Count the numbers

Clayton Littlejohn

1. Easy questions about the morality of rescue

We often skip past the easy questions about the morality of rescue. Let's stop to consider a few. Please fill in the blanks. Use a pen. These are easy questions.

> **Simple island rescue (boat types)**: There's one person trapped on Alpha Island and two on Beta. (Pop quiz: how many would be on Gamma, Delta, Epsilon, …, Delta to the Delta, etc. …?) You can't reach both islands in time. Which boat should you take, the sailboat (.6 chance of getting to your chosen island in time) or the motorboat (.9 chance)?
> Easy, we say: ___________________________.
>
> **Simple island rescue (boat size)**: There are two trapped on Beta. You can send the boat that's big enough to bring one back or the boat large enough to bring back two. You're certain that neither boat is more likely to reach the island than the other. Which boat do you send?
> Easy, we say: ___________________________.
>
> **Simple island rescue (routes)**: Heading towards your island, you could take the route that creates a wake that might (e.g., one in ten chance) push needed medical supplies to the other island or take the route that won't create any such wake. Which way do you go?
> Easy, we say: ___________________________.
>
> **Mission improbable**: You can take the boat of your choice to Alpha or Alpha'. There is almost no chance of reaching Alpha' in time and almost no chance that you wouldn't reach Alpha in time. Which way do you go?
> Easy, we say: ___________________________.

I hope we agree that these were easy questions. I also hope our answers look alike.

With some luck, we might also draw the same lessons from this. We might agree, for example, that if we're going to try to rescue some person(s), we shouldn't choose the less reliable means (**No Daredevils**).[1] If we're sending the surgeon in for surgery and they're known to drink over lunch, it's better to schedule those surgeries before lunch. We also agree that when we're choosing between two missions where we don't think it's more desirable to succeed in one as opposed to the other, we should choose one mission over another if we're nearly certain to fail in one case and not the other

[1] Less informally, we can describe lotteries with two outcomes (where the most desirable outcome(s) will be described as prizes) like so. Understand 'L(p, X, Y)' as a lottery in which the probability of getting X (but not Y) is p, the probability of getting Y (but not X) is 1-p. If, say, there's some operation that is certain to save either Brenda or Bernhard, certain not to save both, and equally likely to save either, we could represent this as 'L(.5, Brenda, Bernhard)'. If they are both on Beta island and we're looking at a lottery in which there's equal chance of saving both and of saving neither, we could represent this as, 'L(.5, β2, β0). I'll speak as if Brenda's successful rescue and Bernhard's successful rescue are prizes (because equally desirable) and won't describe β0 (i.e., the outcome in which neither is saved) as a prize. We can say that **No Daredevils** says that if $p > p'$ and X and Y are possible outcomes where X is more desirable than Y, L(p, X, Y) > L(p', X, Y). This follows from Resnik's (1987) better chances condition.

(**No Heroes**).[2] Even if we think petitionary prayer might raise the probability of Alice's recovery a smidge, if we can use that time to pack and send medicine that's nearly certain to save Adella, that's what we should do. We also agree that we should save everyone in a group and not just some if it's just as easy to save everyone (**No Monsters**). If we have pills enough to save Adella and Alfonse, we should send both, not just send enough for one. Finally, if we think we should create the wake and use the motorboat, we seem to think that options can be sweetened in two ways—by increasing the probability of a successful mission or by adding causal pathways to increase the number of lives saved if an option is chosen (**Double Sweetening**).

My plan isn't to collect a few handfuls of obvious answers to easy questions and then down my tools. My plan is to use put these answers to good use. I want to use them to try to settle a difficult question about the morality of rescue. The questions above, in my experience, do not divide opinion. The lessons taken from them don't strike us as particularly controversial. We should be able to use them to test theories about the morality of rescue. Some theories don't struggle at all to explain why our answers are correct. Some will be unable to explain why they're correct. Other things equal, we should prefer the former theories to the latter on the grounds that they explain obvious points about what we should do to help others that rival theories do not.

2. The difficult question

Readers must have known that this paper wasn't about whether it's okay to abandon Bernhard when rescuing Bernadette (thereby violating **No Monsters**) or whether it's okay for surgeons dim the lights before attempting a difficult operation (thereby violating **No Daredevils**). This paper is really about numbers. Specifically, it's about the moral relevance of the numbers of strangers who stand to gain or lose when we choose between our options. It's about the way we answer questions like this one:

> **Simple Island Rescue**: One person is trapped on Alpha island. Two are trapped on Beta. While there's time to reach one island, there's not time to reach both. A choice must be made. Assuming that we should head to one island or the other, which island should we head to?
> You say: __________________________.

