

The Rules of Rescue: Cost, Distance, and Effective Altruism, by Theron Pummer. New York, Oxford University Press, 2023. Pp. x + 253.

Must we devote our lives and incomes to helping distant strangers? Is it wrong to be an ‘ineffective altruist’, supporting causes that help relatively few people? Is it worthy of blame?

You might expect *The Rules of Rescue*—Theron Pummer’s excellent new book—to opt for one of two views: either we are burdened with pervasive and demanding requirements to help, or else we enjoy extensive freedoms to do as we please. Instead, Pummer develops both views at once.

Every year, millions of people die for want of cheap medicine and basic resources. Many of us have the power to help, either through private donations or political action. If we do have this power, we confront a ‘ubiquity’ of reasons to rescue strangers (p. 181). The more we can rescue, the stronger our reasons (pp. 39–58). These reasons do not diminish when the strangers are far away (105–09), when their plight is not salient to us (pp. 109–13), or when we have already rescued many others (p. 135). It makes no great difference whether those in need are victims of injustice or of misfortune (pp. 115–17), whether we are their unique potential saviours (pp. 113–15), whether there is some tiny chance of failure (pp. 117–18), and whether we are unable to know in advance whom we will end up rescuing (pp. 118–121). As long as we know that our efforts will probably save lives (unlike in ‘Risky Diffusion’, pp. 120–23)—and that we are unlikely to cause significant harm (p. 118)—we have powerful reasons to do what we can.

On the other hand, there is more to being human than making oneself useful. We are autonomous beings, free to determine the course of our own lives (pp. 27–30), especially when altruism costs us ‘limbs, comfort, safety, time, and money’ (p. 27). Even the inconvenience of being ‘switched on’ as a vigilant altruist counts as a cost (p. 141). Our permissions are further ‘amplified’ as we expect costs to pile up over the course of a lifetime (pp. 138–44). When we ‘responsibly acquire’ special connections—such as friendships and voluntary associations—this may ‘increase the lifetime cost’ we must pay to help others (pp. 174–75). But connections that are ‘fully non-responsibly acquired’—such as relations to family—do not have this effect (p. 175). Commitments to the pickleball club cannot justify giving less to charity, but a parent’s medical bills can. On top of all this, we may often choose to help a particular individual as such, even if others need our help somewhat more (p. 46). The ‘ubiquity’ of reasons is thus met with a ubiquity of freedoms. The results are sensible and subtle—effective altruism with a human face.

In light of all this, I recommend *The Rules of Rescue*. Each chapter has something fresh to say. As a whole, the book stands out for its clarity, breadth, and elegant structure. The first half develops the ‘rules’ through a series of twists on a thought experiment; the second deals with real-world complexities. Both parts of the project are easy to praise and enjoyable to read.

Now then: objections. In my view, some of Pummer’s arguments do not go into enough depth, and some of the views should have been worked out more systematically. This was not such a problem with Chapters 3 (The All or Nothing Problem), 4 (Praiseworthiness), 6 (Frequent Rescues), and 7 (Special Connections). But it is a drawback of the treatment of permissions (§1), aggregation (§2), and especially risk (§3).

The gloves are off for these topics. Then I will end with more praise (§4).

1. Reasons and requirements

Must you rescue people from lethal harm? It depends. Contrast two cases:

Costless Rescue

'A stranger faces a deadly threat. You can either do nothing, allowing them to die, or you can, at no cost to yourself, save their life.' (p. 19)

Costly Rescue

Same as before, but you can save the life only 'at great cost to yourself' —not anything so bad as death, but something still 'very painful and disruptive', like losing your legs. (p. 20)

Only monsters forgo costless rescues, but you don't have to save a life at the cost of your legs.

Whence the difference? In a word: reasons. In *Costless Rescue*, your reason to save the stranger is unopposed, so it grounds a requirement. In *Costly Rescue*, your reason to save is opposed by a reason not to give up your legs, with the result that giving is optional.

