even inexplicably, but we must not charge an agent with a failure to act permissibly if we have no idea why the agent acted as she did.

As a side note, I do not think that the actions of the ordinary altruist, by and large, are impermissible,[17] but I will not discuss this issue further. The more important point is that the actions of the ordinary altruist, being rationally intelligible, are a proper subject of moral evaluation, and it *may* turn out that at least some of them are impermissible. In the case of Pummer's altruist, by contrast, it is just not clear what the motive is, so evaluation is problematic as it must be done partly in the dark.

Having said all that, I will set the issue of motivation aside and assume, for the sake of argument, that the agents' motives *are* something like what Pummer imagines. Does the argument go through then?

3. Action Individuation

3.1. No-conflict scenarios

Consider a version of Red/Green No-Conflict in which, before I drop into the fiery pit, I must sign a document to the effect that I am sacrificing my life of my own volition (perhaps, that's necessary to ensure that the people I save don't face murder charges later). I must then deposit the form in one of two urns – a red urn or a green one. If I place the document into the green urn, an eleventh stranger's finger would be saved. If, instead, I place it into the red urn, the eleventh stranger would lose a finger. The urns are equidistant from me, and there is no other reason to deposit the document in the red urn. I place the document in the red urn. (We can add the additional provisions Pummer makes: I am not motivated by malice, and so on. My motive is just like the motive of the agent in Pummer's original scenario.)

How are we to evaluate my conduct in this case? A natural way to characterize it is to say that I did one highly praiseworthy thing – agreed to sacrifice my life to save ten strangers – and one arguably impermissible thing, namely, I failed to save an eleventh stranger's finger at no additional cost. The impermissible component of my behavior can be described as either an action or an omission. For present purposes, it does not matter which description we pick. The important point is that my overall conduct involves two separate pieces with different moral statuses that must be evaluated separately. What I wish to suggest is that Pummer's Red/Green No-Conflict case is exactly parallel to the modified scenario above, except that the agent performs multiple actions and/or omissions with different moral statuses simultaneously, rather than in short succession. In Red/Green No-Conflict, the praiseworthy and the impermissible thing are, as it were, collated. But this collation is an accident and inessential to the evaluation. We could, if we wish, regard Pummer's scenario as presenting a modification of the urns version, a modification such that the sequence chosen by the agent is executed via the same set of bodily movements, without any temporal separation.

The important point is that in both cases, the agent does one praiseworthy thing – saving 10 strangers at the cost of his life – and one arguably impermissible thing, namely, omitting to save an eleventh stranger's finger at zero cost.[18] In Pummer's version of the case, the two parts are "glued" together as the agent is given an option to perform both simultaneously.

More generally, all Pummer-style cases which involve a Pareto-superior alternative are best seen as composites of two actions (or one action and one omission) with different moral statuses: one praiseworthy (and permissible) and one arguably impermissible, which happen to be performed via the same sequence of bodily motions. Still, one might ask why we should think that this is the best way to see Pummer's scenarios. There are, after all, cases in which a single action has more than one aspect subject to evaluation. Why regard cases such as Red/Green No-Conflict as involving composites of two different *actions*, rather than as single actions with more than one *aspect* subject to evaluation?

The answer is that in Pummer-style cases, the two actions can, in principle, be separated from each other, as I just did. We can redescribe the scenarios so that they involve two actions. This is unlike scenarios involving a single multi-faceted action, as when you fulfil a duty in delivering a piece of unwelcome news, but you do so unkindly or with schadenfreude. In this latter case, we have one single action that cannot be split into two: news delivery, on the one hand, and unkindness or schadenfreude, on the other. The unkindness is unkindness in delivering this particular piece of news, so unkindness and news delivery are organically connected. Not so for Pummer's scenarios.

Of course, we may still wish to evaluate a person's overall conduct on such occasions, but this is a familiar problem which often arises in relation to conduct that involves more than one action (and/or

omission). Did Marge behave well at the party? Well, she was very kind to Sue but mean to Tommy. Was Percy a good president? He reduced the level of corruption significantly but refused to admit mistakes he made. If Marge was praiseworthily kind to Sue and impermissibly mean to Tommy, it would be misleading to say that Marge's conduct at the party was simultaneously praiseworthy and impermissible. Rather, the overall conduct involves some praiseworthy bits and some impermissible bits. Likewise in Percy's case. He may have praiseworthily reduced the level of corruption but impermissibly omitted (or even deliberately refused) to acknowledge any mistakes he made. We may say in such cases that the person being evaluated was a complicated person or that the person's legacy was mixed. There are questions about how someone can be this way. What were this person's values? *Who* was that person? These questions, however, concern the proper evaluation of fragmented human characters. They belong to the domain of moral psychology rather than that of normative ethics. Pummer's no-conflict scenarios are just like these, but the agent is presented with an alternative whereby two actions, or an action and an omission, with different moral statuses are performed via the same set of bodily motions.

