

If You'll Be My Bodyguard: Agreements to Save and the Duty to Minimize Harm*

Helen Frowe

This article explores how agreements to preferentially save can ground an exception to the duty to minimize harm when saving. A rescuer preferentially saves if she knowingly fails to minimize harm among prospective victims, even though minimizing harm would not have imposed greater costs on the rescuer herself. Allowing rescuers to act on agreements to preferentially save is justified by the reasons we have to respect the agreements that agents form as a means of pursuing their own ends.

I. INTRODUCTION

This article considers the permissibility of preferential saving. A rescuer preferentially saves if she knowingly fails to minimize harm among prospective victims, even though minimizing harm would not have imposed (significantly) greater cost on the rescuer herself.¹ For example, a rescuer

* Versions of this paper were presented at the Society for Applied Philosophy Annual Conference at Queen's University, Belfast; the Society for Applied Philosophy Panel at the 2017 Pacific APA; the "Moralities of Warfare" Conference at the University of Vienna; the Law Research Seminar at the University of Cambridge and department colloquia at Leeds University, the Central European University, and Manchester Metropolitan University; the Research School of Social Science at Australian National University; and Aarhus University. I am grateful to the audiences for their comments. I am especially grateful to two anonymous referees and the associate editors for *Ethics*, whose comments significantly improved the paper. Thanks also to Christian Barry, Cécile Fabre, Nic Southwood, and James Wiloughby for helpful comments and discussion. Special thanks to Seth Lazar and Jonathan Parry for helpful comments and (multiple) discussions. Work on this paper has been supported by a Wallenberg Academy Fellowship from the Knut and Alice Wallenberg Foundation (grant no. 1521101).

1. Clearly, varying the cost to the rescuer can affect what she ought to do. I set aside those complications here, focusing on cases in which rescuing imposes no cost on the rescuer whether she minimizes harm or not.

Ethics 129 (January 2019): 204–229

© 2019 by The University of Chicago. All rights reserved. 0014-1704/2019/12902-0002\$10.00

preferentially saves if she prevents someone from losing one leg rather than preventing someone else from losing two legs, when preventing the loss of two legs would have been no costlier for the rescuer.

Plausibly, preferential saving is typically impermissible because one is ordinarily under a duty to minimize harm when rescuing, at least when the harms are set to befall equally innocent people.² It is this duty that underpins Effective Altruists' claim that one ought to give to the charities that do the most good, even when giving is supererogatory.³ It has also led some writers to argue that humanitarian wars are impermissible, since the resources that are spent trying to aid people through war would do considerably more good if they were spent on alleviating poverty.⁴ And yet, equally plausibly, one is not always required to minimize harm when rescuing. For example, it seems permissible for a mother to save her own child rather than two strangers, and for members of the British Armed Forces to save British citizens rather than an equal or somewhat greater number of foreigners.

Given that much hangs on the scope of the duty to minimize harm, we would do well to better understand its contours. This article contributes to that project by exploring the role of agreements in rendering preferential saving permissible. I defend the following thesis:⁵

Permissible Saving: Within proportionality limits, a prospective rescuer's agreement to preferentially save A grounds a duty to save A from an equal or lesser harm than faces B if, and only if, the prospective rescuer is under no duty to rescue (and no other conflicting duty) when she forms the agreement to save A, and the agreement is uncoerced. Contracts to preferentially save—understood as agreements in-

2. I assume throughout that all the prospective victims and rescuers are innocent in the relevant sense—that is, nobody is culpable, morally responsible, or negligent with respect to what threatens the prospective victims.

3. Theron Pummer, "Whether and Where to Give," *Philosophy and Public Affairs*, 44 (2017): 77–88; William MacAskill, *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference* (London: Faber & Faber, 2015). See also Iason Gabriel, "Effective Altruism and Its Critics," *Journal of Applied Philosophy* 34 (2017): 457–73. Of course, the claim that there is a duty to minimize harm when rescuing predates Effective Altruism. See, e.g., Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon, 1989), 16; Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), 225. Further, not everyone endorses this view; see, e.g., John Taurek, "Should the Numbers Count?," *Philosophy and Public Affairs* 6 (1977): 293–316.

4. Victor Tadros, "Unjust Wars Worth Fighting For," *Journal of Practical Ethics* 4 (2016): 52–80; Kieran Oberman, "War and Poverty," *Philosophical Studies* (forthcoming), <https://doi.org/10.1007/s11098-017-1012-4>.

5. These are merely sufficient conditions for preferential saving; there may be other grounds of preferential saving, such as agent-relative prerogatives. But they are necessary and sufficient conditions for permissible agreements to preferentially save.

volved payment—are valid only if A also satisfies her general duties to aid.⁶

By “within proportionality limits,” I mean to rule out the permissibility of preferential saving in cases in which there is very significant disparity between the harms facing the prospective victims. It is, for example, always impermissible to save Alice from a broken wrist at the cost of failing to save Betty from the loss of both her legs. We can judge these proportionality limits by drawing on more familiar intuitions about proportionate collateral harms. For example, even though the loss of two legs is worse than the loss of one, it is not sufficiently worse that a rescuer may inflict the loss of one leg on Alice to prevent the loss of two legs to Betty. In contrast, it would be proportionate for a rescuer to break Alice’s wrist as a side effect of preventing the loss of Betty’s legs. Taking this as a guide, then, we can say that if it is not proportionate to inflict harm *X* in order to prevent larger harm *Y*, then preventing harm *X* rather than larger harm *Y* does not violate the proportionality limit on preferential saving.⁷ In cases of significant disparity, the person facing the lesser harm is under a duty to accept this cost for the sake of preventing the much greater harm to someone else.⁸

Permissible Saving applies to cases in which two individuals face harms of different magnitude, but within the proportionality limits just described. Although I believe we can extend these arguments to cases in which the members of smaller and larger groups face the same type of harm, I do not defend that extension here. Those who are sympathetic to the view that one ought to save the greatest number can apply my arguments directly to such cases. Those who are skeptical can still endorse what I say about minimizing harm in cases of different harms to different people.⁹ I take my argument to be compatible with, and supportive of, the

6. By “satisfied her general duties to aid,” I mean that the person has given away (or made plans to give away) her resources to the extent that morality demands, such that her remaining resources are hers to use pursuing her own ends. This is not to say that she would lack any further duties to aid in specific rescue cases, where the cost would be (say) some physical harm to herself.

7. There are doing and allowing differences here: other things being equal, I may cause less harm than I may allow, and thus may cause less collateral damage to prevent two broken legs than I may fail to prevent when saving someone from two broken legs. Hence, my claim is that preferential saving would be permissible in at least these cases. I set aside here the difficult question of precisely how we should judge proportionality.

8. This need not entail that she is under a duty to inflict such harm on herself: it could be permissible for a bystander to divert a trolley from two to one, causing the one to lose a foot, even if the one would not be required to inflict the loss of her foot on herself for the sake of the two. But since our interest here is in what third-party rescuers may do, the collateral damage standard seems appropriate.

9. As Michael Otsuka points out in “Skepticism about Saving the Greater Number,” *Philosophy and Public Affairs* 32 (2004): 414.

view that it is typically wrong to fail to minimize harm when aiding, at least when one can prevent serious harm at reasonable cost to oneself. My interest is in the atypical cases in which preferential saving is permissible.¹⁰

The conditions that the prospective rescuer is under no other conflicting duty and that the agreement is uncoerced are intended to capture familiar ways in which an agreement might be rendered invalid. For example, my agreement to meet you for lunch might be rendered invalid by my earlier promise, and resultant conflicting duty, to feed my neighbor's cat at lunchtime. Similarly, an agreement to preferentially save could be rendered invalid if the content of the agreement conflicts with some independent duty the prospective rescuer has. For example, one may not agree to serve as A's lifeguard for the afternoon if going to the beach will prevent one from fulfilling one's duty to collect one's child from school. I shall assume that these general constraints on valid agreements straightforwardly apply to agreements to preferentially save, and I shall say no more about them here.