I can't assume we answer this in the same way. My view is we should head to Beta to save two (β2) rather than head to Alpha to save one (α1).[3] My answer is controversial, so this question is difficult. Some will say that it's not wrong to choose α1 straight off or use a coin flip to determine who will be saved. Taurek (1977) seems inclined to flip a coin, but I get the sense that he might be fine with deciding 'straight off' to head to Alpha or to Beta. It's this disagreement between people who think we should save the greater number and the number sceptics ('sceptics' henceforth) that's the focus of this paper.[4]

---

[2] For some n, if neither X nor X' is strictly preferred to the other, if p-p' > n, L(p, X, Y) > L(p', X', Y). If, for example, we have to choose between pressing a button that's nearly certain to save Adella and using thoughts and prayers when that is nearly certain not to help Alfonse, assuming that we care about both to some similar degree, we should press the button.

[3] The numbers don't always settle the issue for me. If someone is on the smaller island because of some sort of injustice or the person on that island stands in some sort of significant relation to me (e.g., it's a parent, a friend, a friendly barista, etc.), that changes things, but then I don't see this as a pure rescue case. Even in the pure cases where the parties are perfect strangers, historical considerations don't seem relevant, the people don't differ in terms of their future potential, etc., some will say that it's not true that we have to choose β2 over α1.

[4] For discussions of the sceptical view that are sympathetic (to some degree), see Anscombe (1967), Doggett (2013), Fischer (2021), Munoz-Darde (2005), Taurek (1977, 2021), Tenenbaum (forthcoming), and Timmermann (2004). For some (non-consequentialist) arguments that we should save the greater number, see Kamm (2005) and Scanlon (1998).

Part of my interest in this paper derives from the fact that the answers to easy questions don't, in my experience, give us much information about how they'll answer this difficult question about numbers. I think it's interesting, then, to think about whether the attitudes that we all have towards the prospect of failure and success and the subsequent attitudes we take towards uncertainty and risk might be difficult to reconcile with the sceptic's attitude towards the desirability of different outcomes, options, and choices. The sceptics don't see themselves as ineffective altruists (i.e., people who profess to care about others but concede that they don't really prefer to take the more effective means to their ends). They might think, for example, that it's fine to flip a coin to decide whether to choose β2 or α1, but they don't think that it's fine to flip a coin to decide whether to take a boat that can hold at least three passengers on missions to Beta, whether to roll a die deciding whether to send surgeons to drinks before or after their surgeries, etc. Their attitude seems to be that some kinds of numbers matter. The numbers that describe the probabilities matter. Other numbers, they insist, don't matter. It doesn't matter, they say, whether we rescue two on Beta or one on Alpha. Coins here are fair to flip.

When I first started thinking about these issues, I started to wonder whether we could make sense of this view. Here's a naïve thought. Someone who has decided to head to Beta and knows that they should take a big enough boat sees full success (i.e., β2) as more desirable than the partial success that comes from saving just one (i.e., β1), and both are more desirable, in turn, than complete failure (i.e., β0). This explains why they care about the means (e.g., why it's better to bring drugs or use boats that might help both). This assumes they care about the ends in a way that suggested, to me, that the numbers of lives saved matter. At the very least, such numbers mattered sometimes. Ultimately, I think there is a tension between our intuitions, our tools for thinking about uncertainty, and the sceptic's commitments. I'll need to convince you that there is by convincing you that there's no combination of decision-theoretic tools and sceptical views that remains true to the sceptic's core commitments and justifies our answers without generating incoherent or contradictory results.

In §3, we'll begin by seeing whether the sceptic can use the most familiar decision-theoretic tools to justify our answers. This approach might initially seem promising because claims like **No Daredevils** seem to be closely related to axioms some might use to justify expected value theory (EVT).[5] Upon closer examination, we'll see that the sceptic cannot use EVT to justify any of our answers to the easy questions. Once we see why, we'll see that the sceptic might want to relax some of the assumptions about desirability and rational preference operative in §3 and find some alternative decision-theory that might let them relax these assumptions. In §4, we'll see whether the sceptic might be able to use some propsectist decision theory to try to justify our answers to the easy questions. We'll look at two prospectist theories and see that neither theory gives the sceptic what they need. One prospectist theory gives the sceptic too much leeway. The other doesn't give them enough. At this point, it should be clear that we've done due diligence and we don't have to continue looking for new decision-theoretic machinery that the sceptic might use to justify our answers on the questionnaire above. In §5, I shall sketch a positive proposal about how we should approach these cases.

3. The simple view

The simple view (as I'm calling it) combines the sceptic's moral view with some familiar decision-theoretic machinery. We can imagine (or try to) that some sceptical agent's preferences can

[5] When we look at views that assume EVT, we'll assume that the possible outcomes of actions have a kind of value or desirability that meets ordering conditions (e.g., every pair of outcomes is either equally desirable or one is superior), that some sceptical agent might have preferences that mirror or encode these value relations, and, if this is possible, that this agent's preferences meet the conditions necessary for that agent to be representable as an expected utility maximiser. We'll assume that this agent is permitted to choose an option if there's no alternative with greater expected value and required to choose an option if it alone maximises expected value. Readers who think that it would be unwise for the sceptic to make use of EVT might be right. I ask for their patience. We'll look at alternatives later.

be represented by some utility function. This agent will prefer, for example, saving one person on Alpha ($\alpha1$) to saving nobody from Alpha ($\alpha0$), saving each person on Beta (β2) to saving just one (β1) or neither (β0). Of course, this agent will not prefer β2 to $\alpha1$ or vice-versa. Since the agent's preferences are complete (i.e., for every pair of options, X and Y, the agent weakly prefers X to Y or vice-versa), this agent equi-prefers β2 and $\alpha1$ (i.e., $\alpha1$~β2).

