Impermissible Yet Praiseworthy*

Theron Pummer

tgp4@st-andrews.ac.uk

Penultimate version—final version forthcoming in *Ethics* (July 2021)

Abstract:

It is commonly held that unexcused impermissible acts are necessarily blameworthy, not praiseworthy. I argue that unexcused impermissible acts can not only be pro tanto praiseworthy, but overall praiseworthy—and even more so than permissible alternatives. For example, there are cases in which it is impermissible to at great cost to yourself rescue fewer rather than more strangers, yet overall praiseworthy, and more so than permissibly rescuing no one. I develop a general framework illuminating how praiseworthiness can so radically come apart from deontic status.

* For very helpful written comments on earlier drafts, I am grateful to Brian Berkey, Joseph Bowen, Matthew Clark, Roger Crisp, Johann Frick, Peter Graham, Joe Horton, Shelly Kagan, Frances Kamm, Per Milam, Kian Mintz-Woo, Daniel Muñoz, Julia Nefsky, Douglas Portmore, Justin Snedegar, Travis Timmerman, Erik Zhang, and especially the editors and two referees at *Ethics*. Many thanks also to audiences at the Princeton DeCamp Seminar, the Princeton Workshop in Normative Philosophy, the Rutgers Value Theory Reading Group, and the Yale Moral Philosophy Working Group.

Impermissible Yet Praiseworthy

Suppose you are morally required to adopt a vegan diet, but you adopt a lacto-vegetarian diet instead. Although what you do is impermissible, blaming you for not going all the way to veganism could be counterproductive. Perhaps the effects of blaming you are even bad enough that we ought not to do so. But the effects of blaming you do not settle whether you are blame*worthy*, in the sense of it being appropriate to have a negative reactive attitude toward you for your conduct, or in the sense of your being deserving of discredit for your conduct.[1] To be praiseworthy or blameworthy is to be worthy of praise or blame, in some such sense, for one's conduct.[2]

---

[1] I am thus here setting aside consequentialist accounts of blameworthiness. See, for instance: Richard Arneson, "The Smart Theory of Moral Responsibility and Desert," in *Desert and Justice*, edited by Serena Olsaretti (Oxford: Oxford University Press, 2003): 233–58. It also seems that you can be worthy of blame even when we lack the standing to blame you. On the latter, see, for instance: R. Jay Wallace, "Hypocrisy, Moral Address and the Equal Standing of Persons," *Philosophy and Public Affairs* 38 (2010): 307–341; Macalester Bell, "The Standing to Blame: A Critique," in *Blame: Its Nature and Norms*, edited by Justin Coates and Neal Tognazzini (New York: Oxford University Press, 2013), 263–281.

[2] For a sample of relevant literature, see: Michael Zimmerman, *An Essay on Moral Responsibility* (Totowa: Rowman and Littlefield, 1988); Holly Smith, "Varieties of Moral Worth and Moral Credit," *Ethics* 101 (1991): 279–303; R. Jay Wallace, *Responsibility and the Moral Sentiments* (Cambridge: Harvard University Press, 1994); Nomy Arpaly, *Unprincipled Virtue: An Inquiry Into Moral Agency* (New York: Oxford University Press, 2003); T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge: Harvard University Press, 2008); Julia Markovits, "Acting for the Right Reasons," *Philosophical Review* 119 (2010): 201–42; Michelle Mason, "Blame: Taking It Seriously," *Philosophy and Phenomenological Research* 83 (2011): 473–81; Amy Massoud, "Moral Worth and Supererogation," *Ethics* 126 (2016): 690–710; Dana Nelkin, "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness," *Noûs* 50 (2016): 356–78; David Shoemaker, "Response-Dependent Responsibility; or, A Funny Thing Happened on the Way to Blame," *Philosophical Review* 126 (2017): 481–527;

It is commonly held that unexcused impermissible acts are necessarily blameworthy, not praiseworthy.[3] On some views, what it *means* for an act to be impermissible is that it is blameworthy, absent excusing conditions. On other views, all unexcused impermissible acts are blameworthy, even if what it means for an act to be impermissible is not that it is blameworthy, but, for example, that it *mustn't be done*.[4]

In this paper, I defend the following claim:

> ***Impermissible Yet Praiseworthy*:** There are situations in which it is overall praiseworthy to do an act *A*, and more so than it is to do a permissible and entirely blameless alternative *B*, even though *A* is impermissible, and all relevant responsibility conditions are met (so that *A* is unexcused).

It may, for example, be impermissible to at great cost to yourself rescue fewer rather than more strangers, yet overall praiseworthy, and more so than permissibly rescuing none. No one to my knowledge has defended *Impermissible Yet Praiseworthy*.[5] Not only is this claim interesting

---

Elinor Mason, *Ways to be Blameworthy: Rightness, Wrongness, and Responsibility* (Oxford: Oxford University Press, 2019); Pamela Hieronymi, *Freedom, Resentment, and the Metaphysics of Morals* (Princeton: Princeton University Press, 2020).

[3] For example, see: Allan Gibbard, *Wise Choices, Apt Feelings* (Oxford: Clarendon Press, 1990), chapter 3; John Skorupski, *The Domain of Reasons* (Oxford: Oxford University Press, 2010), part III; Stephen Darwall, *Morality, Authority, and Law: Essays in Second-Personal Ethics I* (Oxford: Oxford University Press, 2013), 21; Brian McElwee, "Demandingness Objections in Ethics," *Philosophical Quarterly* 67 (2017): 84–105.

[4] See Derek Parfit, *On What Matters, Vol. 1* (Oxford: Oxford University Press, 2011), chapter 7.

[5] The impermissible yet praiseworthy is in some sense the mirror image of the suberogatory, that is, the permissible yet blameworthy. See Julia Driver, "The Suberogatory," *Australasian Journal of Philosophy* 70 (1992): 286–295. In "Morally Permissible Moral Mistakes," *Ethics* 126 (2016): 366–393, Elizabeth Harman recognizes moral

and surprising, but it is also of potentially great dialectical significance. If it is correct, then we cannot conclude that an act is permissible from the fact that it is praiseworthy, or from the fact that it is more praiseworthy than a permissible alternative.[6]

In section I, I present cases in which praiseworthiness appears to attach to impermissible acts. These impermissible acts seem pro tanto praiseworthy, and some might even be overall praiseworthy. However, since these impermissible acts are not *more* praiseworthy overall than permissible alternatives, the cases presented in section I are not of

---

mistakes that are praiseworthy, but, crucially, these are morally *permissible* moral mistakes. In "Praise, Blame, Obligation, and Beyond: Toward a Comprehensive Framework for the Classical Conception of Supererogation and Kin," *Deontic Logic in Computer Science*, edited by Ron van der Meyden and Leendart van der Torre, Lecture Notes in Artificial Intelligence, 5076 (Berlin: Springer Verlag, 2008), Paul McNamara offers a case (15) in which an act is impermissible yet praiseworthy, but crucially the act is excused in virtue of non-culpable ignorance. In "Distinctive Duress," *Philosophical Studies* 177 (2020): 1007–1026, Craig Agule defends the existence of a kind of praiseworthy impermissibility, though his view and argument are quite different from mine (for example, unlike Agule's, my view crucially relies on a distinction between being pro tanto praiseworthy and being overall praiseworthy). Arguably the closest predecessors of my work here can be found in: Ulla Wessels, "Beyond the Call of Duty: The Structure of a Moral Region," *Royal Institute of Philosophy Supplement* 77 (2015): 87–104; and Daniel Muñoz, "Three Paradoxes of Supererogation," *Noûs* (2020) https://doi.org/10.1111/nous.12326. Using cases similar to those explored in section III, Wessels and Muñoz argue that impermissible acts can be *better* than permissible alternatives. However, neither defends my claims about praiseworthiness.

[6] Insofar as a supererogatory act is a non-required act that is praiseworthy, or more praiseworthy than a merely permissible alternative, I can be read as arguing that some supererogatory acts are impermissible (though it is not my aim here to defend a view of the supererogatory per se). For relevant discussion, see: Joseph Raz, "Permissions and Supererogation," *American Philosophical Quarterly* 12 (1975): 161–168, 164; David Heyd, *Supererogation: Its Status in Ethical Theory* (Cambridge: Cambridge University Press, 1982); Paul McNamara, "Supererogation, Inside and Out: Toward an Adequate Scheme for Common Sense Morality," *Oxford Studies in Normative Ethics* 1 (2011): 202–235; Alfred Archer, "Are Acts of Supererogation Always Praiseworthy?" *Theoria* 82 (2016): 238–55; Claire Benn, "Supererogation, Optionality, and Cost," *Philosophical Studies* 175 (2018): 2399–2417.

the form picked out by *Impermissible Yet Praiseworthy*. In section II, I offer a general framework illuminating how praiseworthiness can come apart from deontic status in this more radical way. The framework showing this conceptual possibility consists of a distinction between *requiring reasons* and *justifications*, and of views about how they relate to praiseworthiness and blameworthiness. In section III, I move beyond mere conceptual possibility. I argue there plausibly *are* cases of the form picked out by *Impermissible Yet Praiseworthy*, in which you can at great cost to yourself rescue more or fewer people. The intended purpose of this section is to provide, by way of illustration, a more robust defense of the central claim that there are *some* cases of this general form (even if none of the particular ones I discuss are). In section IV, I suggest—as a supplement to the framework presented in section II—that there may be limits on when impermissible acts can be overall praiseworthy. The holding of such limits will constrain the form that cases picked out by *Impermissible Yet Praiseworthy* can take. In section V, I conclude with a brief discussion of some potential practical implications.