Section II makes some preliminary remarks about the duty to rescue and the duty to minimize harm. Section III outlines the familiar idea that agent-relative prerogatives can render preferential saving permissible. I grant this in order to focus on a different claim—namely, that agent-relative prerogatives can be outsourced to third parties by forming appropriate agreements with them, and that these third parties may then engage in preferential saving on the outsourcing agent's behalf. Versions of this claim have been defended by Seth Lazar, Cécile Fabre, and, more cautiously, Jonathan Quong.¹¹ In Section IV, I argue that the notion of outsourcing is incompatible with any plausibly robust duty to minimize harm when rescuing. Moreover, its proponents fail to show how an agree-

10. Following others, including those whose views I criticize here, I'll assume that the duty to minimize harm is a presumptive duty to which there are exceptions (we can think of the duty as applying when all other things are equal, and this article as giving an account of one way in which things might not be equal). We might instead think that there is no presumptive duty, but rather differences in permissibility depending on whether one is saving strangers, saving relatives, saving clients, etc. I doubt that anything substantive turns on which of these models we adopt.

11. Seth Lazar, "Associative Duties and the Ethics of Killing in War," *Journal of Practical Ethics* 1 (2013): 3–48; Seth Lazar, "The Justification of Associative Duties," *Journal of Moral Philosophy* 11 (2014): 28–55; Seth Lazar, "Authorisation and the Morality of War," *Australasian Journal of Philosophy* 94 (2016): 211–26; Cécile Fabre, *Cosmopolitan War* (Oxford: Oxford University Press, 2012), 61–63; Cécile Fabre, "Permissible Rescue Killings," *Proceedings of the Aristotelian Society* 109 (2009): 149–64, 157–58; Jonathan Quong, "Agent-Relative Prerogatives to Do Harm," *Criminal Law and Philosophy* 10 (2016): 815–29. Quong concedes that his view—that agent-relative prerogatives can only be outsourced to those with whom one has a special relationship—stands in need of elaboration and defense, and thus I do not address it here.

ment between an outsourcing agent and a recipient agent could vitiate other people's claims to be rescued. However, the failings of the Outsourcing View are instructive since they also give us reason to doubt the widely accepted view that one may agree to preferentially save by, for example, contracting to serve as someone's bodyguard.

In Section V, I argue that agents may form agreements to preferentially save only when, at the time of forming the agreement, nobody has a claim that the rescuer save her instead. In Section VI, I argue that allowing a rescuer to act on an agreement to preferentially save is justified by the moral reasons we have to respect the agreements that agents form as a means of shaping their lives in accordance with their own ends. This includes respecting their entitlements to benefits that they secure for themselves through agreements with others. Preferential saving is permissible when the rescue is a benefit to which the less imperiled person is entitled.

My account thus proposes a principled way to reconcile the plausibility of a duty to minimize harm with the broader nonconsequentialist project. Only if we grant agents claims to the benefits of their agreements do we adequately accommodate the concern that agents be permitted to sometimes pursue their own ends at the expense of minimizing harm. Thus, I argue, rescuers are not merely permitted but also required to act on valid agreements to preferentially save. I then consider the scope of agreements to preferentially save, including the implications of my account for agreements to enact preferential supererogatory rescues.

In Section VII, I anticipate and reject the objection that agreements to preferentially save can be valid only against a backdrop of an initially fair distribution of resources. Section VIII concludes.

II. DUTIES TO RESCUE

I assume that one has a duty to rescue at least in cases in which one can prevent serious harm to others at little cost to oneself, as in *Pond*:¹²

Pond: Alice is drowning in a shallow pond. Passer-By can save her, but doing so will involve getting his shoes muddy.

The duty to rescue is owed by the rescuer to the prospective victim, who has a claim against the rescuer that he save her. Passer-By would wrong Alice if he refused to muddy his shoes by rescuing her. I assume here that the duty to rescue is evidential—that is, Passer-By is not under a duty to

12. See Peter Singer, "Famine, Affluence and Morality," *Philosophy and Public Affairs* 1 (1972): 229–43, for a classic defense of this view.

rescue Alice if he has no reason to think that she is in danger, or that he is able to save her.¹³

I propose that, ordinarily, one has a duty to minimize harm when rescuing. Consider *Busy Pond*:

Busy Pond: Alice is being chased by a tiny piranha that will bite off one of her fingers. Betty is drowning in the same pond. Passer-By can save either at no cost to himself, but he cannot save both.

I'll assume that the loss of a finger is a sufficiently serious harm that it usually triggers a duty to rescue in those who are in a position to prevent it at comparatively low cost to themselves. Passer-By is certainly under a duty to rescue in *Busy Pond*. But he cannot discharge that duty by rescuing Alice. Discharging the duty to rescue requires more than simply preventing some harm to someone, even if the harm one thereby prevents would, ordinarily, trigger a duty to rescue. Rather, Passer-By must rescue in a way that minimizes harm. After all, the duty to rescue just is a duty to prevent harm. It's hard to see how such a duty could be discharged by intentionally preventing less harm, when preventing more harm is equally onerous (assuming, as we are here, that all the prospective victims are relevantly innocent). This seems true at least in cases where one can prevent either greater or lesser harm to the same person, or greater or lesser harms to different people.¹⁴ If I can, at equal cost, pull either a person's hand or both her legs out of the path of an oncoming train, I do not plausibly satisfy my duty to rescue if I save her hand.

When a rescuer is under a duty to minimize harm, he owes a duty of rescue to the prospective victim of the greatest harm. Passer-By would wrong Betty if he failed to save her life in order to prevent the loss of Alice's finger. In the event that Alice and Betty face threats of equal harm, Passer-By should plausibly toss a coin to give each a fair chance of rescue.¹⁵

13. It is a standard background assumption of discussions of the duty to rescue that potential rescuers know that someone is in peril and that they may be able to save her (this evidence constraint is implicit in the idea that one could "easily rescue"). The literature on rescuing is vast, but for influential discussions that make this assumption see, e.g., Singer, "Famine, Affluence and Morality"; Peter Unger, *Living High and Letting Die: Our Illusion of Innocence* (New York: Oxford University Press, 1996); F. M. Kamm, "Does Distance Matter Morally to the Duty to Rescue?," *Law and Philosophy* 19 (2000): 655–81.

14. As above, some people might think that things are different when we can prevent the same harm to differently sized groups of people.

15. Even if both are strangers to Passer-By, we have enough evidence about people's biases—implicit or otherwise—to want him to engage in a fair decision-making procedure if he can.