Since we'll be looking at cases involving uncertainty, we also need to know how the sceptic thinks we should choose between options where it's not certain that things will turn out for the best. It's here that they might appeal to EVT. I'll be using 'value' in some very extended sense to capture what's morally desirable according to the sceptical view. I'll assume that the sceptic thinks that we should only be concerned with the interests of potential beneficiaries. I'll also assume that their indifference between $\alpha1$ and β2 is compatible with a broadly anti-inegalitarian outlook.[6] I'll assume that the desirability of an option derives from the desirability of its possible outcomes and the probability that they'll eventuate.

If we use EVT, it seems we can make quick work of claims like **No Heroes**, **No Monsters**, and **No Daredevils**. For example, if the sceptic sees something desirable about rescuing each person, it's not surprising that they'd have these preferences: β2 > β1 > β0. If they deny that it's better simpliciter or period to head to Beta just because we could save Bruna and save Bernhard rather than simply save Adella on Alpha, it's not surprising that their agent has this equi-preference: $\alpha1$~β2. If EVT is both powerful and sufficiently neutral on substantive questions to be widely applicable, shouldn't things work out well for the sceptical view?

No. Remember the game isn't to see whether we can draw out coherent consequences from the combination of the sceptical view and EVT. The game is to see whether the consequences of this combination of views fit with the moral facts as we see them. Problems emerge when we think about situations like these:

> **Pretty Simple Island Rescue (a):** One person is trapped on Alpha island. Two are trapped on Beta. While there's time to reach one island, there's not time to reach both. A choice must be made. It's certain that we'd save one person if we made it to Alpha ($\alpha1$), the probability that the first Beta will survive if they get the drug is .5 and the same is true for the second. (These probabilities are independent. We might imagine Asclepius flips a fair coin to determine whether the medical intervention will succeed in each case.)
>
> **Pretty Simple Island Rescue (b):** Same as the above, but this time it's certain the intervention will work for one Beta and not for the other. While each person in the story has a credence of .5 that the drug will save the first Beta and .5 that it will save the second, each person is certain that if we head to Beta we'll bring about β1 and $\alpha1$ if we head to Alpha.

Personally, I don't think these cases differ. I also don't think the sceptic thinks that these cases differ. Set the Alpha option aside for one moment. Let's say that Brenda and Bernhard are two strangers located on Beta. If it helps, imagine them asleep. Suppose that when we try to intervene out of concern for each of them, we had to choose between these two lotteries. I, personally, wouldn't see any ground for preferring L(.5, Brenda, Bernhard) to L(.5, β2, β0) or vice-versa. I also cannot believe that this is where the sceptic wants to put up a fight. If one lottery were better for Brenda or better for Bernhard, they would presumably be willing to sacrifice something of value to swap. They wouldn't be willing to do that. Brenda's prospects are equally good either way. Concern for her makes me wish we had better options than these, but not that one of these is a better option. The same goes for Bernhard and my concerns for him. If you'd like a principle to refer back to, this should do:

[6] For example, if Adella is on Alpha and Bernhard and Bernadette are on Beta, I'll assume that our agent would remain indifferent between a mission to Alpha and a mission to Beta islands if, say, Adella is swapped with one of the two Betas. Remember that the people are strangers and that the sceptic harbours no bizarre bias that favours Alphas or disfavours Betas, Gammas, etc.

1/n & 1/n: For any (non-zero) number of potential beneficiaries, n, we can be indifferent between L(1/n, n, 0) and a lottery, L', in which these n individuals is each given 1/n chance of survival.

I want to use these cases to make two points. The first point is that if readers agree that the cases don't differ in any interesting way, we can just focus on lotteries with two possible outcomes where the more desirable outcome is that everyone lives and the less desirable one is that nobody does (e.g., L(1/2, β2, β0)). This seems like a philosophically harmless way of justifying a decision to use lotteries that are easier to represent. I don't think the sceptic should have to defend the view that these cases differ in some important way.[7] The second is that the following seems intuitively obvious. I think an agent would be permitted to choose $\alpha 1$ and permitted to choose L(1/2, β2, β0) out of this option set: {$\alpha 1$, L(1/2, β2, β0)}.[8] This is not intended to be any sort of counterexample to the sceptical view. So that we might refer to this later:

1&1: For any (non-zero) number of potential beneficiaries, n, we can be indifferent between the certainty of saving one person (e.g., $\alpha 1$) and the certainty of saving one person from a (non-overlapping) group of n people.