3.2. Conflict scenarios

Perhaps Pummer can restrict his argument to cases that involve conflicts of interest between beneficiaries and no Pareto-superior alternative. On the face of it, it may seem that the act of saving X and the omission to save Y and Z are inseparable since it is in choosing to save X that I omit to save Y and Z. But actually, given what Pummer assumes about the importance of numbers and the irrelevance of the identity of beneficiaries, we have one action and one omission here as well. Consider Red/Green Conflict. What the agent does in choosing to save ten strangers plus stranger X rather than ten strangers plus strangers Y and Z is to save eleven strangers and omit to save a twelfth stranger at no additional cost. We can therefore redescribe the case on the model of the urns case above: The

agent signs a form saying he is sacrificing himself of his own volition to save eleven strangers. If he then deposits the form in the green urn rather than the red, a twelfth stranger will be saved. He deposits it in the red urn, and the twelfth stranger dies. The agent's overall conduct here will, once again, involve an action that is praiseworthy and also permissible, and another one that is impermissible and not praiseworthy, rather than a single action that is both impermissible and praiseworthy.

Given Pummer's assumptions, then, it does not matter whether there is a Pareto-superior alternative or not. Conflict cases are, morally speaking, on a par with no-conflict cases, since only the numbers matter.[19] We can imagine a version in which the agent has no idea whether the available options, beyond saving ten strangers, involve saving X rather than X and Y, or saving X rather than Y and Z. All she knows is that she can save either one more person or two more people. Whether she thinks that the additional people are X, on the one hand, and Y and Z, on the other, or she thinks they are X, on the one hand, and X and Y, on the other, is of no consequence when it comes to morally evaluating her conduct. Given Pummer's assumptions, this case too would be evaluatively identical to his.

Since in Pummer's scenarios – as well as, we may add, in scenario Benighted discussed at the end of the previous section – the two actions or the action and the omission with different moral statuses are executed indivisibly, via the same set of bodily motions, there may be a temptation to speak of them as one action, but doing so is loose talk. Such talk is fine in ordinary contexts, but in the context of inquiry relevant here – namely, a context in which our aim is to understand morality – we must endeavor to be more precise.

I conclude that the types of cases Pummer offers are not, in fact, cases in which an agent does something that is overall praiseworthy yet impermissible. Does it follow from here that no action can be both praiseworthy and impermissible? This is the question I turn to now.

4. Truly Impermissible Altruism

I wish to suggest that there *are* types of altruistic acts which may be both praiseworthy and impermissible, but they create additional difficulties instead of saving the day. I have in mind acts of altruism that are too extreme to be permissible. A person who donates a kidney and saves a stranger's life is doing something highly admirable and praiseworthy. Suppose, however, that this person is extremely altruistic and attempts to donate *all* of her organs: her second kidney, her heart, her liver, etc.[20] Doing so with the intention to save several lives is praiseworthy, but it is impermissible.[21] Importantly, this is not a case in which a person intends to sacrifice her life arbitrarily. An arbitrary sacrifice of one's life, altruistic though it may be, would be impermissible but not praiseworthy (and would be arguably blameworthy). For instance, suppose Bill is so in love with Bella that he accepts a significant harm to his leg in order to prevent a minor inconvenience to Bella. This action is, arguably, not praiseworthy and may in fact be blameworthy. But the extreme altruist case is not of this sort. All organs can save lives, and several lives can be saved if the surgery is performed. Still, donating all of one's organs is impermissible on widely shared assumptions. Thus, suppose you are the administrator of the hospital where the patient wants to have this surgery. One of your surgeons comes to you and asks for permission to do this surgery. Do you approve? Surely not. But if it's permissible to do the surgery, and it would save 5 lives, wouldn't you be obligated to approve it?

One can argue here that perhaps, as administrator, you would not approve the surgery, because you may then be legally liable and possibly subject to criminal prosecution. But the question here is what

your moral duties are. If a law is unjust – for instance, if you are by law obligated to commit genocide – you would be morally required to disregard the law. By contrast, it does not seem that you would be morally required to disregard a law that requires you to stop an organ harvesting procedure or permits you to charge a surgeon performing that type of operation with a crime.