Several writers have argued that the duty to minimize harm holds even in cases of supererogatory rescues.¹⁶ This claim strikes some people as peculiar: how could one be required to save in a particular way if one is not required to save at all? But consider *Foot*:

Foot: Passer-By can save a hundred lives at the cost of his foot. However, Passer-By is willing to save only fifty people out of the imperiled one hundred. Saving the fifty will also cause the loss of his foot.¹⁷

Assume that Passer-By is not required to save anyone at the cost of his foot. Nevertheless, as Victor Tadros argues, he clearly acts impermissibly if he undertakes the rescue but does not save the most lives that he can.¹⁸ If he is willing to suffer loss of his foot for the fifty, he has a costless opportunity to save a further fifty lives. He is now under a conditional duty to save the hundred. Of course, Passer-By in *Foot* need not choose between saving fifty lives and saving a hundred lives; he can save everyone. But the fact that if he elects to bear the supererogatory cost he must save the hundred shows that supererogatory acts are not generally exempt from the duty to minimize harm. Theron Pummer also argues that even if one is not required to save, one ought not to save in a way that prevents less harm when one could prevent more harm at equal or only a slightly greater cost to oneself.¹⁹

III. AGENT-RELATIVE PREROGATIVES

The foregoing suggests that preferential saving is often prohibited. But there seem to be exceptions to this prohibition. Many people believe that one may sometimes preferentially save either oneself or a loved one. Consider *Crocodiles*:

Crocodiles: Alice is being chased by a small crocodile that will bite off one of her legs. Betty is being chased by a big crocodile that will bite off both her legs. Alice can quickly swim to shore, saving her own leg, or help Betty to get to the shore, saving both of Betty's legs at the cost of losing her own leg.

Plausibly, Alice is permitted to save herself rather than help Betty, even though doing so would not minimize harm overall. One prominent ex-

16. See, e.g., Pummer, "Whether and Where to Give"; MacAskill, *Doing Good Better*.

17. Based on a case in Tadros, "Unjust Wars Worth Fighting For," 53.

18. Tadros, "Unjust Wars Worth Fighting For," 53.

19. Pummer, "Whether and Where to Give," 93.

planation of why Alice may preferentially save herself is that she has an agent-relative prerogative to weight her own interests somewhat more heavily than other people's interests. Similarly, Alice's mother plausibly may (and perhaps must) save Alice rather than Betty, given the extra weight she should accord Alice's interests in light of their valuable relationship.²⁰

On nonconsequentialist views, agent-relative prerogatives reflect the intrinsic value of pursuing our own goals and projects and of developing relationships with other people, both of which rely on being able to prefer our own and special others' interests at least some of the time.²¹ Not all nonconsequentialists believe that there are agent-relative prerogatives of this sort, and I offer no defense of them here. Rather, I grant that we have such prerogatives and focus on their role in justifying preferential saving.

According to their proponents, agent-relative prerogatives function as justifications, rather than excuses. They render permissible actions that would ordinarily be wrong, including actions against which other people would ordinarily have claims. Invoking these prerogatives in the context of preferential saving, then, is an acknowledgment of the general duty to minimize harm in the absence of special justification. Betty, for example, has a claim against other people that they save her rather than Alice, since she will suffer twice the harm that will befall Alice. But she lacks such a claim against Alice, given how costly rescuing her would be for Alice, once we include the extra weight that Alice may attach to her own interests. Our claims to be rescued are typically understood as cost sensitive in this way: if saving a person is too costly, she simply lacks a claim to be saved.²²

20. Jeff McMahan, "The Limits of National Partiality," in *The Morality of Nationalism*, ed. R. McKim and J. McMahan (Oxford: Oxford University Press, 1997), 107–38, 111; Lazar, "Associative Duties." There is disagreement about how special relationships have to be to generate agent-relative prerogatives (see, e.g., Thomas Hurka, "Proportionality in the Morality of War," *Philosophy and Public Affairs* 33 (2005): 34–66; David Rodin, "The Myth of National Self-Defence," in *The Morality of Defensive War*, ed. C. Fabre and S. Lazar (Oxford: Oxford University Press, 2014), 69–89. I take no stance on that issue here.

21. See, e.g., Samuel Scheffler, *The Rejection of Consequentialism* (New York: Oxford University Press, 1982); Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press 1989). There are also consequentialist accounts of agent-relative prerogatives. See, e.g., Peter Railton, "Alienation, Consequentialism, and the Demands of Morality," *Philosophy and Public Affairs* 13 (1984): 134–71; Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101 (1991): 461–82.

22. See, e.g., Joanna M. Firth and Jonathan Quong, "Necessity, Moral Liability and Defensive Harm," *Law and Philosophy* 31 (2012): 673–701, 695; Jonathan Quong, "Rights against Harm," *Supplement to the Proceedings of The Aristotelian Society* 89 (2015): 249–66, 252. We might also understand agent-relative prerogatives as justifying infringements of other people's rights. However, this looks implausible given how we normally treat justified infringements of rights: e.g., it seems unlikely that Alice owes Betty compensation or an apology.

IV. OUTSOURCING AGENT-RELATIVE PREROGATIVES

A. *The Outsourcing View*

Consider *Help*:

Help: Alice is being chased by a small crocodile who will bite off one of her legs. Betty is being chased by a big crocodile who will bite off both her legs. Passer-By is able to save either Alice or Betty at no risk to himself. Alice asks Passer-By to save her instead of Betty.

Is Passer-By permitted to engage in preferential saving at Alice's behest? According to the Outsourcing View, one can authorize otherwise impartial third parties to act on one's agent-relative prerogatives, thereby making it permissible for those third parties to engage in preferential saving on one's behalf. Note that outsourcing is distinct from the idea that, for example, a parent might permissibly favor her own child.²³ In those cases, the parent has her own agent-relative prerogatives to weight her child's interests more heavily, generated by her role in their valuable relationship. Outsourcing, in contrast, is supposed to take place between an agent who has an agent-relative prerogative and an agent who otherwise lacks such a prerogative. Proponents of outsourcing grant the general duty to minimize harm when saving: the role of outsourcing is precisely to defeat or vitiate this duty.

Lazar argues that outsourced agent-relative prerogatives help to justify combatants' actions in war.²⁴ A state's combatants act not only on their own agent-relative prerogatives to protect their loved ones but also on the agent-relative prerogatives of all their citizens to do likewise, which affects the proportionality of the harms the combatants inflict.²⁵ Fabre believes that the capacity to outsource agent-relative prerogatives is part of an agent's defensive rights. If an agent's interest in her survival is sufficiently important to be protected by a right to defend herself, then, Fabre argues, "surely it is important enough on principle to be protected by a *prima facie* power to transfer that right to third parties. To claim otherwise is to impose an arbitrary restriction on [the agent's] ability to promote this fundamental interest of hers."²⁶

At first glance, the claim that one can outsource agent-relative prerogatives looks conceptually confused. The reasons underpinning these prerogatives and duties are, after all, grounded in one's special stake in

23. See, e.g., Quong, "Agent-Relative Prerogatives," 824–25.

24. Much of Lazar's discussion focuses on agent-relative duties, rather than prerogatives, but this doesn't affect my arguments here.

25. Lazar, "Associative Duties," esp. 30–33. Adil Ahmad Haque tentatively endorses Lazar's suggestion in *Law and Morality at War* (Oxford: Oxford University Press, 2017), 80.

26. Fabre, *Cosmopolitan War*, 62–63. See also Fabre, "Permissible Rescue Killings," 158.

one's own welfare, or in one's valuable relationships. It is unclear how such reasons could be transferred to someone who lacks a similar stake in one's welfare or relationships. Perhaps a better way to understand the view is as an expansive account of the possible bases of agent-relative prerogatives. On this interpretation, if Passer-By agrees to act on Alice's behalf, that agreement itself generates an agent-relative prerogative for Passer-By to save Alice. We can call this the Generating View. However, as I shall argue, the Generating View fares no better against important objections to the Outsourcing View.