This seems perfectly compatible with the sceptical view. Still, it's here that we should see why the simple view is hopeless.

I hope that didn't seem harsh. Please remember that the view I've called hopeless isn't defended by anyone. The simple view isn't *the* sceptical view. It's *a* sceptical view. It's a view that combines a perfectly respectable moral view with another perfectly respectable view about how we should choose in the face of uncertainty. We know that we can combine independently plausible views and get something implausible. That's what the simple view is. Choices between degenerate lotteries (e.g., $\alpha 1$, which is equivalent to L(1, $\alpha 1$, $\alpha 0$)) and non-degenerate lotteries (e.g., L(1/2, β2, β0)) will quickly reveal that the sceptic cannot make sense of our intuitions using EVT. I think it might show that the sceptic will, if they assume EVT, value options and outcomes inconsistently.

Let's start with **No Heroes**. To get a feel for the intuition, imagine Adella is stranded on a nearby island and we're nearly certain to succeed if we try to save her. Agnes, we might also imagine, is stranded on an island but it's very, very far away. It's nearly certain that she'll perish before we can reach her (e.g., we might have only 1/64 chance of reaching her in time). Our boat isn't that fast and thoughts and prayers aren't that effective. We should try to save Adella. If Agnes alone were in peril and it cost little to try, our duty would be to try to save her, but that's not the case we face. It matters how we use our resources. Given our constraints, we shouldn't use the little time we have to try the very improbable when it's nearly certain we could succeed in saving a life by directing our efforts elsewhere.

Stating **No Heroes** less informally, the idea is that there should be some difference between probabilities p and p' that's sufficiently great to make it the case that L(p, $\alpha 1$, $\alpha 0$) > L(p', $\alpha' 1$, $\alpha' 0$) even if we don't agree that $\alpha 1 > \alpha' 1$ provided that we don't think that the desirability of $\alpha' 1$ is much greater

[7] Think of 'L(1/2, β2, β0)' as a lottery in which there is a probability of .5 of the first outcome (i.e., β2) and 1-.5 probability of the second outcome (β0). Any outcome that's not worse than the first outcome will be called a 'prize'. Our lotteries have two possible outcomes. Please also remember that β2 is, in this context, equivalent to saving two from Beta while the person on Alpha dies, β0 is equivalent to nobody on Alpha or Beta survives, and that similar things hold for Alpha options.

[8] This fits with my intuition and the views of Reibetanz Moreeau (1998). I think it's in line with the sceptic's view of things, too. Suppose the sceptic were to say that we can flip a coin or decide straight off to head to either island. If so, it seems that their attitude is this: $\alpha 1 \sim$ β2 ~ L(1/2, $\alpha 1$, β2).In the lottery, each person is given .5 chance of survival but this presumably doesn't make this option dispreferred to, say, the option of heading straight to Alpha' to bring about $\alpha' 1$. If this lottery gives Adella on Alpha .5 chance of survival and .5 chance of survival to Bernhard and Bernadette and it's still equi-preferred to saving Alfonse on Alpha', it seems the sceptic must agree that $\alpha 1$ is neither more desirable nor less desirable than L(1/2, β2, β0).

than $\alpha 1$. Since we expect the sceptic to agree that shuffling people around from one island to another shouldn't make much difference to the desirability of the options, it might seem the sceptic can easily explain **No Heroes**. Since they presumably deny both $\alpha 1 > \alpha' 1$ and $\alpha 1 < \alpha' 1$, they have to say that $\alpha 1 \sim \alpha' 1$. Since we can sour some mission to Alpha by decreasing the probability of its success, can't the sceptic justify **No Heroes** like this? Since $\alpha 1 \sim \alpha' 1$, we should accept that L(1/64, $\alpha 1$, $\alpha 0$)~L(1/64, $\alpha' 1$, $\alpha' 0$). Since $\alpha 1$ > L(1/64, $\alpha 1$, $\alpha 0$), doesn't it follow that $\alpha 1$ > L(1/64, $\alpha' 1$, $\alpha' 0$)?

That seems fine, but problems emerge when we think about the sceptic's attitude towards $\alpha 1$ and lotteries like L(1/2, β2, β0) or L(1/64, $\delta^{\delta}64$, $\delta^{\delta}0$). Remember that we're assuming that the sceptic agrees that $\alpha 1$ ~ L(1/2, β2, β0) (in keeping with **1/n & 1/n**). So what could the sceptic say about the following lotteries: L(1/64, $\alpha' 1$, $\alpha' 0$) and L(1/64, $\delta^{\delta}64$, $\delta^{\delta}0$)? Here, our choice is between a mission to Alpha' that's designed to save one and a mission to Delta to the Delta that's designed to save 64. EVT tells us that these lotteries differ in desirability iff there is a difference in the probabilities of outcomes or differences in the values of the outcomes. The probabilities of the more desirable and less desirable outcomes are the same. The sceptic's view is that $\alpha' 1 \sim \delta^{\delta}64$. If the sceptic agrees L(1/64, $\alpha' 1$, $\alpha' 0$) ~ L(1/64, $\delta^{\delta}64$, $\delta^{\delta}0$) and agrees L(1/64, $\delta^{\delta}64$, $\delta^{\delta}0$) ~ $\alpha 1$, they have to accept L(1/64, $\alpha' 1$, $\alpha' 0$) ~ $\alpha 1$. That means they have to deny **No Heroes**.