It is worth bearing in mind that my focus in this paper is *moral* normativity, that is, moral requirement, moral permissibility, and so forth (as opposed to rational requirement, rational permissibility, and so forth). I will often omit the term "moral" for brevity.

I. Praiseworthiness and Impermissibility

Many acts that are required are not praiseworthy.[7] You are not praiseworthy for using turn signals while driving, though you are required to do so. Many acts that "go beyond" what is

[7] Speaking of *acts* as being praiseworthy is only shorthand; ultimately it is agents who are praiseworthy. My primary focus here is the praiseworthiness of agents *for* performing certain acts for certain motivating reasons, as opposed to their praiseworthiness for having or displaying certain dispositions (see, for instance: Phillip

required are praiseworthy. You are praiseworthy for risking your life to save a stranger trapped in a burning building, though you are not required to do so. But sometimes praiseworthiness attaches to other sorts of acts. Some acts that are required are praiseworthy. You are praiseworthy for missing a nonrefundable international flight to take a seriously injured cyclist to the hospital, and required to do so.[8] And some acts that fall short of what is required—that is, that are impermissible—are praiseworthy, at least to some extent. Here are three putative examples.

Suppose you are out for dinner at a restaurant. There are just three items on the menu: a meat dish, a vegetarian dish, and a vegan dish. You most want the meat dish, the vegetarian dish is a distant second, and the vegan dish you would really rather do without. We can suppose that ordering the vegan dish is both required and praiseworthy. It might then seem that making an effort but still falling short—ordering the vegetarian dish—is at least to some extent praiseworthy. Much of what would make it praiseworthy to order the vegan dish would appear to be present in such a case of ordering the vegetarian dish. It is at least arguable that the impermissibility of not fully meeting morality's demand does not entirely cancel your praiseworthiness. Intuitively you remain worthy of at least *some* credit.

Here is another example: Suppose that you are required to go on strike. But, out of concern for your students, you decide to cross the picket line to go in and teach. Given the circumstances, teaching comes at a significant cost to you. It would appear we can fill in the

Montague, "Acts, Agents, and Supererogation," *American Philosophical Quarterly* 26 (1989): 101–11; Arpaly 2003, chapter 3; Markovits 2010; Massoud 2016).

[8] We can suppose you are the only person around to take the cyclist to the hospital. See: Massoud 2016 (692). For earlier discussion, see: J. O. Urmson, "Saints and Heroes," in *Essays in Moral Philosophy*, edited by A. I. Melden (Seattle: University of Washington Press, 1958), 198–216.

particular details of the case so that it is plausible that teaching is at least somewhat praiseworthy, even though it is impermissible.

And a third example: You rob a bank at significant cost to yourself in order to give the money to the homeless stranger you encountered on your morning walk.[9] Even if what you do is impermissible, and on the whole blameworthy, it still seems you are at least somewhat praiseworthy for your conduct. While it may be odd to say your act is a step in the direction of what morality demands, you do pay a significant cost in responding to a moral reason to help the stranger (your motivating reason for robbing the bank coincides with this normative reason—you did not rob the bank just for the thrill of it).[10] It is arguably the fact that you put in significant effort or paid significant costs in responding to moral reasons that explains why you are at least somewhat praiseworthy in the two earlier examples as well. Though you would have been more responsive had you ordered the vegan dish instead of ordering the merely vegetarian dish, you are nonetheless in each case significantly responsive to reasons.[11]

[9] For related examples, see: Michael Zimmerman, "A Plea for Ambivalence," *Metaphilosophy* 24 (1993): 382–89; Markovits 2010 (240–41).

[10] I here build upon Massoud 2016 and Apraly 2003 (chapter 3).

[11] Raz 1975 (165–66) can be read as claiming that an act cannot be *at all* praiseworthy unless there is on balance most reason to do it, as opposed to merely a reason to do it. This would leave us unable to accommodate the "ambivalent" or "mixed" response it seems appropriate to have toward you for robbing the bank to help the homeless stranger (there is most reason to refrain from robbing the bank). We cannot satisfactorily accommodate this ambivalence by suggesting, as Raz 1975 (166) does, that we are simultaneously admiring one aspect of your character while regretting another. For we can suppose that your motivation to help the homeless stranger, though genuine, is out of character (see, for instance, Arpaly 2003 (94–5)). The most plausible explanation of the ambivalence is that you are praiseworthy for an aspect of your conduct, and blameworthy for another aspect of your conduct. We can instead read Raz as offering the more plausible claim that an act cannot be *overall*

These examples suggest that we need to distinguish between being *pro tanto* praiseworthy and being *overall* praiseworthy. Let us say an agent is pro tanto praiseworthy for performing a token act for a token motivating reason when they are praiseworthy for some aspect of performing this act for this motivating reason (and likewise for being pro tanto blameworthy). An agent can at once be pro tanto praiseworthy for performing an act for a motivating reason—for example, robbing the bank to help the homeless stranger—and be pro tanto blameworthy for performing this same act for this same motivating reason. An agent is overall praiseworthy for performing an act for a motivating reason when, taking into account all of the praiseworthy and blameworthy aspects of doing so, they are on balance praiseworthy for doing so (and likewise for being overall blameworthy). No agent can at once be overall praiseworthy for performing an act for a motivating reason and be overall blameworthy for performing this same act for this same motivating reason.[12]

What exactly it is to be overall praiseworthy will depend on what it is to be worthy of praise or blame for one's conduct.[13] We might, for instance, conceive of the latter as being deserving of credit or discredit. In the bank robbery case, you deserve some credit for helping

---

praiseworthy unless there is most reason to do it; the arguments for *Impermissible Yet Praiseworthy* offered below serve as arguments against this more plausible claim.

[12] Even if there are genuine moral dilemmas in which act *A* is required, act *B* is required, and you cannot do both, neither *A* nor *B* can at once be overall praiseworthy *and* overall blameworthy. In certain asymmetric dilemmas, it may be that *A* is overall praiseworthy and *B* is not. However, these would not be cases of the sort picked out by *Impermissible Yet Praiseworthy*, in which impermissible act *A* is more praiseworthy overall than permissible alternative *B*. In a genuine dilemma, no alternative is permissible. For putative examples, see: Michael Walzer, "Political Action: The Problem of Dirty Hands," *Philosophy and Public Affairs* 2 (1973): 160–180; Thomas Nagel, "War and Massacre," reprinted in his *Mortal Questions* (Cambridge: Cambridge University Press, 1979), 53–74.

[13] For some of the various possible accounts, see footnote 2.

the homeless stranger, but you deserve even more discredit for stealing money and harming innocent bystanders. Overall you deserve discredit for your conduct, which corresponds to the difference between the credit deserved for the good aspects of your conduct and the discredit deserved for the bad aspects of your conduct.[14] We might instead conceive of being worthy of praise or blame as being the appropriate target of a positive or negative reactive attitude. It is appropriate to have a positive reactive attitude toward you for helping the homeless stranger, but it is also appropriate to have an even stronger negative reactive attitude toward you for stealing money and harming innocent bystanders. The appropriateness of such an ambivalent response may be compatible with the appropriateness of having a single higher-order attitude toward you for your conduct as a whole (overall admiration, perhaps). But it may be that there is no single attitude it is appropriate to have toward your conduct as a whole, which corresponds to a "difference" or "average" between the reactive attitudes it is appropriate to have toward each aspect of your conduct. Even in the absence of any single overall appropriate attitude, we could say you are overall blameworthy as shorthand for the claim that the appropriate negative reactive attitude is relevantly stronger than the appropriate positive reactive attitude.[15] I intend to remain neutral on these different possible conceptions of being worthy of praise and blame, and on the corresponding differences in how exactly to conceive of being overall praiseworthy. In what follows, I will occasionally write "praiseworthy" or "more praiseworthy"

[14] The idea of overall credit deserved fits well with (though is not entailed by) Zimmerman's 1988 (chapter 3) "ledger" metaphor.

[15] I take it that it would still be an interesting claim that there are situations in which the positive reactive attitude it is appropriate to have toward you for some aspect of doing *A* is relevantly stronger than the negative reactive attitude it is appropriate to have toward you for some aspect of doing *A* (*A* is overall praiseworthy), even though *A* is impermissible and unexcused.

unaccompanied by either "pro tanto" or "overall." The "overall" reading is intended only when explicitly indicated.

It seems fairly uncontroversial that some acts are impermissible yet pro tanto praiseworthy. The vegetarianism, teaching, and bank robbery examples suggest so. More controversial is the claim that some acts are impermissible yet overall praiseworthy. The bank robbery example, at least, is not a case of this sort, as robbing the bank to help the homeless stranger is overall blameworthy. Are *all* unexcused impermissible acts overall blameworthy?

I do not believe so. The aforementioned views which posit a necessary connection between unexcused impermissibility and blameworthiness do not support the claim that unexcused impermissible acts are necessarily overall blameworthy.[16] To the extent that these views are plausible, they may support the weaker claim that all unexcused impermissible acts are necessarily pro tanto blameworthy, but they are orthogonal to the neglected question of how an impermissible act's blameworthy aspects balance against its praiseworthy aspects. Robbing the bank is pro tanto blameworthy, perhaps partly in virtue of being impermissible, perhaps partly in virtue of (some of) the facts that make it impermissible. It is also pro tanto praiseworthy. But this act's blameworthy aspect more than amply counterbalances its praiseworthy aspect. It is overall blameworthy.