B. *Failing to Aid and Preventing Saving*

First, though, we should note that the Outsourcing View assumes a fairly permissive stance toward the justificatory scope of agent-relative prerogatives. I have granted that these prerogatives can justify failing to save, which is in line with the orthodox view of agent-relative prerogatives.²⁷ However, for the Outsourcing View to succeed, agent-relative prerogatives must also be able to justify preventing other people from being saved. Contrast *Parental Rescue* with *Dinghy*:

Parental Rescue: Alice is being chased by a small crocodile who will bite off one of her legs. Betty is being chased by a big crocodile who will bite off both her legs. Mother is able to save either Alice or Betty at no risk to herself. Alice is Mother's daughter. Betty is a stranger to Mother.

Dinghy: Alice is being chased by a small crocodile who will bite off one of her legs. Betty is being chased by a big crocodile who will bite off both her legs. Alice is Mother's daughter. Betty is a stranger to Mother. Betty is heading toward a small dinghy that will enable her to safely get to shore. Using a fishing rod, Mother can pull the dinghy toward Alice, saving her. Betty will then lose both her legs.

Parental Rescue is a case of failing to aid. Mother does not make Betty any worse off than she would have been in Mother's absence. She merely denies Betty the opportunity to benefit from her presence. *Dinghy*, in contrast, is not a case of failing to aid. Mother's means of saving Alice—moving the dinghy—deprives Betty of access to a resource that she could have used to save herself, and to which she had at least an equal and plausibly a greater claim. Mother prevents Betty from being saved. This is no less true if Mother persuades Passer-By to save Alice instead of saving Betty. This

27. See, e.g., Samuel Scheffler, *Boundaries and Allegiances: Problems of Justice and Responsibility in Liberal Thought* (Oxford: Oxford University Press, 2002); McMahan, "Limits of National Partiality," 132–33; Thomas Pogge, "Cosmopolitanism: A Defense," *Critical Review of International Social and Political Philosophy* 5 (2010): 86–91.

too deprives Betty of access to a resource—Passer-By's assistance—to which Betty had at least an equal and plausibly a greater claim. Once outsourced, prerogatives to fail to save necessarily become instances of preventing others from being saved.

If the justificatory scope of agent-relative prerogatives is limited to failing to aid, and preventing someone from being saved is not morally equivalent to failing to aid her, agent-relative prerogatives cannot be outsourced. And even if the justificatory scope of agent-relative prerogatives is not limited to failing to aid, preventing someone from being saved is plausibly harder to justify than failing to save her.²⁸ This will affect what it is proportionate for rescuers to do: it will not follow from the fact that Mother may save Alice that she may enlist Passer-By to save Alice, thereby preventing Betty from being saved.

However, for those who believe that agent-relative prerogatives have a wide justificatory scope and can permit even straightforward harming, this observation will not pose an insurmountable problem.²⁹ They can concede the differences in proportionality without undermining the notion of outsourcing altogether. Let us consider, then, two independent difficulties with the notion of outsourcing. The first is that outsourcing is incompatible with the general duty to minimize harm when rescuing those with whom one has no special relationship. The second is that we lack a coherent explanation of how outsourcing can defeat the claims of others to be rescued.

C. *The Duty to Minimize Harm*

If all it takes to defeat Passer-By's obligation to minimize harm in *Help* is that he agrees to save Alice instead, the duty to minimize harm is implausibly weak. Moreover, insofar as we can assume that Alice wants to be rescued, it is unclear why there must be an explicit agreement between them before Passer-By may act on Alice's agent-relative prerogatives. Indeed, Fabre argues that outsourcing does not require explicit authorization: a rescuer may act on a victim's agent-relative prerogatives if she rea-

28. On how we ought to categorize preventing someone from being saved, see Mathew Hanser, "Killing, Letting Die and Preventing People from Being Saved," *Utilitas* 11 (1999): 277–95; Jason Hanna, "Enabling Harm, Doing Harm and Undoing One's Own Behavior," *Ethics* 126 (2015): 68–90; Samuel C. Rickless, "The Moral Status of Enabling Harm," *Pacific Philosophical Quarterly* 92 (2011): 66–86; Haque, *Law and Morality at War*, 79.

29. Several writers argue that agent-relative prerogatives can justify harming. See Jonathan Quong, "Killing in Self-Defense," *Ethics* 119 (2009): 507–37; Quong, "Agent-Relative Prerogatives"; Nancy A. Davis, "Abortion and Self-Defence," *Philosophy and Public Affairs* 13 (1984): 175–207; Lazar, "Associative Duties"; Lazar, "Justification of Associative Duties"; Lazar, "Authorisation and the Morality of War"; Fabre, *Cosmopolitan War*, 61; Fabre, "Permissible Rescue Killings."

sonably believes that the victim would authorize her to do so if she were able.³⁰

But if Passer-By may act on Alice's agent-relative prerogatives provided that he is reasonably confident that Alice wants to be rescued, this seems to do away with the duty to minimize harm altogether. All it would take is for Alice to wave her arms in a bid to draw Passer-By's attention, and Passer-By would be somehow released from his duty to rescue Betty, the victim of the greatest prospective harm.

This result has significant implications for the widely held view that, other things being equal, one ought to save a greater rather than smaller number of people when forced to choose between them. It entails that if (one reasonably believes that) the members of the smaller group want to be saved, one may agree to save them and thereby defeat the duty to save the greater number. But I doubt that even those who believe that one may save the smaller number believe that this permissibility turns solely on whether the members of the smaller group want to be saved and the rescuer agrees to save them.

The only way for the Outsourcing View to leave space for a duty to minimize harm is to restrict the range of possible recipients of outsourced prerogatives. On Fabre's account, authorization can be ad hoc—that is, one can simply pick anyone and ask her to act on one's behalf.³¹ Requiring more robust relationships would not answer Fabre's concern about placing arbitrary restrictions on an agent's capacity to protect her interests, since whether the agent can outsource will depend on how well she happens to know the nearby third parties. Lazar, in contrast, thinks that it is unlikely that outsourcing can be ad hoc in this way. Outsourcing requires an "appropriate connection" between the parties: either a preexisting relationship or a "forward-looking compact." If Passer-By is a stranger who will disappear from Alice's life after the rescue, Lazar suggests that he cannot be the recipient of outsourced prerogatives. But if, at the time of rescue, Passer-By agrees to be Alice's guardian, then this newly formed "standing relationship" would make outsourcing possible.³²

Lazar's more restrictive view would leave some space for a duty to minimize harm when rescuing, since unsuitable recipients of outsourced prerogatives will still be subject to such a duty. I suspect, however, that Lazar will struggle to sustain the more restrictive view. If we want to avoid outsourcing's being ad hoc, Passer-By will need to be more than Alice's passing acquaintance. But for the notion of outsourcing to get off the ground, the relationship between the parties must be sufficiently weak

30. Fabre, "Permissible Rescue Killings," 160.

31. Fabre, *Cosmopolitan War*, 63.

32. Lazar, "Authorisation and the Morality of War," 13.

that the third party has no agent-relative prerogatives of her own regarding Alice. If Passer-By is already Alice's godfather, for example, it seems plausible that he simply has his own agent-relative prerogative to save Alice, and there would be no role for outsourcing to play.