It might be helpful here to notice that the sceptic is committed to this principle, which is similar to **1/n & 1/n**:

> **1 & 1/n**: For any number of potential beneficiaries, n, we can be indifferent between a degenerate lottery in which one person (not in n) is rescued and L(1/n, n, 0).

The sceptic cannot deny, for example, that just as the certainty of saving one on Alpha and one of two from Beta isn't better or worse than the lottery in which each Beta is given .5 chance of survival, the same would hold for $\alpha 1$ and any lottery that gives n other individuals 1/n chance of survival. That might seem fine, but given the sceptic's views about the desirability of saving one from Alpha' and 64 from Delta to the Delta, the sceptic must reject EVT if she's going to avoid saying that the 1/64 chance of saving one person on Alpha' is as desirable as the certainty of saving one on Alpha since that, given **1 & 1/n**, is equi-preferred to the 1/64 chance of saving 64 on Delta to the Delta.

This is bad. They also have to deny **No Daredevils**. The above established the former. Remember that the above established $\alpha 1$ isn't more desirable than L(1/64, $\alpha' 1$, $\alpha' 0$). Since L(1/64, $\alpha 1$, $\alpha 0$) ~ L(1/64, $\alpha' 1$, $\alpha' 0$), we cannot consistently say that $\alpha 1$ > L(1/64, $\alpha 1$, $\alpha 0$). But that's just the denial of **No Daredevils**.

As for **No Monsters**, remember that the sceptic agrees that $\alpha 1$ ~ L(.5, β2, β0). They also insist that $\alpha 1$ ~ β2. Thus, by the transitivity of indifference, they have to accept that L(.5, β2, β0) ~ β2. Given **1/n & 1/n** gives us that L(.5, β2, β0) is equi-preferred to the lottery in which, say, we flip a coin to decide which Beta to save, we get that β2 is equi-preferred to flipping a coin to decide which Beta to save or deciding to take the boat that's not big enough to bring each Beta back to safety. The simple view cannot justify **No Monsters**.

We can come at this from a different angle. Given the resources of EVT, we cannot distinguish between an agent who is indifferent between $\alpha 1$ and β2 and an agent who thinks that the additional lives we save once we've saved one have zero marginal value.[9] That's a monstrous attitude, however. Since that's surely not what the sceptic thinks, something has gone terribly wrong.[10]

---

[9] Assuming that this agent has the sceptic's preferences, there's an indifference curve that runs through $\alpha 1$ and β2. We then see that an indifference curve runs through L(.5, β2, β0) and $\alpha 1$. If we assume that it's better to save two Betas as opposed to just one, these curves cannot intersect. The items on the curve with β2 should (for Pareto reasons) be strictly preferred to any on the curve with L(.5, β2, β0), so we cannot put $\alpha 1$ on both curves.

[10] Here, as Kripke might put it, is something like a proof that we should save the greater number if we assume EVT, **1&1**, **No Monsters**, and **No Daredevils**. According to **1&1**, $\alpha 1$~L(.5, Bernhard, Bernadette). According to **No Daredevils**, L(.5, Bernhard, Bernadette) < L(.6, Bernhard & Bernadette,

3.1 Modifying the simple view with buckets

This project began with this rather naive worry about the sceptical view. If I were to explain why, say, in heading to Beta, it's better to use the drugs that are quite likely to save Bernadette and quite likely to save Bernhard as opposed to, say, a drug that's quite likely to save Bernadette but not very likely to help Bernhard, I would say that it's better to save both than to save just one or save neither. And then I'd do some simple math. This kind of explanation only works if the numbers of beneficiaries sometimes matter. It was then quite tempting for me (who admittedly isn't a sceptic) to think that since β2 > β1, β2 > α1 because there are some missions that are sure to result in β1 such that β1 ~ α1. (Think about L(.5, Bernadette, Bernhard).) This, obviously, is a problem for the sceptic. It then seemed that someone would be rightly indifferent between α1 and β2 iff they might be rightly indifferent between β2 and β1. Since it's monstrous to be indifferent between β2 and β1 and irrational to be uninterested in the chances of success in choosing between missions, it seems that the only sensible non-monstrous views will say that all the numbers matter. Given the constraints imposed by EVT, it seemed we were stuck saying that the numbers that represent probabilities matter iff the numbers of potential beneficiaries matter, too.