But this is not just any clash of reasons. In *Costly Rescue*, there remains a sense in which you have *most* reason to save the life, even though you are not *obligated* to do so. Saving the life is the better thing to do, the favoured 'supererogatory' option. Contrast this with a humdrum clash:

Harmful Rescue

Same as before, but you can save the life only at cost to an unwilling third party—say, by crushing their legs or taking their spare kidney.

Here, as before, you are not required to save the life. But nor is saving the life 'favoured'. The reason against harming *disfavours* saving the life—indeed, it makes saving the life *wrong*. This is quite unlike a heroic *Costly Rescue*.

What makes *Costly Rescue* special, for Pummer, is that it involves a special kind of reason. Life-saving and harm-avoidance are favoured by good old 'requiring reasons', which tend to make actions morally required (and therefore also permissible). Self-interest is favoured by 'permitting reasons', which merely tend to make acts permissible (pp. 24–5). That is why the reason to keep your legs does not forbid a *Costly Rescue*.

Pummer's distinction is familiar. Requiring reasons are 'prima facie duties', and permitting reasons moonlight as 'permissions', 'options', or (as I prefer) 'prerogatives'. But Pummer has at least two new things to say about permitting reasons.

First, he gives a novel analysis of what permitting reasons *are*. In a sense, they are not really reasons: they do not (as such) count in favour of actions. Instead, they are a kind of defeater (p. 25). They prevent requiring reasons from requiring. For example, in *Costly Rescue*, the cost of losing your legs prevents your reason to save the life from grounding a requirement. This idea strikes me as essentially right; it is surprising that no one has said it so clearly before. That said, it raises some questions that Pummer does not get into. Why should harm to oneself be a defeater rather than a normal reason? How could it undercut the requiring force of opposing reasons but not their ability to 'favour'? (Defeaters *never* selectively undercut favouring—why not?) Why think that permitting reasons subtract from would-be requirements, rather than adding a presumption of permission?

These questions might leave us wanting a deeper theory of where permitting reasons come from. But this is no objection to Pummer's analysis, which is useful in part because it is

theory-neutral. Whether ‘permitting reasons’ are fundamental or derived, we can understand them as defeaters.

Second, Pummer has a novel view of *which* permitting reasons we have, going beyond the standard ‘agent-favouring prerogative’ to avoid paying costs. Consider what we might call:

Transformative Rescue

You are deciding whether to become a ‘philanthropist banker’ or a ‘humanitarian doctor’. As a banker, you would save more lives, and you would be better off. But your life would unfold in an utterly different way. (p. 28)

Even though there is no ‘cost’ to becoming a banker, you still have a permitting reason to become a doctor. Pummer calls this an ‘autonomy-based permitting reason’, as it lets you ‘determine your own life’ regardless of self-interest (p. 27). The case isn’t as ‘clean’ as Pummer’s others. But the idea seems promising.

The trouble starts when we ask what ‘autonomy-based reasons’ are supposed to permit. Officially, Pummer has one answer:

Life-Affecting Choices

Autonomy-based reasons let you choose ‘how your life unfolds’, giving you permission when it comes to choices that affect the course of your life. (p. 28)

He also suggests:

Rights-Protected Choices

Autonomy-based reasons concern what is ‘rightfully yours’, giving you permission to use your things ‘as you choose’. (p. 29)

But these are not the same, and Pummer’s official answer is hard to defend.

Consider a spin on Harmful Rescue:

Harmful Self-Rescue

You can save *your own* life only at cost to an unwilling third party—say, by taking their spare kidney.

Few deontologists would abide by organ theft. But here it would transform ‘how your life unfolds’. Wouldn’t ‘autonomy’ permit you to take the kidney?