Even if we could find the middle ground of "appropriate" relationships that are sufficiently robust to avoid the charge of being ad hoc, and yet sufficiently weak to avoid generating prerogatives of their own, Lazar's inclusion of forward-looking relationships is still problematic. All that is needed to form such a relationship is the agreement of the relevant parties, who may be strangers. It is not as if, as Alice drowns in the pond, she will carefully assess Passer-By's suitability as a source of guidance for the rest of her life before asking Passer-By to become her guardian. Alice will be willing to outsource to whoever happens to be passing by. And anyone who agrees under such circumstances to serve as a guardian presumably would have done so for any child: as a stranger, it is not as if Passer-By has some special affection for Alice. All this looks unavoidably ad hoc. If outsourcing is possible, then, it looks likely that it must be possible in the ad hoc way described above. And this makes outsourcing incompatible with any plausibly robust duty to minimize harm when rescuing strangers. The same objection applies to the Generating View: if ad hoc agreements generate agent-relative prerogatives to preferentially save, we undercut the idea of a duty to minimize harm when saving.

D. Defeating Claims

Say that we could find a way to avoid ad hoc outsourcing, leaving some scope for a duty to minimize harm. It is still mysterious how an agreement between Alice and Passer-By could defeat Betty's claim to be saved by Passer-By. Alice, by hypothesis, may fail to save Betty on the grounds of her prerogative. But since Betty retains a claim to be rescued against anyone who lacks such a prerogative, the challenge for proponents of outsourcing is to show how Passer-By could be justified in making an agreement with Alice not to save Betty. Passer-By cannot invoke Alice's agent-relative prerogatives to do this. His putative right to invoke these prerogatives is an upshot of the agreement. It cannot, therefore, serve as a justification for making the agreement.³³

This looks like a fundamental, and insurmountable, problem with the notion of outsourcing agent-relative prerogatives. That saving requires the special justification provided by an agent-relative prerogative entails that there is someone else who has a claim against the rescuer—it

33. This is also true, of course, in a case where Alice and Betty face threats of equal harm and each has a claim to a fair chance of rescue. If Betty has a claim that Passer-By toss a coin, this claim cannot be defeated by Passer-By's agreeing to simply save Alice.

is because Betty has a claim to be saved that Passer-By needs to justify failing to saving her. Thus, there cannot be any cases in which one needs to outsource an agent-relative prerogative to a rescuer, and yet the rescuer has no duty to save someone else. And if there is a duty to save someone else, the rescuer cannot escape that duty by simply agreeing not to save that person. This objection tells equally against the Generating View, according to which agreements themselves can be a possible basis of agent-relative prerogatives. One would need to be able to act on these alleged prerogatives in order to justify forming the agreement, which is incoherent. Preferential saving cannot be justified merely on the grounds that one has agreed to preferentially save, whether such an agreement is understood as a means of outsourcing agent-relative prerogatives or as generating new agent-relative prerogatives.

So much for outsourcing as a ground of preferential saving, then. But identifying the shortcomings of the Outsourcing View is instructive, since they reveal constraints on agreements to preferentially save more generally. Consider *Hired Help*:

Hired Help: Alice is being chased by a small crocodile who will bite off one of her legs. Betty is being chased by a big crocodile who will bite off both her legs. Alice's hired protector, Bodyguard, is in a position to save either Alice or Betty at no risk to himself.

This might strike many of us as a case in which preferential saving is obviously permissible. We needn't invoke any thoughts about agent-relative prerogatives to explain this apparent permissibility (we can assume that, ignoring their contract, Alice and Bodyguard do not have a valuable relationship of the sort that could generate such prerogatives). Rather, it seems plausible that Bodyguard may rescue Alice in virtue of the contract that exists between them. If the fairly common practice of serving as a bodyguard is permissible, it looks as if we have identified another ground of permissible preferential saving in addition to agent-relative prerogatives based on valuable relationships.

But we have just seen that mere agreements to preferentially save cannot make such saving permissible. Contracts are vulnerable to the same objections. Otherwise, we are committed to the following implausible view: when one happens across strangers who need rescuing, one is under a duty to minimize harm. But if one is offered money or some other incentive to fail to minimize harm, one can somehow release oneself from that duty in order to profit or otherwise benefit oneself. Imagine that, in *Help*, Passer-By is about to rescue Betty. Seeing this, Alice dangles a blank check out of the water to incentivize Passer-By to save her instead, asking him to be her bodyguard. It's hard to see how this would be much more than a bribe to induce Passer-By to ignore his moral duties to Betty in

Alice's favor. This cannot be a means by which Betty ceases to have a claim against Passer-By.

V. DUTIES TO RESCUE AND AGREEMENTS TO PREFERENTIALLY SAVE

In what follows, I argue that agreements to preferentially save are valid only if, at the time of forming the agreement, nobody has a claim that the rescuer save her instead. The rescuer may act on the agreement only when preferential saving is proportionate, in the sense explained in Section I. I then argue that a further condition applies to contracts to save—that is, to agreements to save that involve some form of payment, such as hiring a bodyguard. Contracts to preferentially save are valid only if the client also satisfies her general duties to aid. When these conditions are satisfied, Betty has no claim to be saved that is violated if a rescuer preferentially saves Alice.

My account is partly motivated by the observation that preferential saving is a broad and familiar phenomenon. We often involve others in our plans to further our own ends rather than minimize harm. For example, consider *Alarm*:

Alarm: After a spate of violent break-ins in her neighborhood, Alice hires a security company to install an intruder alarm in her house. If the alarm is triggered, the company will dispatch a security guard. Betty is Alice's neighbor. The position of Betty's house means that she is more likely to be the victim of a break-in.

It would be impermissible for the security company to decide to install the alarm in Betty's house, even if this would minimize harm. And when Alice's alarm goes off, and it turns out that both Betty and Alice have intruders, the security guard should not protect Betty rather than Alice, even if, for example, Betty's intruder will break both her legs and Alice's intruder will break only one of her legs. Nor should he toss a coin in the event that Alice and Betty face identical threats, giving each a fair chance of rescue. Rather, he should simply rescue Alice. This seems true even though the harm that Betty faces ordinarily suffices to generate a duty to rescue, and even if the guard can prevent this harm at no risk to himself.

These very plausible verdicts could not be correct if agreements made no moral difference to who should be saved. When a rescuer encounters strangers facing identical threats, he should toss a coin to give each a fair chance of rescue. That the guard may simply save Alice is good evidence that agreements to preferentially save can be morally relevant to preferential saving. The task before us, then, is to explain why these fairly

familiar instances of preferential saving can be permissible, but Passer-By may not save Alice in *Help*.

Consider, first, an example of an agreement (rather than a contract) to preferentially save:

Hike. Alice is visiting a mountainous region for a conference. She plans to go hiking, but is nervous of falling. She explains this to Hiker, whom she met at the conference. Hiker offers to come with her, promising to preferentially save her if she gets into trouble. Alice gratefully accepts. While on the hike, Alice slips to where she will fall and will break her leg unless Hiker saves her. Nearby, Betty faces breaking both her legs unless Hiker saves her instead of saving Alice.

Plausibly, Hiker may save Alice rather than abandon her to prevent greater harm to Betty. The permissibility of preferential saving in *Hike* is not grounded in an agent-relative prerogative: we've stipulated that Hiker and Alice are only vague acquaintances. Nor is it grounded in the mere value of agreements: as I have argued above, merely agreeing to prefer someone else's interests cannot make preferential saving permissible. Rather, as I shall argue, it is grounded in the fact that Hiker is under no duty to rescue when he forms the agreement to preferentially save, and thus he violates no claim of Betty's in saving Alice. Since saving Alice meets the proportionality constraint described in Section I, saving Alice is permissible.

At the time Alice and Hiker form their agreement, Betty has a merely potential claim to be rescued. We can see this by noticing that, for example, it would be permissible for Hiker not to go hiking at all. He need not offer to accompany Alice: he may stay at home, and he would not violate anyone's claims in doing so.