In setting up the issue this way, friends have accused me of overlooking an important complication. I've failed to take account of a crucial feature of the sceptical view. I asked readers to consider and compare options like these: α1, β2, and L(.5, β2, β0). In doing so, I said nothing about the conflicts of interests that can arise if we're choosing between some options (e.g., Alpha and Beta options) and not others (e.g., multiple Beta options). The sceptical view, I've been reminded, is a view about *conflict* cases, cases where interests of potential beneficiaries come into conflict. The sceptic agrees, they add, that we should save the greater number when interests don't come into conflict. I've muddied the waters, they say, by ignoring the difference between conflict cases and non-conflict cases.

What can I say to this accusation? I'm guilty as charged. I didn't take account of this complication in setting out my arguments against the sceptical view. Here's a confession. I left this detail out *intentionally*. I thought it would be helpful to see why a view that didn't take account of this distinction would run into trouble and then show that taking account of it didn't sufficiently improve matters.

How should we reformulate the sceptical view if the sceptic thinks (in keeping with **No Monsters**) that we should take the numbers into account when choosing between Beta missions but shouldn't take them into account in choosing between Alpha and Beta missions? We can try to reformulate the sceptical view by introducing buckets. For each option, we can identify the individuals who might benefit or might be spared a harm if that option is chosen and we can say that there's some bucket that contains these and only these individuals. We don't want to say that the numbers matter when choosing *between* buckets, but they matter when choosing what to do for the people *within* a bucket. In this way, we can try to both do justice to the idea that β2 is more desirable than β1 (e.g., because it is Pareto superior) without any commitment to the view that β2 is more desirable than α1. Thus, the bucket theorist can try to convince us that we're allowed to choose β2 or α1 out of the choice set {β2, α1} but not permitted to choose β1 out of {β2, β1}.

It does seem that the introduction of buckets helps with this problem, but how should we flesh out the details of the view. Remember that the sceptic must, if she accepts **No Daredevils**, think that α1 is superior to any lottery L(p, α1, α0) where p < 1, that β2 is superior to any lottery L(p, β2, β0), and so on. Given that the bucket theorist accepts EVT, she must then also tell us how to compare degenerate lotteries to degenerate lotteries (e.g., α1, β2) and lotteries that concern individuals in

---

Bernadette). According to **No Daredevils**, L(.6, Bernhard & Bernadette, Bernadette) < β2. Thus, α1 < β2. Similar reasoning will establish that β2 < γ3, γ3 < δ4, etc. The argument makes some lightweight assumptions about diminishing marginal utility (e.g., that, in keeping with **No Monsters**, β2 > β1). We can block this argument by distinguishing indifference from equi-preference, but only if we abandon EVT.

different buckets (e.g., L(p, β2, β0), L(p, α1, α0)). It looks as if the bucket theorist thinks that the most desirable missions will be the degenerate lotteries and that every non-degenerate lottery must be strictly dispreferred. From here, the precise details aren't obvious, but we have enough before us to see why the bucket theory won't do justice to our intuitions.

In the section above, I argued that the simple view cannot justify **No Monsters** relying on this sort of argument:

P1. α1 ~ β2.
P2. L(1/2, β2, β0) ~ α1.
P3. If α1 ~ β2 and L(1/2, β2, β0) ~ α1, β2 ~ L(1/2, β2, β0).
C. So, β2 ~ L(1/2, β2, β0).

Because of EVT, we cannot deny the transitivity of indifference. The bucket theorist must accept P3 and reject either P1 or P2. The bucket theorist has, so far, only insisted that β2 > L(1/2, β2, β0), but they haven't said which premise they'd reject in the argument for this conclusion or offered any grounds to justify that rejection.

They cannot reject P1 because that's the sceptical view. P2 remains. I can't see how the sceptic can plausibly deny P2. Imagine an angel tells you that if you head to Alpha or to Beta you can save a stranger's life but confesses that there are two people on Beta and they're not sure which person you'll save. They recommend trying to save both if we head to Beta while preparing for partial success. Once you reconcile yourself to the fact that a mission to Beta will only be partially successful, I don't see why you should disprefer L(.5, Bernhard, Bernadette) to α1 unless you've decided that your primary concern is with the difference between partial and full success rather than the interests of the individuals who need your help.

The bucket theorist can only reject the argument's conclusion by rejecting P2. This, in my view, is implausible. Here, I think the sceptic and I would agree. Remember that the sceptic thinks that we're supposed to look at things from the perspective of potential beneficiaries. If it's a fact that we cannot save Bernard while trying to save Bernard and Bernadette on Beta, how would this make it less important to try to save Bernadette? Why would the presence of someone who cannot be saved give us reason to head to Alpha instead of Beta? The bucket theorist makes the numbers matter, but in a perverse way. Instead of seeing the presence of additional people on Beta as a reason to give missions to Beta preference, it sees the addition of persons as people we might fail to save as a reason to focus our efforts on smaller groups where complete success is more likely. This makes numbers matter in a strange way. Given normal assumptions about the correlation between chance of complete success and number of people in a bucket (e.g., that the chance of losing one will increase as the numbers of those in need increase) the numbers function like a disincentive to try to save the larger group.