Pummer’s reply is that we lack ‘autonomy-based permitting reasons to act in ways that contravene the rights of others’ (28). By contrast, ‘it seems plausible that there can be cost-based permitting reasons’ to contravene rights; he illustrates with the example of rolling a ‘boulder onto a stranger’s hand’ as a ‘necessary side-effect’ of saving one’s own life (28).

Three problems.

First, even if you may roll the boulder, you have *no* special permitting reason to steal the kidney. Why not?

Second, Pummer’s view has controversial implications when it comes to unowned resources. Suppose you and I are being chased by a shark, and we come across an unowned one-person canoe; whoever gets in will be safe, and whoever is left out will lose a leg. In Helen

Frowe's (2021, pp. 74–5) view, even if I can outswim you, I may not simply take the canoe: I do not own it, and I am not entitled to acquire it on the basis of a speedy freestyle. Ideally, we should flip a coin. Frowe's view strikes me as plausible (though itself controversial). But Pummer's view entails that I may use the canoe as I wish—even if you need it more. Autonomy and self-interest both permit me to arrogate unowned resources. This entailment needs defence.

Third and finally, although it is plausible that autonomy stops where others' rights begin, Pummer cannot tell us *why* this should be so. Infringing rights affects 'how my life unfolds'. Why shouldn't this be a matter of 'autonomy'?

Pummer could have had an answer. If he had linked autonomy to rights, he could have said that your autonomy extends only to what you own. You would then have a permission to keep your kidney, but not to steal one. This view also provides a principled take on unowned resources: you have no rights over them, and therefore no special permissions to take them.

In light of the above problems, I think Pummer should have focused less on costs and more on rights. Autonomy-based permissions could be grounded exclusively in rights, and cost-based permissions could be scrapped altogether. (See Muñoz 2021 for a view like this.) After all, cost-based permissions are not needed to explain Costly Rescue, in which you give up your legs, and cost-based permissions seem inapt in Harmful Self-Rescue and Frowe's canoe. Why not stick with rights?

One last thought. I wonder if we can derive Pummer's 'life-affecting' permissions from more familiar cost-based permissions. Isn't there some cost to changing lifestyles? It might not be costly *all things considered*, but it will be costly *along some dimensions of value*, which might be enough to ground a permission. This would make Pummer's view more parsimonious. But given the problems with cost-based permissions, we might prefer to get our parsimony in another way: by grounding all permissions in rights.

2. Saving and summing

On to Chapter 2: 'Numbers Count'. Doesn't everyone agree that we should save more people? It depends. Contrast two cases:

Costless No-Conflict

'Two strangers face a deadly threat. You can do nothing, save one stranger's life at no cost to yourself, or save both their lives at no cost to yourself.' (p. 40)

Costless Conflict

Three strangers face a deadly threat. You can do nothing, save one stranger's life at no cost to yourself, or save the other two at no cost to yourself. Tragically, you cannot save all three.

Again, only monsters forgo costless rescues—so, in Costless No-Conflict, you have to rescue everybody—but there is *some* reason not to save the many in Costless Conflict: you would be letting someone die. Some views take this reason to justify saving the few outright. Others just say you may flip a coin (or roll a die).

Pummer, like most ethicists, believes that you *must* save the many: 'the numbers count'. This idea of 'interpersonal aggregation' greatly helps the case for effective altruism. Why believe it? Some arguments appeal to sums of suffering, maximizing rationality, transitivity

principles, and fairness. Pummer (probably wisely) does not attempt a full survey. He focuses on three other arguments in favour of aggregation, along with three objections. Let us start with the arguments in favour.

First, if we do not count the numbers when it comes to rescues, we cannot count the numbers when it comes to *harms*. But even if you may save your life in a way that harms one person, few would say you may save your life in a way that harms many (p. 56). This point is interesting but, as Pummer acknowledges, familiar (Kamm 2005).