If the scales can tip one way, can they not tip the other? It seems they can. Consider again the vegetarianism example. Even if ordering the vegan dish is what is required, it might be more praiseworthy to order the vegetarian dish (rather than order the meat dish) than it is blameworthy to order the vegetarian dish (rather than order the vegan dish). It could in this way be overall praiseworthy to order the vegetarian dish, even if impermissible. For another putative example, suppose you are an avid traveler who each year takes dozens of flights for

[16] See footnote 3 for literature on these views.

leisure. You have a tradition of booking all your flights for the upcoming year on New Year's Eve. But you have recently become concerned about climate change, and this year book only one flight. Now suppose that the reason to reduce our carbon emissions is strong enough that each of us is required to eliminate air travel for leisure. It is therefore impermissible to book even just the one flight for next year. At the same time, perhaps this is overall praiseworthy. More generally, it seems that acts can be praiseworthy even if they are not *most* responsive to the reasons present—you can be worthy of "partial credit" for being only somewhat responsive, particularly if this is costly or difficult for you. Perhaps sometimes this partial credit is relevantly greater than the discredit deserved for falling short of what is required.

One might object that responding to reasons at cost to yourself is insufficient for being pro tanto praiseworthy. One might claim that, to be praiseworthy at all, you must do as much as required, if not more (what is required may be less than what there is most reason to do). If this claim is correct, the impermissible acts in the cases offered so far are not even pro tanto praiseworthy, let alone overall praiseworthy. In reply, we could note that it would seem a stingy view of praiseworthiness that implied you could sacrifice and sacrifice in response to moral reasons and deserve *no* praise whatsoever for this if you still fall a bit short. But I need not insist on this reply. For in sections II and III I turn to acts that go beyond what is required, yet are at the same time impermissible.

In this section, I suggested that praiseworthiness can attach to impermissible acts. I presented examples in which impermissible acts seem pro tanto praiseworthy. Some of them might even involve acts that are impermissible and yet overall praiseworthy. But even if these are cases in which impermissible acts are overall praiseworthy, they are relatively tame ones: the impermissible acts in question are not *more* praiseworthy overall than permissible alternatives. The cases discussed in this section are therefore not of the form picked out by *Impermissible Yet Praiseworthy*, that is, situations in which it is overall praiseworthy to do an

act *A*, and more so than it is to do a permissible and entirely blameless alternative *B* (*B* is not pro tanto blameworthy), even though *A* is impermissible and unexcused. In the following section, I offer a general framework illuminating how praiseworthiness can come apart from deontic status in this more radical way.

II. Impermissible Yet Praiseworthy: A Framework

In this section, I introduce a general framework illuminating how there can be cases of the form picked out by *Impermissible Yet Praiseworthy*. It consists of a familiar distinction between *requiring reasons* and *justifications*, and of views about how they relate to praiseworthiness and blameworthiness. I do not claim that this is the only framework within which cases of the form picked out by *Impermissible Yet Praiseworthy* can be generated and explained, but it strikes me as a fairly natural and useful one.

First, requiring reasons. 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. But you may sometimes also have an even stronger requiring reason *not* to perform this act.[17] The requiring reason not to let an innocent stranger die can be stronger than the requiring reason to keep your promise. I will say that there is "more requiring reason" (mass noun) to do act *B* than there is to do an alternative *A* when the respective requiring

[17] Requiring reasons can thus be likened to Rossian prima facie (pro tanto) duties, which play a contributory role in the determination of all things considered duties. See: W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930). Ross himself apologizes for the misleading phrase "prima facie" (20), as he intends to be referring not to appearances but to actual features that play a pro tanto or contributory role.

reasons to do *B* and *A* on balance make it the case that act *A* is impermissible, absent a sufficiently strong justification to do *A*. I will likewise take it that *B* is required whenever there is most requiring reason to do *B*, absent a sufficiently strong justification not to do *B*.[18]

Second, justifications. I will here treat justifications as considerations that justify without requiring.[19] As I construe them, justifications tend to make acts permissible only indirectly, by tending to *prevent* requiring reasons from making acts impermissible. A justification is sufficiently strong when it successfully prevents the balance of requiring reasons from making an act impermissible. For instance, if saving the life of a stranger unavoidably involved the painful loss of your legs, I take it there would be a strong enough justification to

[18] I will here be using "more requiring reason" in a way according to which if there is more morally requiring reason to do *B* than there is to do *A*, then *B* is morally better than *A*. For an important discussion of how there can be more moral reason to do *B* than there is to do *A* even though *A* is morally better than *B*, see: Douglas Portmore, *Opting for the Best: Oughts and Options* (Oxford: Oxford University Press, 2019). On Portmore's view, there is more reason to do *B* than there is to do *A* if the best "maximal option" that entails doing *B* is better than the best maximal option that entails doing *A*. But it can be that the best maximal option that entails doing *B* is better than the best maximal option that entails doing *A* even when *A* is itself better than *B* (see Portmore 2019, section 6.4, especially 210–13).

[19] On purely justifying considerations, see: Douglas Portmore, "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11 (2008): 369–88; Thomas Hurka and Esther Shubert, "Permissions to Do Less Than the Best: A Moving Band," *Oxford Studies in Normative Ethics* 2 (2012): 1–27; Joshua Gert, "The Distinction Between Justifying and Requiring: Nothing to Fear," in *Weighing Reasons*, eds. Errol Lord and Barry Maguire (New York: Oxford University Press, 2016), 157–172; Seth Lazar, "Accommodating Options," *Pacific Philosophical Quarterly* 100 (2019): 233–255. Purely *morally* justifying considerations are widely endorsed, and recall that my focus here is moral normativity (by "a requiring reason" I mean "a morally requiring reason," and so on).

refrain from saving the stranger.[20] You would be permitted but not required to refrain from saving them.[21] If instead saving the stranger's life involved nothing more than the loss of your shoes, there might be a weak justification to refrain from saving the stranger's life, but it would not be sufficiently strong in the face of the strong requiring reason to save the stranger's life. You would be required to save them. When an act is supported by the balance of requiring reasons—that is, when there is no more requiring reason to do anything else—no justification is needed for the act to be permissible. (Or, if you prefer, a justification of zero strength would be sufficient.) Finally, the fact that there is a justification to do act *A* strong enough to prevent the fact that there is more requiring reason to do *B* than there is to do *A* from making *A* impermissible is not always enough to make *A* permissible. After all, such a justification to do *A* might not be strong enough to prevent the fact that there is more requiring reason to do *C* than there is to do *A* from making *A* impermissible.[22]

[20] On agent-relative prerogatives, see: Samuel Scheffler, *The Rejection of Consequentialism* (New York: Oxford University Press, 1982); Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989); Hurka and Shubert 2012; Lazar 2019.

[21] This is not to deny that you are *prudentially* required to keep your legs. Nor is it to deny that there is *a* morally requiring reason to keep your legs (even if the loss of your legs is costly for no one but you). The picture is instead that there is significantly more morally requiring reason to save the stranger at the cost of your legs than there is to do nothing, but that you are for all that not morally required to incur this cost, since there is a sufficiently strong moral justification to keep your legs.

[22] For example, suppose you can do nothing, save a stranger's index fingers at the cost of your pinky finger, or save the same stranger's index fingers along with another stranger's life at the cost of your pinky finger. It seems the fact that doing nothing avoids the loss of your pinky is a sufficiently strong justification to prevent the fact that there is more requiring reason to save the stranger's index fingers than there is to do nothing from making it impermissible to do nothing. But, as the requiring reason to save a life is considerably stronger than the requiring reason to save fingers, it seems that the fact that doing nothing avoids the loss of your pinky is not a sufficiently

The remainder of the framework consists of two connections between requiring reasons, justifications, praiseworthiness, and blameworthiness. First, consider the following view about praiseworthiness.

> ***Praiseworthiness*:** Assuming relevant responsibility conditions are met, act *A* is praiseworthy if there are requiring reasons to do *A* and there is a justification not to do *A*, and is more praiseworthy, the stronger these reasons to do *A*, and the stronger this justification not to do *A*.

Praiseworthiness is not offered as a complete account of praiseworthiness, but as a view identifying two factors that can amplify an act's praiseworthiness.[23] Perhaps there are other such factors.[24] We may wish to add to *Praiseworthiness* that *A* is done for sufficiently good motivating reasons (that, for instance, your motivating reasons for doing *A* coincide with

---

strong justification to prevent the fact that there is more requiring to save a stranger's index fingers along with another stranger's life than there is to do nothing from making it impermissible to do nothing.

[23] Unlike *Praiseworthiness*, Massoud's 2016 account of praiseworthiness is only indirectly sensitive to requiring reasons (704–7). This difference does not matter for my main purposes here.

[24] For instance, perhaps the strength of the *non-requiring* reasons to do an act can also amplify its praiseworthiness, or perhaps (only) *unusual* responsiveness to requiring reasons is praiseworthy. On non-requiring reasons, see: James Dreier, "Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't," in *Satisficing and Maximizing: Moral Theorists on Practical Reason*, edited by Michael Byron (Cambridge: Cambridge University Press, 2004), 131–154; Terry Horgan and Mark Timmons, "Untying a knot from the inside out: reflections on the 'paradox' of supererogation," *Social Philosophy and Policy* 27 (2010): 29–63; Margaret Little and Coleen Macnamara, "For better or worse: commendatory reasons and latitude," *Oxford Studies in Normative Ethics* 7 (2017): 138–160. On unusualness, see: Urmson 1958, Julia Markovits, "Saints, Heroes, Sages, and Villains," *Philosophical Studies* 158 (2012): 289–311; McElwee 2017.

normative reasons to do *A*). *Praiseworthiness* is offered as a coarse-grained view identifying two factors relevant to praiseworthiness. It is intended to be compatible with a range of fine-grained views about the nature and specific amplifying roles of these factors. For instance, on one view, only *cost-based* justifications not to perform *A* (to perform some alternative instead of *A*) are relevant to praiseworthiness; on other views, other sorts of justifications are relevant too. On one view, you cannot be pro tanto praiseworthy for performing *A* unless it is both to some extent costly for you to do so and recognized by you as such; on other views, you can be. On one view, the reasons to do *A* fix an upper limit on *A*'s praiseworthiness, however costly it is for you to do *A*; on another view, the costliness of doing *A* can amplify *A*'s praiseworthiness without any such upper limit. And so forth.