One can argue also that donating all of one's organs is permissible and perhaps even required on certain versions of utilitarianism. [22] There are three points I would make in response. First, most discussions in ethics – of rights, resources allocation, and so on – take seriously common-sense intuitions, which happen to be inconsistent with extreme versions of utilitarianism. Second, since the organ harvesting example is meant to suggest a new difficulty for Pummer, as I will explain shortly, it is worth noting that Pummer cannot appeal to utilitarianism in order to argue that the action in this case is permissible. His project in the target paper has the ultimate aim of defending a conditional obligation to pursue altruistic aims effectively if we pursue altruistic aims at all and to mount such a defense starting from common-sense moral premises. If Pummer assumed utilitarianism, then he would be trivializing his project since it is trivial to get the obligation to act as an effective altruist would from utilitarianism. Third and finally, it is possible to give an example along the lines I have in mind tailored to utilitarian sensibilities. We can imagine an agent who makes a sacrifice for another such that the cost to the agent exceeds significantly the benefit to the other. The gap is not so big as to render the action blameworthy, as in the Bella case above, but it is big enough to make it overall impermissible though still praiseworthy. But let us go back to the organ harvesting case.

The case presents difficulties for Pummer. Either Pummer has not given us resources to explain what makes such actions impermissible, or else there is a problem with his definition of praiseworthiness. Consider the question of impermissibility first. How, following Pummer, would we explain the

impermissibility in the extreme altruism case? On Pummer's view, an action is impermissible when the balance of requiring reasons goes against it and there is no sufficient justification to prevent the requiring reasons from making the action impermissible.[23] Pummer defines "requiring reason" thus:

> I will use 'a requiring reason' to refer to an aspect of an act (in a set of alternatives) that tends to make this act required. Even when such a reason does not succeed in making an act required it nonetheless pulls in that direction, contributing to it being the case that the act is required. That an act involves keeping a promise, for instance, tends to make this act required.[24]

The extreme altruist is responding to a requiring reason – to save a person's life. If there are no other requiring reasons, then the balance of reasons favors saving the stranger, and we have no explanation of impermissibility.

Perhaps, we can say, on behalf of Pummer, that there is *also* a requiring reason to preserve one's own life, and that this reason is stronger here than the reasons to save the strangers. But if we tried to handle the impermissibility of extreme altruism by simply postulating a requiring reason to save one's life, that would create a problem for Pummer's account of praiseworthiness. Pummer writes:

> **Praiseworthiness**: Act A is praiseworthy to some extent when there is a requiring reason to do A and a justification not to. Act A is praiseworthy to a greater extent, the stronger the requiring reasons to do A, and the stronger the justifications not to.[25]

The proposal is meant to be compatible with a range of views about the precise nature of the elements identified in the definition – requiring and justifying reasons – and of the role they play in praiseworthiness. (Are only cost-based justifications relevant to praiseworthiness, or are there other

possible justifications? Does an action have to be at least somewhat costly to you to be praiseworthy? And so on.)

The problem I see here, however, is not in the details but in the overall framework. If we postulate a requiring reason to save one's life, it will follow, on Pummer's definition of praiseworthiness, that a person who acts so as to save her life when she has a justification not to is doing something morally praiseworthy. For instance, if a cancer patient accepted treatment – assuming she has a right to refuse on grounds of self-ownership – she would be acting in a morally praiseworthy way. That is an unfortunate consequence, and it suggests that the framework needs substantial revision and not simply fine-tuning. So either the impermissibility of donating all of one's organs remains unexplained, or the account of praiseworthiness is found lacking.[26]

I have argued elsewhere that the permissibility of altruism in general, which we usually take for granted but which actually raises a moral question, has to do with a special relationship we have to ourselves: we are permitted to impose costs on ourselves in excess of what can be objectively required of us. (Altruism's *virtuous* character resides in our lack of general inclination to be altruistic and the social benefits associated with altruism.) Objective morality, however, places constraints on the costs we can permissibly impose on ourselves. Going too far in the direction of altruism often remains praiseworthy but can become impermissible, as it does in the multiple organ donor case. Importantly, the moral constraints on altruism are constraints on a virtuous inclination. They function quite differently from constraints on selfish inclinations. Violating constraints on selfish inclinations is not praiseworthy, but violating constraints on virtuous inclinations may well be, sometimes even past the point of permissibility. But this is a topic for another occasion.

* I would like to thank Mike Huemer and Jon Tresan for stimulating discussions of the material covered in this article as well as two anonymous referees for *Ethics* and the journal editor, Doug Portmore, for written feedback. Thanks are due also to CU-Boulder graduate student Kyle York who, after bringing up Pummer's idea in my seminar and hearing my response to it, reacted with enthusiasm and suggested that I write a formal reply.

[1] Theron Pummer, "Impermissible yet Praiseworthy," *Ethics* 131(2021)): 697–726, 710.

[2] Ibid., 713.

[3] Ibid., 719.

[4] Pummer writes in this regard, "Although heroic life-saving acts are paradigmatically praiseworthy, they can fail to be overall or even pro tanto praiseworthy—and can be pro tanto or even overall blameworthy instead—if done for sufficiently bad motivating reasons. It is plausible that you would be overall blameworthy for saving the lesser number in *Costly No-Conflict* if you acted out of a combination of concern for the one stranger and unjustifiable hate for the other.," Ibid., 712.