Second, an argument from ‘moral equivalence’. Saving both of the two is clearly better than saving only one of them—it is better for someone, equally good for everybody else. But saving only one of the two is ‘morally equivalent’ to saving the lone other. Saving the two, therefore, must also be better than saving the lone other.

A hopeless argument, as Pummer explains:

It is true that, whether you save *A* [the lone one] or save *B* [one of the two], you save exactly one person. But reasons to rescue numbers of people are built of reasons to rescue particular individuals. Because *A* and *B* are different particular individuals, there is a moral difference between saving *A* rather than *A* and *C* in *Costless No-Conflict* and saving *A* rather than *B* and *C* in *Costless Conflict*. (p. 44)

To put it another way, no self-respecting deontologist would say that human lives have fungible value. The whole *point* of deontology is to respect the ‘separateness of persons’, as opposed to combining their pleasures and pains into a homogenous blob. (I independently came up with a version of this in Muñoz 2022.)

The third argument, which Pummer endorses, is that reasons combine:

What makes it wrong to save the lesser number in cases like *Costless Conflict* is simply that there is more requiring reason overall to save the greater number (*B* and *C*) and no sufficiently strong permitting reason to save the lesser number (*A*). (p. 44)

...the requiring reasons to save *B*’s life and *C*’s life together outweigh the requiring reason to save *A*’s life.... (p. 44)

This is Pummer’s main argument for aggregation.

Now on to the objections. The first attacks this idea about combining reasons:

If the requiring reasons to save *B*’s life and *C*’s life together outweigh the requiring reason to save *A*’s life, then presumably the requiring reasons to save *B*’s life and *C*’s *finger* together outweigh the requiring reason to save *A*’s life. But it’s not wrong to save *A*’s life instead of saving *B*’s life and *C*’s finger. (p. 44; see also pp. 52–3)

In other words, if reasons combine, then ties should be sensitive even to mild sweetening—like the prospect of saving an additional finger. But Pummer does not think that moral ties are so easily broken. So, he has to find some way to make sense of the permissibility of saving *A*’s life rather than *B*’s life plus *C*’s finger.

Here, I was expecting Pummer to say that insensitivity to sweetening is the mark of ‘parity’ (Chang 2002), the relation that holds between ‘roughly balanced reasons’ (as he puts

it in Gordon-Solmon & Pummer 2022, p. 641). Surprisingly, *The Rules of Rescue* mentions this relation only once (at p. 209n9, suggesting that top charities might be on par).

Instead of parity, Pummer appeals to a new concept: ‘individualist permitting reasons’. For each person, you have a requiring reason to save them *and* a further permitting reason to save them. You therefore have *more* permitting reason than requiring reason to save a person’s life. The result is moral wiggle room. Although you have more requiring reason to save *B*’s life and *C*’s finger, your ‘individualist permitting reason’ to save *A* can compensate for the difference (p. 46).

This idea is interesting and original—an agent-neutral prerogative.

Some will object that, on Pummer’s view, saving the life plus the finger comes out as the better thing to do. That does seem odd (to a parity fan like me). But here I want to focus on a different problem: the idea needs to be more systematically worked out. I will limit myself to three questions.

First, why just individual people? Why don’t we have permitting reasons whenever we choose between things with nonfungible value?

Second, if we have individualist permitting reasons, presumably we also have them as regards ourselves. But then why do we need further ‘cost-based permitting reasons’? Perhaps Pummer wants these to account for a ‘self-other asymmetry’, where *self*-interest grants uniquely broad permissions. But then I would have liked to see an argument for the asymmetry, since the opposition rejects it (Taurek 1977, pp. 296–302).

Third, why do ‘individualist permitting reasons’ cover interests but not autonomy? While I am at it, why assume that autonomy-based permissions are self-other asymmetric—why not a permission to refrain from disturbing the lives of others?