Recall from Section II that we are assuming that the duty to save is evidential, such that a rescuer is under no duty to save if his evidence is either that nobody needs rescuing or that he is unable to rescue at reasonable cost to himself. We can thus explain Hiker's lack of a duty to save in one of two ways. First, we might say that Hiker is under no duty to save (and hence has no duty to go hiking) because he has no reason to think that anyone needs saving. Even if Hiker would be obliged to go hiking if his evidence suggested that he could thereby prevent harm at reasonable cost to himself, he is not so obliged when his evidence suggests otherwise.

Or, we might say that Hiker's evidence is that he is not in a position to rescue at reasonable cost to himself. It may be well known that hikers occasionally fall on the mountain, but we do not typically demand that people go hiking on the off chance that they might be able to rescue people. Even when we know that we might be able to save in such cases, the

low probability makes it too demanding to insist that people make themselves available to serve as rescuers. As long as the probability that he can rescue at reasonable cost is sufficiently low that Hiker is under no duty to rescue, his forming an agreement to preferentially save Alice does not conflict with anyone else's claim to be rescued.

Either explanation entails that Hiker is under no duty to save Betty when he agrees to preferentially save Alice.³⁴ Rather, his agreeing to save anyone is supererogatory. The same is true of the members of the security company in *Alarm*. Prior to forming the relevant agreements, the company is under no obligation to install alarms in people's houses, including Betty's. Since they are under no duty to rescue, their agreements with Alice cannot conflict with such a duty. Contrast these cases with *Help*. When Alice asks him to save her, Passer-By knows that he can rescue either Alice or Betty at reasonable cost to himself. He may not decide to simply keep on walking past the lake, saving nobody. Rather, at the point at which Alice asks Passer-By to save her, Passer-By already has a duty to save Betty. Thus, he cannot form agreements that conflict with that duty, such as an agreement to save Alice instead. This is why preferential saving is impermissible in *Help*.

The foregoing holds that agreements to preferentially save are valid only if, at the time of forming the agreement, the rescuer is under no obligation to save anyone. But, as we saw above, even supererogatory rescuers can be required to minimize harm. Thus, the fact that Hiker agrees to look after Alice might permit him to do so only insofar as nobody else faces a greater harm. In other words, their agreement shows only that he ought to help Alice if she falls; it doesn't secure the stronger, but plausible, claim that he may do so at the cost of failing to save Betty from a somewhat greater harm. Likewise in *Alarm*. We thus need a more substantial explanation of why a rescuer might be permitted to fail to minimize harm.

VI. JUSTIFYING PREFERENTIAL SAVING

A. *Agreements and Entitlements*

In *Hike*, Betty will suffer more harm than Alice, and Hiker is in a position to rescue Betty without excessive cost to himself or bystanders. How, then, could an agreement to save Alice prevent Betty's having a claim to be res-

34. We might think that the duty to rescue is a kind of general background duty: that one is always under a duty to rescue whenever one can do so at reasonable cost to oneself. But this kind of general duty cannot plausibly correlate to particular people's having particular claims against the potential rescuer, such that Hiker fails in his duty and violates those claims by deciding not to go hiking. Thus, this account of the duty to rescue does not undercut my argument here, which requires only that there is nobody else who has a claim that Hiker rescue her at the time he forms the agreement with Alice.

cued? The answer lies within the nonconsequentialist framework that we are assuming here, according to which an agent need not spend her life minimizing harm or promoting the good. Such a framework gives an agent considerable scope to shape her own life and pursue her own goals, even when these goals are suboptimific impartially considered. Even those who are skeptical about agent-relative prerogatives can agree that we have reason to let people pursue their own ends and to give moral weight to the plans they have made.

The permission to pursue her own projects and shape her own life includes allowing an agent to form agreements that further her own ends. Once Alice satisfies her duties to aid, she need not make plans or agreements for Betty's benefit. On the contrary: she may make agreements for her own sake, directing her surplus resources to her own projects. This is the crux of the nonconsequentialist view. Of course, Alice may not form agreements that conflict with Betty's claims. She may not, for example, hire Hitman to kill Betty, even if this would secure her Betty's job. But Alice's forming an agreement that Hiker will rescue her does not conflict with Betty's claims. Hiker's capacity to rescue Betty is a benefit of Alice's plan to ensure rescue for herself. Nonconsequentialists cannot plausibly grant that Alice is allowed to form agreements to shape her own life and promote her own ends while denying that she is entitled to the benefits arising from those agreements. Securing benefits for ourselves is often the point of such agreements.

Recall that we are assuming in each of our cases that Alice would be under no duty to save Betty if she were able to do so: her agent-relative prerogative would, for example, permit her to save herself in *Hike* and, presumably, to use her own resources to this end. Demanding that the benefit of Alice's plans—including Hiker's capacity to rescue—be redirected to Betty is thus to comprehensively undermine the nonconsequentialist project, undercutting the limit on the duty to aid that protects the space in which agents may pursue their own ends. It's akin to taking a life preserver that Alice has brought to protect herself and giving it to Betty. If we agree that, in virtue of the harm she faces, Alice need not give the life preserver to Betty, I doubt we may demand that a rescuer take it from her, even to prevent greater harm to Betty.

Note that we cannot meet this concern by arguing that benefits and resources are to be redirected only when others are in urgent and greater need, as if such cases will be few and far between. It will nearly always be the case that others have such needs: this is precisely why nonconsequentialists insist on a space in which the needs of others do not trump an agent's claims. Nonconsequentialists hold that Alice is entitled to her surplus resources and benefits even when she does not need them to avert harm to herself. Given this, it is hardly plausible to requisition them when she does need them to avert harm to herself, as in *Hike*.

If Hiker's capacity to rescue is a good to which Alice is entitled, Hiker cannot save Betty without depriving Alice of a benefit to which she is entitled, and that she needs to prevent serious harm to herself. Insisting that Hiker minimize harm by saving Betty thus fails to appropriately respect Alice's right to an autonomous life, impermissibly transferring to Betty a benefit that belongs to Alice.

B. The Scope of Agreements to Preferentially Save

I have suggested that the reasons we have to respect agents' entitlements to the benefits of their agreements plausibly underpin the validity of agreements to preferentially save. Exactly how permissive this account of preferential saving turns out to be will depend on when one thinks Alice is entitled to the benefits of her agreements, which in turn depends on the correct general theory of ownership rights. Answering those questions is beyond the scope of this article, of course. However, we can still usefully consider whether particular examples of preferential saving will likely be justified by any plausible account of entitlement, in order to give us an idea of the probable scope of agreements to preferentially save.

In *Hike*, Hiker is in a position to save Betty only because of his agreement with Alice. Alice's claim to his help seems especially strong in this case. But it seems unlikely that Alice is entitled to the benefits of her agreements only when those benefits are counterfactually dependent on her agreements. For example, imagine that Betty in *Alarm* would have hired the same security company had she not known that Alice had hired them (perhaps she thinks that the company sign on Alice's house will be enough to deter thieves on the whole street, or perhaps she doubts Alice's judgment and assumes that any company Alice hires must be cheap and unreliable). That the guard's presence is not counterfactually dependent on Alice's agreements doesn't undermine Alice's claim to the guard's protection. It's enough to ground Alice's claim to be saved that the guard is in fact there because of their agreement.