[5] Ibid.

[6] Ibid., 718.

[7] Ibid., 714.

[8] Ibid., 725.

[9] John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6 (1977): 293–316, quoted at Ibid., 716n47. In the same footnote, Pummer references a number of other authors who have defended views along those lines, including Anscombe, Munoz-Dardé, Broome, Bradley, Parfit, and Kamm. As a side note, I suspect a morally better procedure would involve randomization but in such a way that numbers are taken into account, for instance, in a choice of one or two, the group of two has a two-thirds chance of being picked.

[10] Perhaps, one can argue that even if Taurek style arguments are reasonable and the agent is persuaded by them and acts on that persuasion, if Taurek is wrong as a matter of fact, then the agent is acting impermissibly. But I think this line of response is not sufficiently sensitive to the agent's epistemic state. If the response succeeded, it would follow that many of us are likely acting impermissibly much of the time now, since in all likelihood, our future descendants will discover that a significant proportion of our reasonable moral beliefs are false.

[11] He does not make an attempt to refute the arguments of Taurek and others, but he suggests briefly that that there is more requiring reason to save the larger number than the smaller, and no sufficiently good justification to save the smaller. For details regarding Pummer's notion of "requiring reason," see Section 4.

[12] Ibid., 712–13.

[13] Thanks are due to an anonymous referee for pressing this line of argument.

[14] Cf. footnote 4 above.

[15] Many thanks to an anonymous referee for this example.

[16] Ibid., 713, footnote 37.

[17] That is particularly true when one is making smaller contributions. If you intend to donate a billion dollars, then there is a strong reason to spend some time thinking things over, though even then, it is not clear that you'd be doing anything *impermissible* by choosing charities that are significantly sub-optimal. It is worth mentioning also that Will McAskill explicitly says that effective altruism is a "project," not a "set of normative claims." See https://forum.effectivealtruism.org/posts/9wYa8BqSTMcx9j2tK/defining-effective-altruism.

[18] In some of the cases Pummer considers, the omission involves a minor cost compared to the good achieved, for instance, muddying your shoes to save a person's life. This detail does not matter for present purposes.

[19] Of course, one can dispute Pummer's assumptions and argue that it matters greatly whether there is a Pareto-superior alternative or not. But this line of defense against the problematic action individuation objection is not available to Pummer. What is worse, if Taurek-style arguments succeed, and we come to accept that while saving X is not as good as saving X and Y, saving X is just as good as saving Y and Z (perhaps, sacrificing an innocent life is an infinite bad), the argument can be saved from the action individuation objections – at least in conflict cases – but it would fail for a different reason: if the numbers don't matter, then Pummer's agent would not be doing anything impermissible in going for the smaller number in conflict cases.

[20] A case like this is portrayed in the Television Show *House*. See David Shore (writer), Sara Hess (writer), and Greg Yaitanes (director), "Charity Case", *House, M.D.*, season 8, episode 3, aired Oct. 17, 2011 (NBCUniversal Television Distribution).

[21] There is another complication here: It seems permissible to sacrifice one's life in emergency situations but not in non-emergency situations, for instance, you can give your life vest to your companion (thereby saving a single person), but you cannot permissibly donate all of your organs and save several lives. A full discussion of the reasons for this, however, is beyond the scope of this paper.

[22] Thanks to an anonymous referee for pressing this point.

[23] Pummer, "Impermissible yet Praiseworthy," 704–705.

[24] Ibid., 704.

[25] Ibid., 706.

[26] One may ask here whether the unintelligible motives objection I pressed against Pummer's cases cannot be pressed against the extreme altruist case I describe. I thank an anonymous referee for this point. My response is that while the motivation of an extreme altruist may be difficult to relate to, it is not unintelligible, and there is no problem with morally evaluating the action. What makes extreme altruism difficult to relate to is that most of us are not inclined to be very altruistic. But we know exactly what the agent's motive is – to save lives – so we can pass a moral judgment. Compare: it is difficult to relate to a thrill-seeker's desire to climb a skyscraper without any protection or equipment. But we know what the motive is – a desire for thrill. In both cases, we may not be able to identify with the desires and motives of the agents, because the agents are at the end of the relevant spectrum while most of us are not. The problem with Pummer's agents is quite different. It is not clear what the motives of those agents are, what story the agents tell themselves about their own actions, or what they would say in response if anyone were to ask them what their reasons were. We have to imagine that the agents are simultaneously motivated by what seem to be psychologically incompatible tendencies: to help some strangers at a great personal cost, on the one hand, and to remain utterly indifferent to the fate of similarly positioned strangers that can be helped at no additional cost, on the other hand. Are the agents altruistic or callously indifferent?