Enough about permissions. On to the second objection to aggregation: that it will ‘run amok’, forcing us to sum up small harms across many people (p. 50). Some will see this as a disaster. No number of headaches can add up to something worse than death. If that is your view, Pummer says you should not give up on aggregation; you should ‘limit’ it to cases where harms are ‘sufficiently similar’ (p. 51). But limited aggregation has problems—diachronic inconsistency, intransitivity, arbitrariness. Of these, only intransitivity gets a mention (p. 58n7).

The final objection, due to Anscombe (1967), links wrongness to owing: we do not owe it to any individual among the many to save them, and so we cannot do wrong by saving the few instead. Pummer convincingly shows that this argument is not persuasive (pp. 48–9).

The goal of Chapter 2 is to ‘defend the claim that it is wrong to save the lesser number’ (p. 42). I think Pummer has laid some bricks in the fort. But he has not gone on the *attack*. There is no new argument for aggregation, which is a problem, if the goal is to establish a ‘ubiquity’ of reasons for effective altruism. Indeed, Pummer later glosses ‘effectiveness’ as ‘helping the most’ (p. 197), understood as ‘*helping a greater number of people*’ (p. 196). Without aggregation, the case for ‘effectiveness’ will not be so persuasive.

In my view, Pummer’s argument needs a supplement: either a more aggressive case for aggregation, or an argument that even ‘numbers skeptics’ ought to give effectively. That said, I am more skeptical of aggregation than many other deontologists. They might not need to be persuaded to count the numbers.

3. Expected value and ex ante views

My main objection has to do with risk.

In Chapter 5 (Distant Rescues), Pummer makes a powerful case that salience and physical proximity do not affect our duties to help. But the sections on risk mostly consist of raw intuitions, some of which his opponents will reject for principled reasons.

Consider a choice between two lotteries: L_1 will give Alan a 99% chance of being saved ('Concentration'), and L_2 will randomly select two people from a group of 10 billion and give *those two* a 99% chance of both being saved ('Closed Diffusion'). To Pummer, there seems to be 'more requiring reason' to choose lotteries like L_2 , which are almost certain to save more lives (p. 121). But some would disagree. Some think we should focus on the expected benefits to each individual considered 'ex ante' (that is, at the time of choice). Clearly, ex ante, Alan has *much* more at stake than does anybody among the ten billion—his chances go up by 99% if you pick L_1 , whereas the others gain only a sliver of a percent if you pick L_2 . So, we must weigh Alan's big claim against the tiny claims of the billions. The 'limited' aggregationist will say that Alan's claim wins: no number of tiny claims can outweigh a very big one. Since Pummer is open to limited aggregation, as we saw in §2, it seems Pummer must either reject ex ante views, or else give up his view that we should pick lotteries like L_2 over those like L_1 .

The stakes are high. If Pummer backs down here, he will undercut the case for giving to effective charities. Even if a large donation is near certain to save lives, it might not give anyone in particular a big chance of being saved. In such cases, the duty to donate ends up favouring charities that concentrate their efforts on some identifiable individuals, as opposed to those that save the most 'statistical lives'.

Again, I think the argument needs a supplement: either an objection to ex ante views (when combined with limited aggregation), or an argument that even those who subscribe to such views ought to give effectively. (A third option: Pummer could qualify his claims about the 'ubiquity of requiring reasons' in favour of altruism; see, for example, the list of causes on p. 187, many of which do not give any identifiable person a hefty expected benefit.)

Pummer's sympathies, I suspect, lie with 'ex post' views, which tell us to focus on the claims that people will predictably have once the dust has settled. (For example, if L_3 randomly selects and kills someone from the 10 billion, there is a single life-sized claim against L_3 ex post—as opposed to 10 billion tiny claims ex ante.) But *The Rules of Rescue* does not make the case for ex post views, grapple with the case for (thoroughly) ex ante views, or look for a hybrid. The terms 'ex post' and 'ex ante' do not come up.