Thus, it seems that preferential saving is permissible at least when a rescuer's agreement with Alice explains their actual presence, that is, when it features in their deliberations that result in their capacity to rescue. We can still think of the rescuer's presence as an upshot of the agreement in such cases—we might call this *explanatory* dependence. Compare, for example, how we can sensibly describe my presence at the ball as an upshot of my having been invited by Prince Charming even if, had Charming not invited me, Prince Harming would have invited me and I would have gone with him instead. The invitation from Prince Charming is what in fact features in, and explains, my decision to attend. My attendance is thus an upshot of his invitation, even though it's not counterfactually dependent on that invitation.

Granting this broader scope of preferential saving allows us to capture some other intuitively permissible cases. The fact that, for example, Bodyguard might well have attended a large public event even if he had not been going as Alice's bodyguard does not entail that he may not preferentially save Alice when he attends because of their agreement. His presence is still sensibly described as an upshot of their agreement in this case, just as my attendance at the ball is an upshot of Prince Charming's invitation.

We might plausibly grant an even broader account of when Alice is entitled to be saved. Consider *Beach*:

Beach: Before going into the sea, Alice hires Sunbather to preferentially save her. Sunbather was planning to stay on the beach all day anyway. His agreement to save Alice does not feature in his deliberations about whether to stay on the beach.

If we grant that Bodyguard may save Alice at a public event that he would have anyway attended, it might seem implausible to deny that Sunbather may preferentially save Alice in *Beach*. Sunbather is under no duty to rescue anyone else when Alice asks him to preferentially save her. Thus, if we think Alice can secure benefits for herself merely by agreement, even when the agreement doesn't explain the presence of the benefit, then it seems that preferential saving is also permissible in *Beach*. Our intuitions about the permissibility of preferential saving in *Beach* might be strengthened if we stipulate that Alice will go swimming only if Sunbather agrees to preferentially save her—that is, if her need for rescue is an upshot of their agreement, even if Sunbather's presence is not. But the foregoing suggests that this is not necessary to make preferential saving permissible: it's enough that Sunbather has no duty to save anyone else when he agrees to save Alice. (However, as I suggest below, these cases seem especially vulnerable to complaints about unfair distribution of resources.)

The above arguments suggest that a rescuer who has entered into a valid agreement to preferentially save is not only permitted but also required to preferentially save—that, since Alice has a claim to the benefit of her agreements, Hiker would wrong Alice if he decided to renege on their agreement in order to save Betty. This seems correct to me, and it is in line with the intuitive impermissibility of preferentially saving Betty in *Alarm*. The security guard is not merely permitted to protect Alice rather than Betty. He ought to do so: he would wrong Alice by fending off Betty's attacker instead of Alice's. The guard's capacity to rescue is a good that Alice has secured for herself, via an agreement she formed to protect her own interests. Failing to save Betty under these conditions does not violate any claim of Betty's, who had no claim that Alice form agreements or produce benefits for her sake. But failing to save Alice does violate her

claim to the benefits of her legitimate plans and is thus impermissible. That saving Alice is required also reflects the fact that the lack of a duty to save Betty is grounded in respect for Alice's, and not the rescuer's, autonomy. The rescuer's autonomy is appropriately respected by the fact that forming the agreement to preferentially save Alice is morally optional. He can easily avoid incurring an obligation to preferentially save by not entering into such agreements.

Grounding a permission to preferentially save in our obligations to respect agents' claims to the benefits of their agreements explains why a rescuer may preferentially save someone only if she faces the sort of harm from which he agreed to save her. That Hiker has agreed to preferentially save Alice if she falls while they are hiking does not permit him to preferentially save her if, for example, they are in a traffic accident as they drive to the mountain. Their agreement does not confer a claim to that good on Alice. Thus, respecting Alice's claim to the benefits of her agreements does not demand that we allow Hiker to preferentially save her from this sort of harm.

That we are trying to appropriately respect Alice's plans, and her use of her surplus resources, also explains why Alice is entitled to be rescued only when the rescue is the benefit of an agreement to preferentially save her, not just of any agreement she happens to make. Imagine that in a variation of *Help*, Alice's cleaner appears at the lake to return Alice's house keys, as per their earlier agreement. The fact that the cleaner's capacity to rescue is a result of his agreement with Alice does not make it permissible for the cleaner to preferentially save Alice. That an agreement results in the capacity to preferentially save matters only if the agreement was made to that end. Only then could the rescue be something to which Alice is entitled.

C. *Supererogatory Rescues*

Note that the argument advanced here does not undermine the view (that I endorsed in Sec. II) that those performing even supererogatory rescues can be under a duty to minimize harm. The reason why the most imperiled person in a supererogatory case has a merely potential claim to be rescued is that the rescue will impose such a high cost on the rescuer that it is too demanding to be required. We object to forcing someone to compromise their own interests to such a degree. But when a rescuer volunteers to bear that cost, the worry of overdemandingness disappears. The rescuer is not being forced to bear a cost. And if they are not prepared to bear the cost for the most imperiled person—to minimize harm—then they need not rescue at all. Our usual ground for denying the most imperiled person a claim to be rescued in a supererogation case—that the rescue would

force someone to bear an unacceptable cost—thus evaporates when the rescuer volunteers to bear that cost for the sake of preventing the harm.

But this is not true when the more imperiled person lacks a claim to be saved because the rescuer agreed to preferentially save someone else at a time when he was under no duty to rescue. The reason we permit Bodyguard to fail to save Betty in *Hired Help* is not that saving Betty is too costly to be required, but rather that Bodyguard formed an agreement to save Alice when he was under no duty to rescue anyone. This explanation holds in supererogatory cases as well. And, unlike the demandingness worry, this explanation persists even once the rescuer elects to bear the cost of rescue.

Note, too, that the view defended here permits a rescuer to form an agreement to save Alice at the time she is endangered, provided that saving Betty is supererogatory. This still does not threaten our verdict that it would be wrong for Passer-By to save fifty of a hundred lives in a case like *Foot*. In *Foot*, the saving of a hundred lives does not conflict with saving fifty lives; Passer-By saves only a subset of the people whom she could save at the same cost. But in a case in which a rescuer must choose between preventing either a lesser or a greater harm, and preventing either will involve bearing a supererogatory cost, the account developed here suggests that he is permitted to form, and act upon, an agreement to save the less imperiled person.

This result challenges strong versions of Effective Altruism that hold that one ought to minimize harm in supererogatory cases quite generally, since it holds that one can agree to prevent less harm (and that this makes preventing less harm permissible), provided that one is under no duty to prevent harm. And, since I've argued that a rescuer who has agreed to preferentially save ought to save Alice, we have identified another exception to the general rule that one must minimize harm when one has a duty to save (along with the plausible agent-relative exceptions that, e.g., parents ought to save their children rather than minimize harm).³⁵ However, my argument leaves intact the claim that insofar as we are under a duty to aid, we should typically direct our resources to where they will do the most good.

VII. CONTRACTS AND UNFAIRNESS

Contracts, I have stipulated, are agreements that involve some sort of payment. *Alarm* is but one example of the many ways in which we contract with others to preferentially save us, and it suggests that such contracts can be valid even if they are impartially suboptimific. And yet, we might think,

35. Whether a rescuer is under a duty to save Alice from an otherwise supererogatory rescue depends on the content of their agreement, of course. But he has a duty to save her when he can do so at reasonable cost to himself, at least.

such contracts are not merely suboptimific but also unfair. Perhaps Betty, unlike Alice, cannot afford to install a security system or hire a bodyguard. We might be uncomfortable with the result that those who are better off get rescued more often. We might, therefore, accept the moral force of contracts against a background assumption of an initially fair distribution of resources while being skeptical of their force when Betty unfairly lacks sufficient surplus resources to form the contracts that are available to Alice.