There's a further problem with the bucket view. Let's imagine three islands, Alpha, Alpha', and Beta where Beta is in the middle and Alpha is off to the West. You can choose between missions that might bring about α1, α'1, or β2 and some missions might result in the rescue each person on Alpha and Beta, some might result in the rescue of each person on Alpha' and Beta, but no mission can reach Alpha and Alpha'. This gives us five buckets:

① The person on Alpha;
② The two on Beta;
③ The person on Alpha';
④ The people on Alpha and Beta;
⑤ The people on Beta and Alpha'.

The sceptic presumably believes① ~ ② ~ ③. Additionally, they accept ④ ~ ⑤. In keeping with **Double Sweetening**, we should see some souring of ④ as more desirable than ① or ② alone and some souring of ⑤ as more desirable than ② or ③ alone (e.g., adding any lottery that gives us a non-zero chance of saving the person on Alpha to ② should sweeten ②, but that's equivalent to some souring of ④ in which the prospect of saving the person on Alpha becomes uncertain).[11] Here's the problem. The sceptic must also agree that saving the people on Alpha and on Beta is *not* more desirable than saving the person on Alpha', so that gives us ④ ~ ③. Given this, some souring of ④, ④$^-$, would be dispreferred to ③ even though this souring of ④ would be strictly preferred to ① and strictly preferred to ②. Given these failures of transitivity, the bucket theorist cannot to use EVT to justify our answers to the easy questions.

3.2 Modifying the simple view with tickets

The bucket view turns out to care about the numbers in a perverse way. When I try to think about how a sceptic might try to take account of the probabilities and avoid the difficulties that arise for the bucket view, one natural thought is to build the account around the idea that the strength of a reason to try to save a person waxes and wanes as the probability of being able to successfully save them if we try increases or decreases. (This is in keeping with **No Heroes** and **No Daredevils**.) If we want to focus on individuals rather than groups, we could rank options by first finding the option or options in which the person 'easiest' to save is located. To make this vivid, imagine each person has a ticket with a number representing their best prospect of rescue and we choose options by first locating the ticket with the largest number and then committing to saving as many as we can co-located with the person with this ticket (in keeping with **No Monsters**). How would the ticket view fare?

Poorly. First, it seems to tell us that there's an important difference between missions in which we're certain to save someone or other and missions in which it's certain that some particular person will be saved. It's a cost of the view that it says that we cannot choose a lottery in which each Beta gets an equal chance of survival when choosing between this lottery and $\alpha 1$ but can choose, say, a mission in which it's certain that a Beta will be saved if we head there (but uncertain which) when choosing between this lottery and $\alpha 1$.

This argument, in my view, is sufficient to show that the ticket view fails:

P1. We may choose either L(1/2, β2, β1) or $\alpha 1$ out of {L(1/2, β2, β1), $\alpha 1$} iff $\alpha 1$ ~ L(1/2, β2, β0).

P2. We may choose either β2 or $\alpha 1$ out of {β2, $\alpha 1$} iff β2 ~ $\alpha 1$.

P3. We may choose either L(1/2, β2, β1) or $\alpha 1$ out of {L(1/2, β2, β1), $\alpha 1$} and may choose β2 or $\alpha 1$ out of {β2, $\alpha 1$}.

C. So, β2 ~ $\alpha 1$ and $\alpha 1$ ~ L(1/2, β2, β1).

We surely must reject the argument's conclusion. Unfortunately, the ticket view commits us to P3 and EVT gives us P1 and P2.

The only way the sceptic might avoid this is by revising the view again so that the probability of saving the second person helps us break ties. The problem with *this* view emerges when we compare missions to Alpha, Beta, and Beta'. We initially contemplate a choice between $\alpha 1$, a mission to Beta in which we're certain to save Bernhard and have a .6 chance of also saving Bernadette (L_B), and a mission to Beta' in which we're certain to save Binh and there's a .5 chance of saving Benji (L_B'). When we compare $\alpha 1$ to L_B, the sceptic shouldn't say that $\alpha 1$ is superior to L_B. They're either tied or L_B is superior. When we compare $\alpha 1$ to $L_{B'}$, they are either tied or $L_{B'}$ is superior. The sceptic cannot say that successfully saving two on Beta or Beta' is superior to saving one on Alpha, so we're stuck saying that $\alpha 1$ ~ L_B and $\alpha 1$ ~ $L_{B'}$ even though the ticket theory that looks at the second ticket will say that $L_B > L_B$'. This violates EVT.