My guess is that readers' main concern about *The Rules of Rescue* will be its breezy treatment of risk. Perhaps the discussion on pp. 117–23 could have been expanded into its own chapter, covering such issues as 'cluelessness', collective action, and the case against (thoroughly) ex ante views. That said, one virtue of *Rules* is its brevity. I can see why Pummer would not want to get bogged down in internecine debates.

4. Conclusion

The Rules of Rescue is a superb source on a variety of topics such as salience, distance, imperfect duties, and praise. Unsurprisingly, I also agreed with Pummer's take on conditional obligations, a topic on which we have collaborated (Muñoz & Pummer 2022).

I do not think *The Rules of Rescue* shows that we have ubiquitous reasons to become effective altruists, because it leaves open the possibility of ex ante views with limited aggregation (and it does not convert the numbers skeptic). But perhaps this is the wrong way to see the project. Maybe the point is not to get *everyone* to go along, but instead to show that *many* deontologists should do so. So understood, *The Rules of Rescue* is a success. It also

succeeds in developing its pluralist theory of permissions, which I found impressive, despite thinking that it should be more parsimonious and less centred around costs (§1).

At last, as promised: praise. Consider:

Costly No-Conflict

‘Three strangers face a deadly threat. You can do nothing, save one stranger’s life at great cost to yourself, or save the other two at the same great cost to yourself.’ (p. 62)

In Chapter 4 (Praiseworthiness), Pummer argues that it would be wrong to save only one of the strangers (since you easily save more); permissible to save neither (given the costs of heroism); and yet it is not overall *blameworthy* to save only the one—it might even be *praiseworthy*. This seems outrageous. How could a wrong act merit praise? But Pummer has an argument. You are somehow blameworthy for saving one (rather than two); somehow praiseworthy for saving one (rather than zero); and, in this case, the praiseworthiness seems weightier. Such comparisons of weight seem utterly commonplace (the numbers skeptic will say similar things in other cases). Either Pummer has discovered a bizarre moral possibility—one with real-world applications for how to think of suboptimal do-gooders—or he has hit upon a paradox in our commonsense concept of blame.

Whatever we make of Pummer’s discovery, it is a real discovery, and like several other ideas in *The Rules of Rescue*, I expect that it will have an enduring impact on normative ethics.*

References

- Anscombe, G.E.M. (1967). ‘Who is wronged? Philippa Foot on double effect: One point’. *The Oxford Review* 5: 16–7.
- Chang, Ruth (2002). ‘The possibility of parity’. *Ethics* 112 (4): 659–88.
- Frowe, Helen (2021). ‘The limited use view of the duty to save’, in David Sobel, Peter Vallentyne, and Steven Wall (eds.), *Oxford Studies in Political Philosophy, Volume 7*. Oxford: Oxford University Press, pp. 66–99.
- Gordon-Solmon, Kerah and Pummer, Theron (2022). ‘Lesser-evil justifications: A reply to Frowe’. *Law and Philosophy* 41: 639–46.
- Kamm, Frances (2005). ‘Aggregation and two moral methods’. *Utilitas* 17 (1): 1–23.
- Muñoz, Daniel (2021). ‘From rights to prerogatives’. *Philosophy and Phenomenological Research* 102 (3): 608–23.
- Muñoz, Daniel (2022). ‘The many, the few, and the nature of value’. *Ergo* 9 (4): 70–87.
- Muñoz, Daniel and Pummer, Theron (2022). ‘Supererogation and conditional obligation’. *Philosophical Studies* 179: 1429–443.
- Taurek, John M. (1977). ‘Should the numbers count?’ *Philosophy and Public Affairs* 6 (4): 293–316.

DANIEL MUÑOZ

University of North Carolina at Chapel Hill, USA

daniel.munoz@unc.edu

* My thanks to Abbey Burke and Joe Bowen for helpful comments. Special thanks to Theron Pummer, who gave multiple rounds of comments leading to many improvements, especially to §3.