Next consider the following view about the relationship between requiring reasons, justifications, and blameworthiness.

> ***Blameworthiness*:** Assuming relevant responsibility conditions are met, act *A* is blameworthy if it is not the alternative there is most requiring reason (mass noun) to do and there is no sufficiently strong justification to do *A*, and is more blameworthy, the greater the shortfall in requiring reason between *A* and the alternative there is most requiring reason to do, and the weaker the insufficiently strong justification to do *A*.

The basic idea behind *Blameworthiness* is that you are pro tanto blameworthy for falling short of what there is most requiring reason to do, absent a sufficiently strong justification, and that blameworthiness is amplified by the degrees to which these requiring and justifying considerations fall short. We can, for instance, take the degree to which the requiring reason to do an act falls short of what there is most requiring reason to do, and multiply it by the degree to which the justification to do this act is weaker than the weakest sufficiently strong

justification. As long as there *is* a sufficiently strong justification for failing to do what there is most requiring reason to do, you are fully insulated from blame, regardless of the degree of shortfall in requiring reason.[25] There is an important but commonsensical asymmetry between *Praiseworthiness* and *Blameworthiness*. Whereas you are to various degrees praiseworthy for making sacrifices in response to individual requiring reasons, you need not be *at all* blameworthy for failing to respond to individual requiring reasons. Blameworthiness for failing to respond to a requiring reason gets "switched off" if you do what there is most requiring reason to do, or, failing that, if you do an alternative for which there is a sufficiently strong justification.[26]

This general framework consisting of requiring reasons, justifications, *Praiseworthiness*, and *Blameworthiness*—together with the distinction between being pro tanto praiseworthy and being overall praiseworthy—gives us the tools needed to generate cases of the form picked out by *Impermissible Yet Praiseworthy*. To see this, suppose again you have three alternatives: order the meat dish, order the vegetarian dish, and order the vegan dish. As before, you most want the meat dish. But now suppose you are indifferent between the

[25] Some have argued that there are cases in which an act is blameworthy even though there is a sufficiently strong justification to do it (see Driver 1992). *Blameworthiness* can be modified to accommodate this claim by distinguishing between the strength of justification needed to prevent an act from being made impermissible and the strength of justification needed to prevent an act from being made blameworthy. Alternatively or additionally, we could distinguish between different types of justifications, claiming that the fact that it would involve the use of your property makes it permissible for you not to help a stranger in need, while that the fact that it would be costless for you to help them makes you blameworthy for not helping. In other words, to avoid being blameworthy for not helping, you need a sufficiently strong *cost-based* justification to refrain from helping.

[26] Failing to respond to requiring reasons may leave various forms of moral residue, even when not pro tanto blameworthy. For example, while you are not pro tanto blameworthy for breaking a promise in order to save a drowning child, perhaps you have to compensate the promisee (Ross 1930, 28).

vegetarian and vegan options. Let us assume that the relevant reasons and justifications are as follows: it is costly enough for you to order either the vegetarian dish or the vegan dish that there is a strong enough justification to order the meat dish; since it is no costlier for you to order the vegan dish than it is for you to order the vegetarian dish, there is no sufficiently strong justification to order the vegetarian dish; there is no requiring reason to order the meat dish, strong requiring reason to order the vegetarian dish, and somewhat stronger (and most) requiring reason to order the vegan dish. Given all this, it is impermissible to order the vegetarian dish, permissible to order the meat dish, and permissible to order the vegan dish; nonetheless *Praiseworthiness* and *Blameworthiness* can together imply that the praiseworthiness of ordering the vegetarian dish relevantly exceeds its blameworthiness. It can thus turn out to be overall praiseworthy to order the vegetarian dish, and more so than ordering the meat dish, even when the former is impermissible and the latter is permissible. We can in this way have a case of the form picked out by *Impermissible Yet Praiseworthy*, if the relevant reasons and justifications are as stipulated above.

The framework presented in this section shows how cases of the form picked out by *Impermissible Yet Praiseworthy* are conceptually possible. To move beyond mere conceptual possibility, we need cases in which the relevant reasons and justifications are not merely stipulated or assumed. In the following section, I work through a series of "emergency rescue" cases some of which plausibly *are* of the form picked out by *Impermissible Yet Praiseworthy*. However, it is not my primary aim to persuade readers that the *particular* cases I will focus on are of the form picked out by *Impermissible Yet Praiseworthy*. Instead my main aim is to provide, by way of illustration, a more robust defense of the central claim that there are *some* cases of the form picked out by *Impermissible Yet Praiseworthy* (even if none of the particular ones I discuss are). The general framework of this section, together with the particular cases of

the following section, should give readers a clear enough idea of how to construct their own cases that fit the relevant form.

III. Rescue Cases

With our basic framework in place, we can proceed to some emergency rescue cases that are illustrative of the form picked out by *Impermissible Yet Praiseworthy*. To help minimize noise, let us consider idealized cases in which potentially relevant considerations not explicitly mentioned are held constant or bracketed appropriately.[27] We can, for instance, suppose that those in need of rescue are innocent and have serious interests in continuing to live, that potential rescuers are relevantly informed, and that performing a given rescue does not affect one's abilities or opportunities to help others in the future.[28] Let us begin with the following case.[29]

---

[27] Readers are invited to fill in more concrete details of these cases as they like, so long as they take care not to introduce confounding factors.

[28] To appreciate the significance of this last qualification, see, for example, Portmore 2019 (section 6.4).

[29] For similar or relevant cases see: Charles Fried and Derek Parfit, "Correspondence," *Philosophy and Public Affairs* 8 (1979), 393–397; Derek Parfit, "Future Generations: Further Problems," *Philosophy and Public Affairs* 11 (1982): 113–72, 131; Kagan 1989, 16; Douglas Portmore, "Dual-Ranking Act-Consequentialism," *Philosophical Studies* 138 (2008): 409–427, 420–21; Victor Tadros, *The Ends of Harm*: *The Moral Foundations of Criminal Law* (Oxford: Oxford University Press, 2011), 161–62; Theron Pummer, "Whether and Where to Give," *Philosophy and Public Affairs* 44 (2016): 77–95; Joe Horton, "The All or Nothing Problem," *The Journal of Philosophy* 114 (2017): 94–104; Jeff McMahan, "Doing Good and Doing the Best," in *The Ethics of Giving*, edited by Paul Woodruff (New York: Oxford University Press, 2018): 78–102; Thomas Sinclair, "Are We Conditionally Obligated to Be Effective Altruists?" *Philosophy and Public Affairs* 46 (2018): 36–59; Portmore 2019, section 6.4; Helen Frowe, "If You'll Be My Bodyguard: Agreements to Save and the Duty to Minimize Harm," *Ethics* 129 (2019): 204–229, 210; Ralf Bader, "Agent-Relative Prerogatives and Suboptimal

> ***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.

In *Costly No-Conflict*, there is significantly more requiring reason to save the greater number than there is to do either alternative. Nonetheless, there is a strong enough justification to do nothing. Even if the cost is due entirely to immediate pain, transitional factors, and social injustice, the cost of losing a limb is a serious one. I take it this cost is great enough to make it permissible not to sacrifice your legs to save two strangers.[30] At the same time, it seems

---

Beneficence," *Oxford Studies in Normative Ethics* 9 (2019): 223–250; Muñoz 2020; Christian Barry and Seth Lazar, "Supererogation and Optimisation" (unpublished); Kerah Gordon-Solmon, "Between All and Nothing: Or, Defending the Impermissible" (unpublished). Not all of these authors discuss cases with the same structure as those that are of interest to me here. In some of these cases, one faces a series of choices over time: first one can either incur a personal cost to help others, or not, and if one does incur the cost, one can then at a later time choose whether to help more at no additional personal cost. In the cases that are of interest to me here, one faces a single choice in which one can either incur a personal cost to help others, incur a personal cost that is no greater to help more, or incur no personal cost by helping no one at all. The dialectical importance of the distinction between these sorts of cases in part depends on whether it is possible to argue for conclusions about the single choice cases on the basis of claims about the series of choices cases (McMahan 2018 and Bader 2019 discuss the extent to which it is).

[30] We can suppose the cost is greater if we doubt this is enough. My main arguments also hold independently of what it is that is being sacrificed—be it limbs, lives, time, safety, money, effort, or projects. On intuitions about losing limbs (to help others), see: Christian Barry and Gerhard Øverland, "How Much for the Child?" *Ethical Theory and Moral Practice* 16 (2013): 189–204; Kasper Lippert-Rasmussen, "Barry and Øverland on Singer and assistance-based duties," *Ethics and Global Politics* 12 (2019): 15–23; Elizabeth Barnes, *The Minority Body: A Theory of Disability* (Oxford: Oxford University Press, 2016).

impermissible to save the lesser number (one stranger) instead of saving the greater number (both strangers).[31] This would be so if there were no sufficiently strong justification to save the lesser number.