However, as described above, the nonconsequentialist framework that we are assuming grants not only that agents may form suboptimific agreements but also that they need not devote all of their resources to minimizing harm. If an agent satisfies her general duties to aid (e.g., through taxation or charitable giving), we should let her use her remaining resources to further her own ends.³⁶ Whatever such a framework says about unfair distributions of resources, it also holds that there is a limit to the resources that Alice is required to give to others or use for their sake.

This is, of course, compatible with various accounts of when one counts as satisfying one's duties to aid. My argument is simply that at the point at which Alice satisfies those duties, she may use any remaining resources to pursue her own ends. Indeed, this is surely part of what it means to say that an agent satisfies her general duties to aid: that she is now free to do other things with her time and money that do not minimize harm. Granting that one can have such surplus resources seems like a fundamental tenet of nonconsequentialism, since the alternative is that one must devote all one's resources to minimizing harm (at least, perhaps, once one has secured one's own basic interests). After all, whenever Alice uses her surplus resources to buy a theater ticket or a painting, she thereby fails to save the lives of people who are unfairly worse off than her, since she could have donated that money to a life-saving charity. Mere unfairness cannot give Betty a claim to Alice's resources if Alice is under no duty to give those resources away rather than use them to promote her own ends. Assuming, then, that Alice in *Hired Help* satisfies her general duties to aid and has some remaining resources, she does not violate any claim of Betty's if she uses some of her resources to hire Bodyguard and Bodyguard subsequently saves Alice from a lesser harm.

To be clear: the argument here is not that Alice may fail to save Betty because Alice generally satisfies her duties to aid by paying her taxes. Rather, we've been assuming throughout that Alice is under no duty to rescue Betty in either *Crocodiles* or *Hired Help* because that particular rescue is so costly for Alice, given the extra weight that she may attribute to her own interests. Alice need not sacrifice her one leg to save Betty's two

36. Those who think that duties to aid take priority over securing one's basic interests can also apply this analysis to the earlier case, in which Alice is significantly imperiled but has satisfied her duties to aid.

legs. The significance of Alice's satisfying her general duties to aid is rather that this grounds her permission to use her surplus resources to pursue her own ends, including by hiring Bodyguard, even if she and Betty did not begin with fair shares of resources.

Nor does this argument rely on Alice's outsourcing an agent-relative prerogative to Bodyguard. Rather, the argument holds that Betty lacks a claim that Bodyguard rescue her when he has agreed to preferentially save Alice at a time when he was under no duty to rescue Betty. Since Betty lacks a claim against Bodyguard that he save her in these circumstances, Bodyguard does not need some special prerogative to justify failing to save her. Even if Bodyguard would have contracted with Betty had he not contracted with Alice, his contracting with Alice doesn't conflict with any of Betty's claims. Betty has no claim that anyone serve as her bodyguard. The only role that Alice's agent-relative prerogative plays is to ensure that, in saving Alice, Bodyguard is not helping Alice to fail in a duty that she has to save Betty.

However, as indicated above, our account of when agents are entitled to the benefits of their agreements will settle the boundaries of contracts to preferentially save. Consider *Beach* again. Imagine that, through no fault of her own, Betty cannot afford to contract with Sunbather. Since Sunbather's being in a position to save Betty is not an upshot of his agreement with Alice, even in the explanatory sense considered in Section VI.B, we might think that Alice is not entitled to deprive Betty of an important benefit that Betty would have otherwise enjoyed and cannot secure for herself because of an unfair lack of resources. This would be more than mere unfairness in the initial distribution of resources, which cannot itself defeat Alice's claim to use her surplus resources for her own sake. Hiring Sunbather also makes Betty worse off (although Betty, of course, has no claim that Sunbather stay on the beach to rescue her).

Again, I cannot settle the extent of Alice's entitlements here, but I suspect that it is in this sort of case that concerns about initially unfair distributions of resources will have the most bite. One possibility (which I will not explore further) is that an initially unfair distribution of resources affects proportionality—that is, how much more harm a rescuer may allow Betty to suffer in order to save Alice.³⁷ However, concerns about unfair distribution might also be tempered by whether Alice endangers herself only on the condition that Sunbather contracts to save her. We might think it an important part of allowing agents to pursue their own ends that Alice may secure this sort of benefit for herself, and that Sun-

37. We can imagine that something similar is true of inflicting collateral harms: if someone is already badly off, this plausibly affects how much further harm may be inflicted on her.

bather may provide it, provided that Sunbather is under no duty to provide the benefit to Betty at the time he forms the contract with Alice.

In light of the transferring of benefits and resources, contracts to preferentially save might often be more morally significant than mere agreements. The point of surplus resources is that others have no claim that they be used for their sake even if they would thereby do more good. When, for example, the security guard in *Alarm* is in a position to rescue because he has accepted some of Alice's resources in exchange for agreeing to save her, saving Betty not only transfers a benefit to Betty but also effectively transfers Alice's resources to Betty, when Betty had no claim that those resources be used for her sake. This gives us further reason to object to a failure to save Alice.

The significance of the transferring of benefits and resources in justifying preferential saving illuminates one of the most widespread examples of contracting to save that we identified at the outset—namely, the agreement between citizens and the members of their state's armed forces. Through their government, and using their taxes, citizens hire combatants to protect them. These combatants are typically under no duty to rescue when this agreement is formed (enlisting is usually supererogatory). That citizens have formed agreements and used their resources in order to secure their own rescue explains the plausible view that combatants may preferentially save their own citizens even if this does not minimize harm overall. And, attractively, this account does not rely on dubious notions of national partiality, according to which I may care more about my co-citizens merely in virtue of our shared nationality. Rather, it relies on the importance of being able to use our resources and form plans to further our own ends, which is hardly controversial among nonconsequentialists.³⁸ A further advantage of this account is that it can accommodate preferential saving by both armed forces that consist of compatriots and those that consist of foreigners.³⁹

VIII. CONCLUSION

This article has explored the role of agreements in determining the permissibility of preferential saving. I rejected the view that one can, by mere agreement, make it possible for third parties to act on one's agent-relative prerogatives. The notion of outsourcing agent-relative reasons is incompatible with a plausibly robust duty to minimize harm when rescuing. It would permit a rescuer to preferentially save whenever (she reasonably

38. For defenses of the role of national partiality in war, see Hurka, "Proportionality in the Morality of War"; and Rodin, "Myth of National Self-Defence." Note that, on my account, these considerations permit only saving: they do not permit combatants to inflict otherwise disproportionate harms.

39. Thanks to Cécile Fabre for this point.

believes that) the person facing a lesser harm wants to be rescued. It would similarly defeat any duty to save the greatest number of people whenever a smaller number of people want to be saved instead. I further argued that mere agreements to save cannot defeat other people's claims to be saved: Passer-By may not save Alice if he is already under a duty to save Betty.

The failures of the Outsourcing View reveal an important constraint on the kinds of agreements to save that we can form. I have argued that agreements to preferentially save are valid only if the rescuer is not under a duty to save at the time the agreement is formed. This rules out the "on demand" preferential saving in cases such as *Help* that gave us reason to reject the Outsourcing View. While nonconsequentialists can and ought to make space for the duty to minimize harm when rescuing, they must also leave agents adequate scope for pursuing their own ends. This includes allowing agents to reliably secure benefits through forming agreements, which can include the benefit of rescue even when this does not minimize harm.