Is there no sufficiently strong justification to save the lesser number in *Costly No-Conflict*? The most pressing question here is whether there is a justification to save the lesser number that is strong enough to prevent the fact that there is significantly more requiring reason to save the greater number than there is to save the lesser number from making it impermissible to save the lesser number.

There is no plausible *cost-based* justification to save the lesser number. To have a cost-based justification to do act *A*, there must be an alternative that is costlier to you than *A*, so that doing *A* avoids this cost. But neither alternative to saving the lesser number in *Costly No-Conflict* costs you more than saving the lesser number does.

Perhaps instead you have an *ownership-based* justification to use your body (or, more broadly, to live your life) in one way instead of another, even when the cost to you would be the same in every respect, whichever of these two acts you do.[32] First, it is controversial that ownership-based justifications can operate independently of cost in this way. Second, it is particularly controversial that there is any such justification to save the lesser number in *Costly No-Conflict*, since here your altruistic alternatives do not differ at all with respect to *how you*

[31] There is some experimental philosophy that suggests this intuition is widely shared. See: Lucius Caviola and Stefan Schubert, "Do people consider it obligatory to donate effectively?" (unpublished).

[32] On ownership-based justifications or similar, see: Seana Shiffrin, "Moral Autonomy and Agent-Centered Options," *Analysis* 51 (1991): 244–54; Frances Kamm, "Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status," *Philosophy and Public Affairs* 21 (1992): 354–89; Fiona Woollard, *Doing and Allowing Harm* (Oxford: Oxford University Press, 2015); Seth Lazar, "Moral Status and Agent-Centred Options," *Utilitas* 31 (2019): 83–105.

use your body (apart from the effects of how it is used on the number of strangers saved). Third, even if we grant that there is *an* ownership-based justification to save the lesser number, it seems implausible that it is strong enough to prevent the fact that there is significantly more requiring reason to save the greater number from making it impermissible to save the lesser number. After all, ownership-based justifications do not seem strong enough to prevent the fact that there is significantly more requiring reason to save the greater number from making it impermissible to save the lesser number in a variant of *Costly No-Conflict* in which saving either the one stranger or both costs you nothing but involves the use of what is yours (at least as long as it is the same use either way).

In sum, in *Costly No-Conflict*, there is significantly more requiring reason to save the greater number than there is to do either alternative. But while there is a sufficiently strong justification to do nothing, there is no sufficiently strong justification to save the lesser number. It is impermissible to save the lesser number and permissible to do nothing.[33] To determine whether *Costly No-Conflict* is illustrative of the form picked out by *Impermissible Yet*

[33] Horton 2017 argues that the claim that in *Costly No-Conflict* it is permissible to do nothing and impermissible to save the lesser number has the implausible implication that, if you are not going to save both strangers, you are required to do nothing. This inference relies on the principle that, when *A* and *B* are your only permissible alternatives, if you are not going to do *B*, you are required to do *A* (Horton 2017, 96). But we should reject this principle. In *Costly No-Conflict*, it is permissible to do nothing and impermissible to save the lesser number even though there is more requiring reason to save the lesser number than there is to do nothing (as there is a sufficiently strong justification to do nothing but no sufficiently strong justification to save the lesser number). When your alternatives are *A*, *B*, and *C*, and there is more requiring reason to do *C* than there is to do *A*, it is not plausible that, if you are not going to do *B*, you are required to do *A*. For further discussion, see my, "All or Nothing, But If Not All, Next Best or Nothing," *The Journal of Philosophy* 116 (2019): 278–291; and my paper co-authored with Daniel Muñoz, "Supererogation and Conditional Obligation" (unpublished).

Praiseworthy, we need to ask whether it is overall praiseworthy to save the lesser (and more so than doing nothing).

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.[34] By the same token, bad motivating reasons could render you blameworthy for doing nothing or for saving the greater number, even though these acts would remain permissible.[35] For instance, you may be blameworthy for doing nothing out of unjustified hate for the strangers you could save, but not for doing nothing out of concern for your legs. To simplify, I will here limit my focus to cases in which permissible acts are done for the best motivating reasons available. Is the best available motivating reason for saving the lesser number good enough for you to be at least pro tanto praiseworthy for doing so?

Arguably the best—least bad—available motivating reason for saving the lesser number in *Costly No-Conflict* is what I call "completely innumerate altruism." That is, you are moved to *help someone*, but are indifferent to how many. You are fully aware of and moved by the plights of each, but in any given choice situation you are not *moved more* by more plights than fewer.[36] Equally moved to save just one as you are to save both, your choice to save

[34] This is how both Horton 2017 (94) and McMahan 2018 (94–9) characterize *Costly No-Conflict*.

[35] See Scanlon 2008 and Frances Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (New York: Oxford University Press, 2007), 132.

[36] You are thus not moved more by more plights than fewer, whether it is a case of two rather than one, or two million rather than one. A more intermediate possibility is that you are indifferent to how many you help in no-conflict cases only if the (difference in) number is below some finite threshold. I will set aside possibilities like

someone culminates in arbitrarily picking "just one."[37] Your innumerate altruism is anonymous in that you would be equally moved to save just one as you are to save both were the positions of the two swapped. Finally, your innumerate altruistic motivation is quite a powerful one, taking precedence over your strong motivation to keep your legs. Is it irrational to be so moved by the plights of each but not moved more by the plights of more?[38] It would appear so, but not in a way that renders praise and blame inapt. Similarly, it seems irrational to be moved by the plights of each but not moved by plights that would occur on Tuesdays. But for all that you can be praiseworthy for preventing a Monday harm, and blameworthy for allowing a Tuesday harm.

It seems that completely innumerate altruism is not a bad enough motivating reason to prevent saving the lesser number from being pro tanto praiseworthy. *Praiseworthiness* then implies you are pro tanto praiseworthy for saving the lesser number out of completely

---

these and focus on the simpler motivating reason of completely innumerate altruism (when I turn to other sorts of cases—such as conflict cases—I will consider other motivating reasons for saving the lesser number). Completely innumerate altruism may be seen as a motivational reflection of these ancient thoughts: "Whosoever preserves a single soul … scripture ascribes as much [merit] to him as if he had preserved a complete world" (Talmud: Sanhedrin 37a) and "… whosoever saves the life of one, it shall be as if he had saved the life of all of mankind" (Quran 5:32). It may also be seen as a motivational reflection of the (related) thought that each person's life is of infinite moral value, so that the moral value of saving just one person is equivalent to that of saving an arbitrarily large group of people that contains this same person.

[37] Perhaps you cannot be responsible for arbitrarily *picking* one alternative over another, insofar as that does not constitute a *choice*. You nonetheless can be responsible for your choice to save someone by arbitrarily picking.

[38] On rationality and saving the lesser number, see: Véronique Munoz-Dardé, "The Distribution of Numbers and the Comprehensiveness of Reasons," *Proceedings of the Aristotelian Society* 105 (2005): 191–217; Tom Dougherty, "Rational numbers: a non-consequentialist explanation of why you should save the many and not the few," *Philosophical Quarterly* 63 (2013): 413–427; Kieran Setiya, "Love and the Value of a Life," *Philosophical Review* 123 (2014): 251–80.

innumerate altruism. After all, in saving the lesser number you are heroically sacrificing your legs in response to a requiring reason to help someone.

But even if it is pro tanto praiseworthy to save the lesser number out of completely innumerate altruism—as seems plausible—it is not clear that this is overall praiseworthy, or more so than doing nothing (which is not pro tanto praiseworthy). After all, in saving the lesser number you also allow someone else to die *gratuitously*, in that your available alternative of not allowing them to die is at least as favorable in every respect.[39] This would seem to make for a rather substantial shortfall in requiring reason between saving the lesser number and saving the greater number. Given this, and that there is no sufficiently strong justification to save the lesser number, *Blameworthiness* has the plausible implication that you are substantially blameworthy for saving the lesser number out of completely innumerate altruism. It is not clear that the praiseworthiness of saving the lesser number is great enough to counterbalance this blameworthiness.

Costly No-Conflict does not provide clear support for *Impermissible Yet Praiseworthy*. But the foregoing discussion puts us in a good position to consider some further, and more promising, cases. Consider:

> ***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.[40]

[39] By "at least as favorable in every respect," I intend to include non-moral (prudential) respects.

[40] This is a version of my *Hot Death* case (Pummer 2019, 283). Note that dying usually removes all opportunities to help in the future. We can suppose that in this case there are no opportunities to help in the future anyway.

Arguably the best available motivating reason for pressing the red button is what I call "innumerate smaltruism." That is, whenever you are confronted with a mix of big plights (such as deaths) and relatively small plights (such as lost fingers), you are moved by the plights of each, moved more by more big plights than fewer, but not moved more by only a few more relatively small plights than fewer. Innumerate smaltruism seems a better motivating reason than completely innumerate altruism.[41] Given this, and that pressing the red button constitutes making an enormous sacrifice (your life) in response to very strong requiring reasons (to save ten people), *Praiseworthiness* has the plausible implication that pressing the red button out of innumerate smaltruism is very praiseworthy. As pressing the red button also allows someone to lose a finger gratuitously, the shortfall in requiring reason between pressing the red button and pressing the green button is—though significant—relatively small. *Blameworthiness* has the plausible implication that it is only somewhat blameworthy to press the red button. It thus seems plausible that you are overall praiseworthy for pressing the red button, and more so than you are for doing nothing (which is not pro tanto praiseworthy). Insofar as it is impermissible to press the red button and permissible to do nothing, *Red/Green No-Conflict* would therefore appear to support *Impermissible Yet Praiseworthy*.[42]

[41] To the extent that the manifestation of dispositions matters for blameworthiness, a completely innumerate altruist would be more blameworthy for pressing the red button in *Red/Green No-Conflict* than an innumerate smaltruist would be.

[42] One may wonder whether what is doing the intuitive work in this case is the fact that it is not obvious (to you, the agent) that it is impermissible to press the red button. After all, some hold that moral ignorance mitigates blameworthiness in the way empirical ignorance does. On my view, you can be overall praiseworthy for pressing the red button, even if you correctly believe that it is impermissible to do so. As explained in section II, you can be both pro tanto praiseworthy and pro tanto blameworthy for performing a token act. For further discussion of moral ignorance, see, for instance: Gideon Rosen, "Culpability and ignorance," *Proceedings of the Aristotelian*

The explanation of why it is impermissible to save the lesser number yet permissible to do nothing in *Costly No-Conflict* is instructive here. Even though the harm of a lost finger is small relative to the harm of death, it seems there is a significant requiring reason to prevent the loss of a finger, and indeed that in *Red/Green No-Conflict* there is significantly more requiring reason to press the green button than to press the red button.[43] There is a sufficiently strong justification to press neither button.[44] Is there a sufficiently strong justification to press the red button (that is, save the lesser number) in *Red/Green No-Conflict*? That there is a sufficiently strong cost-based or ownership-based justification to save the lesser number is just as implausible in this case as it is in *Costly No-Conflict*, but we might wonder whether some further type of justification is available in *Red/Green No-Conflict*. After all, it seems *very* good to press either button. Perhaps there is a sufficiently strong justification to do *A*, even though there is significantly more requiring reason to do *B* than there is to do *A*, when *A* is *good enough*.[45] And perhaps pressing the red button is good enough. There might be cases in which such a *satisficing justification* to do act *A* prevents the fact that there is more requiring reason to do an alternative from making *A* impermissible. But it does not seem plausible that there is a sufficiently strong satisficing justification to press the red button in *Red/Green No-Conflict*, as pressing the red button allows someone to lose a finger gratuitously. More generally, when

---

Society 103 (2003), 61–84; Elinor Mason, "Moral ignorance and blameworthiness," *Philosophical Studies* 172 (2015): 3037–3057.

[43] If a lost finger seems too small a harm, we can instead make it a lost hand or arm. Also note that Kamm 2007 would not count the finger as an "irrelevant good" here, since the people whose lives you can save by pressing the red button are the same as those whose lives you can save by pressing the green button.

[44] If you doubt this, you can retell the case so that pressing either button would cause you even greater agony, serious non-hedonic harm, the destruction of your life projects, and so forth.

[45] An act might be "good enough" in absolute terms or relative to its alternatives. For discussion, see: Thomas Hurka, "Two Kinds of Satisficing," *Philosophical Studies* 59 (1990): 107–111.

acts *A* and *B* are perfectly alike but for the fact that *A* allows significant harm gratuitously whereas *B* does not, it seems there cannot be a sufficiently strong satisficing justification to do *A*, no matter how good *A* is.[46] In sum, the requiring and justifying considerations present in *Red/Green No-Conflict* together support the claim that it is impermissible to press the red button and permissible to do nothing. So *Red/Green No-Conflict* appears to support *Impermissible Yet Praiseworthy*. There are further cases.

The practical relevance of this paper would be substantially limited if the only plausible cases of impermissible acts that are overall praiseworthy, and more so than permissible alternatives, involved allowing harm gratuitously out of innumerate smaltruism. In most real-world cases we are unable to help everyone at once; when saving the larger group *precludes* (rather than includes) saving the smaller group, saving the smaller group does not constitute allowing harm gratuitously.[47] Consider two cases that have this feature.

> ***Costless Conflict*:** Three strangers face a deadly threat. You can, at no cost to yourself, save one of them or save the other two. Tragically, you cannot save all three.

> ***Costly Conflict*:** Three 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 the other two.

The claim that it is impermissible to save the lesser number in *Costless Conflict* is more controversial than the claim that it is impermissible to save a stranger at no cost to yourself

[46] If a "good enough" act *just is* an act there is a sufficiently strong justification to do, then act *A* cannot be good enough.

[47] See Horton 2017, McMahan 2018, and Sinclair 2018 on charitable giving.

instead of saving both that same stranger and another at no cost to yourself.[48] The claim that it is impermissible to save the lesser number in *Costly Conflict* is accordingly more controversial than the claim that it is impermissible to save the lesser number in *Costly No-Conflict*. Whereas in *Costly No-Conflict* saving the lesser number constitutes allowing significant harm gratuitously, in *Costly Conflict* saving the lesser number does not constitute allowing significant harm gratuitously. In the latter case, the fact that you would save *that* one stranger at least *significantly* favors saving the lesser number without also favoring any other alternative (particular individuals matter above and beyond the amounts of well-being instantiated in their lives).

Nonetheless, I hold that it is impermissible to save the lesser number in *Costless Conflict*. Since a proper defense of this claim lies beyond the scope of this paper, I will just assert it here. But it is important to observe that we need not claim that what makes it impermissible to save the lesser number in cases like *Costless Conflict* and *Costly Conflict* is that it is impermissible to do *A* when *A* fails to produce the impartially best outcome and there

[48] For views according to which you are not required to save the greater number in cases like *Costless Conflict*, see: Elizabeth Anscombe, "Who is Wronged?" *Oxford Review* 5 (1967): 16–17; John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6 (1977): 293–316; Munoz-Dardé 2005; Tyler Doggett, "Saving the Few," *Noûs* 47 (2013): 302–315. On deciding by lot, see: John Broome, "Fairness," *Proceedings of the Aristotelian Society* 91 (1990): 87–101; Ben Bradley, "Saving People and Flipping Coins," *Journal of Ethics and Social Philosophy* 3 (2009): 1–13. For defenses of the popular view that you are required to save the greater number, see: Derek Parfit, "Innumerate Ethics," *Philosophy and Public Affairs* 7 (1978): 285–301; Frances Kamm, *Morality, Mortality, Vol. 1: Death and Whom to Save from It* (New York: Oxford University Press, 1993); T. M. Scanlon, *What We Owe to Each Other* (Cambridge: Harvard University Press, 1998); Michael Otsuka, "Skepticism about saving the greater number," *Philosophy and Public Affairs* 32 (2004): 413–426; Dougherty 2013.

is no sufficiently strong justification to do *A*.[49] Though saving the life of a joyful person may produce a much better outcome than would saving the life of a bored person, it seems you are not required to save the former instead of the latter (at least, assuming their lives are worth continuing and they have serious interests in continuing them).[50] It would also seem you are not required to save the life of the joyful person along with someone else's finger instead of saving the life of the bored person. It is plausible that, when there is a conflict between saving X's life and saving Y's life, adding Z's finger to either side of the scales cannot by itself make it impermissible to save those on the opposite side.[51]

At the same time, I take it that adding Z's *life* to Y's side of the scales *can* make it impermissible to save X's life instead of saving the lives of both Y and Z. What makes it impermissible to save the lesser number in *Costless Conflict* is that there is significantly more requiring reason to save the lives of two people than there is to save the life of one other person, and no sufficiently strong justification to save the lesser number. It seems there is likewise significantly more requiring reason to save the greater number than there is to save the lesser number in *Costly Conflict*. And the claim that there is a sufficiently strong justification to save the lesser number is just as implausible in *Costly Conflict* as it is in previous cases.[52]

---

[49] See Sinclair's 2018 response to me (Pummer 2016) and Horton 2017 on this score.

[50] See Doggett 2013 and Frances Kamm, "Aggregation, Allocating Scarce Resources, and the Disabled," *Social Philosophy and Policy* 26 (2009), 148–97.

[51] Recall that in *Red/Green No-Conflict* the red button and the green button save the lives of the *same* people, so the fact that the green button also saves a finger does make it impermissible to press the red button.

[52] There is a potential disanalogy when it comes to satisficing justifications, as in response to the suggestion that there is a sufficiently strong satisficing justification to save the lesser number in *Red/Green No-Conflict* I appealed to the fact that saving the lesser number allows significant harm gratuitously. But it also seems plausible that when acts *A* and *B* are perfectly alike but for the fact that *A* allows strangers Y and Z to die whereas *B* allows stranger X to die, there cannot be a sufficiently strong satisficing justification to do *A*, no matter how good *A* is. This last

Like *Costly No-Conflict*, *Costly Conflict* is a case in which it is impermissible to save the lesser number, though permissible to do nothing. But there are two important differences that plausibly make it significantly less blameworthy to save the lesser number in *Costly Conflict* than it is to do so in *Costly No-Conflict*. First, because in *Costly Conflict* the requiring reason (count noun) to save the one is in conflict with the requiring reasons to save the two, there is less of a shortfall in requiring reason (mass noun) between saving the lesser number and saving the greater number in *Costly Conflict* than there is in *Costly No-Conflict*. Given this, *Blameworthiness* has the plausible implication that it is less blameworthy to save the lesser number in *Costly Conflict* than it is to do so in *Costly No-Conflict*. Second, while it seems that the best available motivating reason for saving the lesser number in *Costly No-Conflict* is that of completely innumerate altruism—you are moved to *help someone*, but are anonymously indifferent to how many you help—there are better motivating reasons available for saving the lesser number in *Costly Conflict*. You might save the lesser number in the latter case out of "moderately innumerate altruism." That is, you might be moved by the plights of each and indeed moved more by the plights of more, but not moved more by the plights of the many than by the plights of the few, insofar as the few are different from the many (equally moved to save the one as you are to save the other two, you arbitrarily pick "the one"). Or, perhaps better still, you might save the lesser number out of "singular altruism." That is, you might be moved to save the one on the basis of some singular thought about them or their particular

---

claim is compatible with the view that it is permissible to save the lives of ten strangers along with X's life instead of saving the lives of ten *different* strangers along with the lives of Y and Z (acts are less alike when they involve saving different people). It is also compatible with the possibility of cases in which there is a sufficiently strong satisficing justification to do *A*, when *A* and *B* are perfectly alike but for the fact that *A* allows strangers Y and Z to each lose a *finger* whereas *B* allows stranger X to lose a finger (and there is significantly more requiring reason to do *B* than there is to do *A*).

plight, such as "Lorraine is drowning."[53] Arguably the blameworthiness of saving the lesser number in *Costly No-Conflict* is aggravated by the fact that the best motivating reason available for doing so—completely innumerate altruism—involves *disregarding* the number of individual plights. By contrast, moderately innumerate altruism and singular altruism do not involve disregarding the number of plights, but failing to appropriately balance the plights on each side.

While in *Costly No-Conflict* it is not clear that the praiseworthiness of saving the lesser number is great enough to counterbalance the blameworthiness of doing so, in *Costly Conflict* it is plausible that the praiseworthiness of saving the lesser number can be great enough to counterbalance the blameworthiness of doing so. The latter case thus appears to support *Impermissible Yet Praiseworthy*. Some further cases appear more compelling still. Consider:

> ***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.

As in *Costly Conflict*, it seems there is significantly more requiring reason to save the greater number (press the green button) than there is to save the lesser number (press the red button). And the reasons offered above in support of the claim that there is no sufficiently strong justification to save the lesser number in *Costly Conflict* likewise support the claim that there

[53] Obviously the details of *Costly Conflict* will determine whether it is possible for your motivating reason for saving the lesser number to be a singular thought (this may be impossible if those you can rescue are hidden behind a curtain). For further discussion, see: Setiya 2014; Philip Pettit, "Love and Its Place in Moral Discourse," in *Love Analyzed*, edited by Roger Lamb (Boulder: Westview Press, 1997), 153–63.

is no sufficiently strong justification to save the lesser number in *Red/Green Conflict*. So it is plausible to claim that it is impermissible to press the red button in *Red/Green Conflict*. But, at least assuming your motivating reason is something in the vicinity of moderately innumerate altruism or singular altruism, it seems plausible that you are overall praiseworthy for pressing the red button, and more so than doing nothing. After all, the blameworthiness of pressing the red button in this case seems no greater than that of saving the lesser number in *Costly Conflict*, while the praiseworthiness of pressing the red button seems considerably greater than that of saving the lesser number in *Costly Conflict*.

Similar remarks apply to the following variant, which could convince those who believe it is permissible to save the lesser number in conflict cases (assuming they agree that there is more requiring reason to save one stranger's life than there is to save another's legs).

> ***Equinumerous Red/Green Conflict*:** You can do nothing, press a red button thereby saving the lives of ten strangers and saving stranger X's legs, or press a green button thereby saving the lives of these same ten strangers and saving stranger Y's life. Pressing either button will also cause you to drop into a fiery pit and die.

In sum, while *Costly No-Conflict* does not provide clear support for *Impermissible Yet Praiseworthy*, it is plausible that a variety of other cases do—including *Costly Conflict*, but perhaps especially *Red/Green No-Conflict*, *Red/Green Conflict*, and *Equinumerous Red/Green Conflict*.

IV. The Limits of Praiseworthiness

In section I, I presented cases in which praiseworthiness appears to attach to impermissible acts. These impermissible acts seem pro tanto praiseworthy, and some might even be overall

praiseworthy. In sections II and III, I showed how praiseworthiness might come apart from deontic status in the more radical way described by *Impermissible Yet Praiseworthy*—impermissible acts might not only be overall praiseworthy, but more so than permissible alternatives. In this section, I will suggest that there may nonetheless be limits on when impermissible acts can be overall praiseworthy. The holding of such limits would constrain the form that cases picked out by *Impermissible Yet Praiseworthy* can take.

According to the *First Limit*, for an act to be pro tanto praiseworthy, it must meet or exceed what is required.[54] That is, there must be at least as much reason to do this act as there is to do a permissible alternative (you are required to do some or other permissible alternative). In section I, I presented examples of acts that are impermissible and yet seemingly pro tanto praiseworthy (the vegetarianism, teaching, and bank robbery examples). The *First Limit* implies that these acts are in fact not pro tanto praiseworthy—insofar as you are required to order the vegan dish, it is not praiseworthy to order the vegetarian dish. The rescue cases discussed in section III are different. Saving the lesser number in cases like *Red/Green No-Conflict*, *Costly Conflict*, and *Red/Green Conflict* is more responsive to the balance of reasons than is required—there is more reason to save the lesser number than there is to permissibly save no one. These cases provide us with acts that in this way go beyond what is required, yet are impermissible. The *First Limit* does not imply that these impermissible acts cannot be pro tanto praiseworthy, and so does not present an obstacle to my argument for *Impermissible Yet Praiseworthy*. Nonetheless, as mentioned earlier, the *First Limit* seems a somewhat stingy view

[54] We might more extremely claim that, for an act to be pro tanto praiseworthy, it must exceed what is required. This claim conflicts with the intuition (noted in section I) that you are praiseworthy for missing a nonrefundable international flight to take a seriously injured cyclist to the hospital, even though this meets but does not exceed what is required.

of praiseworthiness in that it implies you could sacrifice and sacrifice in response to moral reasons and deserve *no* praise whatsoever for this if you still fall a bit short.

According to the *Second Limit*, then, for an act to be *overall* praiseworthy, it must meet or exceed what is required. Like the *First Limit*, the *Second Limit* is compatible with the argument I provided for *Impermissible Yet Praiseworthy* (in the rescue cases in section III, the impermissible act of saving the lesser number exceeds what is required). But the *Second Limit* is a less stingy view in that it is compatible with the intuition that you can be somewhat praiseworthy for making substantial sacrifices in response to moral reasons even if you fall short of what you are required to do (by, for example, ordering the vegetarian dish rather than the vegan one). The *Second Limit* has seemingly plausible implications in further cases as well.

Consider a variant of the bank robbery example from earlier: you can either at no cost to yourself do nothing, or at significant cost to yourself press a button that saves the homeless stranger's legs yet simultaneously causes two innocent bank employees to lose their legs. These are the only available alternatives. I take it there is a requiring reason to press the button, most requiring reason to do nothing, and no sufficiently strong justification to press the button. *Praiseworthiness* implies that pressing the button is pro tanto praiseworthy, and *Blameworthiness* implies pressing the button is pro tanto blameworthy. It is plausible that this act's blameworthiness is relevantly greater than its praiseworthiness, and that it is therefore overall blameworthy. But, according to *Praiseworthiness*, holding fixed the requiring reasons to press the button, it is more praiseworthy to press the button, the greater the cost to you of pressing it. Couldn't this praiseworthiness then grow very large, if it cost you an awful lot to press the button? The *Second Limit* has the plausible implication that it is *not* overall praiseworthy to press the button, regardless of how costly it is for you to do so. Pressing the button does not meet or exceed what is required, which is to do nothing. There is more reason to do nothing than there is to press the button.

Nonetheless, it might seem that the *Second Limit* remains too stingy a view of praiseworthiness. After all, it implies that, if we are each required to eliminate air travel, you—an avid traveler—cannot be overall praiseworthy for reducing your annual number of flights from 36 all the way down to 1. We might have thought that the praiseworthiness of going so far in response to reasons at substantial cost to yourself is relevantly greater than the blameworthiness of not making it all the way to the level of responsiveness required. There is a third view that accommodates this thought, as well as has the plausible implication that it is not overall praiseworthy to press the button in the variant of the bank robbery example, regardless of how costly it is for you to do so.

According to the *Third Limit*, if *A* is not what there is most reason to do and there is no sufficiently strong justification to do *A*, then, for *A* to be overall praiseworthy, the difference between the reason to do *A* and the reason to do the alternative there is least reason to do must be sufficiently great relative to the difference between the reason to do *A* and the reason to do the alternative there is most reason to do. The basic idea here is that the blameworthiness of *A* for unjustifiably falling short of what there is most reason to do can be counterbalanced by enough praiseworthiness (to render *A* overall praiseworthy) only if *A* rises far enough above what there is least reason to do. The *Third Limit* is compatible with the thought that it is overall praiseworthy to reduce your air travel to once per year, even if you are required to eliminate it altogether. The difference between the reason to reduce your air travel to once per year and the reason to fly whenever you feel like it is much greater than the difference between the reason to reduce your air travel to once per year and the reason to eliminate it altogether.[55] The *Third Limit* also has the plausible implication that it is not overall praiseworthy to press the button in the variant of the bank robbery example, regardless of how costly it is to you to do so. The

[55] This seems plausible even if the fact that an act is impermissible itself provides a reason against doing it. For discussion, see: Stephen Darwall, "But It Would Be Wrong," *Social Philosophy and Policy* 27 (2010): 135–57.

difference between the reason to press the button and the reason to do the alternative there is least reason to do is zero, since pressing the button just is the alternative there is least reason to do. Zero cannot be sufficiently great relative to the difference between the reason to press the button and the reason to do the alternative there is most reason to do (do nothing).

It is not my aim here to defend the *Third Limit*.[56] But if it is correct, it will constrain the general form that cases of the sort picked out by *Impermissible Yet Praiseworthy* can take. Namely they will have to take something close to the general form of three-alternative cases like those offered above in sections II and III. Here is proof. Assume for reductio that *A* and *B* are the only two alternatives, that *A* is impermissible, that *B* is permissible, and that *A* is overall praiseworthy. For *A* to be impermissible, it must be that there is most requiring reason to do *B*, and no sufficiently strong justification to do *A*. Then, according to the *Third Limit*, *A* is overall praiseworthy only if the difference between the reason to do *A* and the reason to do the

[56] In addition to noting that the *Third Limit* has plausible implications in the air travel and bank robbery examples, we might think it can be supported by the following rationale: The fact that *A* responds to *a* reason—for example, a reason to help the homeless person—will not always be enough to render *A overall* praiseworthy, regardless of how costly it is for you to do *A*. It will not be enough when *A* unjustifiably falls short of what there is most requiring reason to do, as then *A* will be more *blameworthy*, the costlier or more difficult it is for you to do *A* than it is to do what you have most requiring reason to do (see, for instance, Nelkin 2016 (368–69)). If scaling up the cost of *A* so amplifies *A*'s blameworthiness, then scaling up the cost of *A* will need to simultaneously amplify its praiseworthiness to a relevantly greater extent, for *A* to be overall praiseworthy. But as *A*'s cost-amplified blameworthiness is based upon a difference between the reason to do *A* and the reason to do the alternative there is most reason to do, it is plausible that the cost-amplified praiseworthiness needed to counterbalance this would similarly be based upon a difference in reason (mass noun)—namely between the reason to do *A* and the reason to do the alternative there is least reason to do. Different precisifications of the *Third Limit* can take different views as to how far above what there is least reason to do *A* needs to rise, for a given shortfall between *A* and what there is most reason to do. Which precisification is correct may depend on the respective rates at which *A*'s praiseworthiness increases and *A*'s blameworthiness increases, as the cost to you of doing *A* is scaled up.

alternative there is least reason to do (*A*) is sufficiently great relative to the difference between the reason to do *A* and the reason to do the alternative there is most reason to do (*B*). Since a difference of zero cannot be sufficient (relative to the difference between the reason to do *A* and the reason to do *B*), *A* is not overall praiseworthy. But this contradicts the assumption that *A* is overall praiseworthy. So, we have a case in which impermissible act *A* is overall praiseworthy, and more so than permissible alternative *B*, only when *A* is made impermissible by the presence of a third alternative *C*—that is, when there is more requiring reason to do *C* than there is to do *A* and no sufficiently strong justification to do *A*. According to the *Third Limit*, then, *A* will be overall praiseworthy only if there is more reason to do *A* than there is to do *B* (assuming *B* is the only alternative to *A* besides *C*). And for *B* to remain permissible, there will in turn have to be a sufficiently strong justification to do *B*. Thus, if the *Third Limit* holds, cases of the sort picked out by *Impermissible Yet Praiseworthy* will have something close to the general form of three-alternative cases like those offered above in sections II and III (imagine *A* is "save the lesser number," *B* is "do nothing," and *C* is "save the greater number").

Even if cases of the sort picked out by *Impermissible Yet Praiseworthy* all share the same general form, it is important to observe that they need not all share the same content. Section III looks at cases of innumerate altruism, in which you can do nothing, rescue the few at great cost to yourself, or rescue the many at the same cost to yourself. To my mind, these are among the most compelling cases of the sort picked out by *Impermissible Yet Praiseworthy* (especially *Red/Green No-Conflict* and *Red/Green Conflict*).[57] But perhaps others will find other cases more compelling still, such as variants of the vegetarianism or air travel examples

[57] In cases like *Red/Green No-Conflict* and *Red/Green Conflict*, not only is it plausible that saving the lesser number (impermissible) is more praiseworthy overall than doing nothing (permissible), but it might also be plausible that saving the lesser number (impermissible) is *nearly as* praiseworthy overall as saving the greater number (permissible).

in which it is very costly to you to respond to the relevant reasons. Or take an example involving promise-keeping: Suppose you have made eleven promises to your business associates. You have a strong requiring reason to keep ten of them, and a significant but relatively weak requiring reason to keep the eleventh. In virtue of some great unexpected cost to you of keeping these promises, there is a sufficiently strong justification not to keep any of them.[58] Now suppose you can do nothing, incur this great cost to keep ten of the promises, or incur this same great cost to keep all eleven. It may well be impermissible yet overall praiseworthy to incur the cost to keep just the ten promises (and more so than doing nothing).[59]

V. Impermissible Yet Praiseworthy in Practice

While it does seem to me that some of the cases of innumerate altruism explored in section III plausibly are of the form picked out by *Impermissible Yet Praiseworthy*, it is not my primary aim to show this. Instead my main aim has been to defend the plausibility of the claim that there exist *some* cases of this general form, whether or not they are the particular ones I discuss. But the particular cases I have discussed, together with the general framework presented in

[58] Suppose the requiring reason to keep these promises is not defeated by the unexpected cost of keeping them.

[59] Arguably you can also go beyond the call of duty in ways worthy of praise without making any significant sacrifice, as when you do a small kindness or forgive a wrongdoer. But while it could be impermissible to keep fewer promises rather than keep more at no greater cost, it is somewhat hard to see how it could be impermissible to do fewer small kindnesses rather than do more at no greater cost, or to forgive fewer wrongdoers rather than forgive more at no greater cost. This is because while there is significant requiring reason to keep promises, and plausibly significantly more requiring reason to keep more promises than there is to keep fewer, it is far from clear whether there is any significant requiring reason to do small kindnesses or to forgive (or more requiring reason to do more such things). Doing small kindnesses and forgiving are often taken to be paradigm examples of acts there is merely *non-requiring* reason to do. See footnote 24.

section II—and supplemented with one or more of the limits of praiseworthiness from section IV—should give readers a good sense of how to construct many additional cases that fit the relevant form. I will now conclude with a brief discussion of some potential practical implications.

I have here focused on somewhat idealized examples in an effort to provide a clear and minimally noisy case for *Impermissible Yet Praiseworthy*. There is good reason to believe there are also plenty of real-world cases that are of the sort picked out by *Impermissible Yet Praiseworthy*. Arguably many can be found in the growing literature on "effective altruism."[60] For instance, that literature suggests there are many cases in which the same monetary donation can save at least twice as many lives (if not ten or even a hundred times as many) if given to one charity rather than another. In a sizable subset of those cases, we have no special or personal connection to either charity, or at least no greater such connection to one than the other, and yet because we have already given an awful lot in the past, because we are choosing between such a large donation to one charity and an equally large donation to the other, or because morality turns out to be quite undemanding, we are not required to make the donation in question. Although ineffective charitable giving is often attributable to ignorance of relevant cost-effectiveness differences, it is also often attributable to a degree of (moderately) innumerate altruism on our part. Such cases of ineffective giving may be relevantly analogous to *Costly Conflict*, in that there is a sufficiently strong cost-based justification for doing nothing, and yet saving the lives of *n* strangers at cost C is impermissible given that we can

[60] For example, see: Toby Ord, "The Moral Imperative towards Cost-Effectiveness in Global Health," *Centre for Global Development* (2013): www.cgdev.org/content/publications/detail/1427016; William MacAskill, *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference* (London: Faber & Faber, 2015); *Effective Altruism: Philosophical Issues*, edited by Hilary Greaves and Theron Pummer (Oxford: Oxford University Press, 2019).

instead save the lives of at least 2*n* different strangers at cost C.[61] Although I have here focused on saving more or fewer lives, numerate altruism is not only about appropriate responsiveness to the number of individuals we can help, but also the degree to which we can help individuals and the probability our acts will actually help them. And giving to charity is but one relevant sort of real-world case—others may include volunteering, choosing a career, or fighting humanitarian wars.[62] To be sure, there are many significant disanalogies between these real-world cases and idealized ones like *Costly Conflict*.[63] Here is not the place to argue that these potential disanalogies do not make inapt the intended analogy.

To the extent that such an analogy holds, there are many real-world cases of impermissible altruism that is more praiseworthy overall than permissible and blameless non-altruism. When we find ourselves engaged in such impermissible altruism, we should not focus entirely on our conduct's negative aspect, and not only because of the counterproductivity of doing so. Our conduct is not only overall praiseworthy, but (considerably) more so than the permissible non-altruism of many others. At the same time, we should not lose sight of the fact that our conduct is impermissible, and has a blameworthy aspect. The ideally appropriate response to such impermissible altruism is mixed, a proportionate reflection of both the praiseworthy as well as blameworthy aspects of our conduct. When such an ideal response is

[61] As suggested earlier (in footnote 52), *perhaps* when the ratio difference between non-overlapping groups is fairly small it is permissible to save the smaller group instead of saving the larger group (for example, saving 11 strangers rather than 12 different strangers, or 99 rather than 100).

[62] On this last sort of case, see: Victor Tadros, "Unjust Wars Worth Fighting For," *Journal of Practical Ethics* 4 (2016): 52–80; Kieran Oberman, "War and Poverty," *Philosophical Studies* 176 (2019): 197–217.

[63] See, for instance: Sinclair 2018; Woollard 2015; Kamm 2007; Caspar Hare, "Obligations to Merely Statistical People," *The Journal of Philosophy* 109 (2012): 378–390; Andreas Mogensen, "The Callousness Objection," in Greaves and Pummer 2019, 227–240.

impossible, or psychologically infeasible, we may have to navigate between under-blaming ourselves and under-praising ourselves. When it comes to our responses to the impermissible yet praiseworthy altruism of others, these matters are arguably more delicate still.[64]

[64] On asymmetries between self-blame and other-blame, see: Dana Nelkin, "How Much to Blame?: An Asymmetry in the Norms of Self-Blame and Other-Blame" (unpublished).