3.3 The Lesson

I've tried to find some way for the sceptic to use EVT to justify our answers to the easy questions. I've failed. Is the failure my fault? Do we need better props now that the buckets and tickets have failed us? I doubt it. We see that β2 must be superior to lotteries like L(1/2, β2, β0) or L(1/2, β2, β1) and see

[11] Think back to your reaction to Simple Island (routes) and the idea that we should choose a route that creates a non-zero chance of rescuing an additional person instead of one that does not.

that options like $\alpha 1$ cannot be superior to options like L(1/2, $\beta 2$, $\beta 0$) or L(1/2, $\beta 2$, $\beta 1$). We cannot stop once we've made these comparisons. EVT forces us to compare $\beta 2$ and $\alpha 1$. It gives us three options to choose from: $\alpha 1 > \beta 2$, $\beta 2 > \alpha 1$, and $\alpha 1 \sim \beta 2$. Of these, only the last option is compatible with the sceptical view. The only view compatible with our intuitions about the easy cases (and the general lessons we draw from them like **No Heroes** or **No Monsters**) is that $\beta 2$ has greater expected value than $\alpha 1$ because it has greater expected value than lotteries that have the same expected value as $\alpha 1$.

We can now see that the sceptic must dispense with EVT.[12] We cannot use EVT unless we're willing to make comparisons that I think the sceptic thinks we shouldn't make. If the sceptic truly believed that the trichotomy thesis held (i.e., every pair of options is either equally desirable or one is superior) and believed that all rational agents have complete preferences, then they would have to say that even the smallest difference between missions to Alpha and Beta might be decisive. That, to me, doesn't seem to capture their outlook. If, in choosing between boats to Alpha, you told the sceptic that this boat is just *slightly* more reliable, they'd choose it. In choosing between whether to take a boat to Alpha or Beta you tell them that it's just *slightly* more likely that they'd reach one destination in time than the other, I doubt that this would decide the matter for them. If, however, they respond to mild sweetening and souring in one case and not in the other, that's a sign that the sceptical agent doesn't equi-prefer $\alpha 1$ and $\beta 2$ even though they don't strictly prefer either $\alpha 1$ or $\beta 2$ (Savage, 1954). They need a decision theory that can tell us what's rational for an agent to do when they're certain that sweetening or souring missions to Beta make those better or worse but has no view about whether the sweetened or soured options are consequently more or less desirable than some missions to Alpha or Gamma.

4. Scepticism and Prospectism

The sceptic cannot use EVT to justify our answers to the easy questions if they deny, as they probably will and should, that every option and outcome can be compared. If some sceptic were to say that their indifference between some Alpha mission and Beta mission remains even after we sweeten one of them (e.g., by increasing the chance of success, by adding causal pathways so that choosing one might save each Beta and, perhaps, a Gamma or two), their preferences cannot be represented by any utility function. It's a sign that they think that $\alpha 1$ and $\beta 2$ aren't comparable (i.e., $\alpha 1 \bowtie \beta 2$) and that the sceptical agent will have a corresponding preferential gap (i.e., $\alpha 1 \parallel \beta 2$). It's not true that the sceptic weakly prefers one of these options. We cannot give a cardinal representation of the desirability of these options.

At this point, there might be some who think that the sceptic's cause is hopeless because the above shows that they cannot use EVT. This isn't how I see things. I won't assume that all rational agents have complete preferences, so I don't assume that EVT is the only game in town. We should consider weakening its constraints. Some of its constraints are controversial and weakening the controversial ones seems to be in keeping with the sceptic's outlook. Prospectist decision theories are designed to accommodate preferential gaps and they seem to explain crucial intuitions about sweetening that caused trouble for the combination of scepticism and EVT. We'll look at proposals from Hare (2013) and Kaivanto (2017) to see if their prospectist views might give the sceptic the tools they need.

Let's start with the simplest prospectist view:

> Simple Prospectism: It is permissible for an agent to choose an option iff there is some utility function, U', that's a coherent completion of this agent's preferences and no alternative to this option that has higher expected U'-utility (Hare 2013: 53).

[12] This feels like something the sceptic should do given Doggett's (2013) arguments and the arguments for thinking that the sceptic must reject transitivity found in Muñoz (forthcoming). It seems to fit with Tenenbaum's (forthcoming) suggestion that we cannot make the necessary comparisons when thinking about the importance of rescuing people on the grounds that we have dignity and not price.

Imagine some morally conscientious agent who doesn't weakly prefer some Alpha mission to some Beta mission or vice-versa (e.g., L(.9, β2, β0)∥L(.9, α1, α0). Imagine, if you can, that this agent lacks this preference because this agent recognises that neither α1 nor β2 is more desirable and they aren't equally desirable (L(.9, β2, β0)⋈L(.9, α1, α0)). When an agent has a preferential gap like this, we expect that they might not be moved by some mild sweetening or souring. If two options are equi-preferred, the slightest sweetening should induce a strict preference. Consider, for example, α1 and $\alpha 1^+$ (where the latter is just α1 with an additional causal pathway that gives us some non-zero chance of bringing about β2). Consider, for another example, L(.8, β2, β0) and L(.9, β2, β0). In both cases, we should prefer the sweetened Alpha option to its unsweetened alternative and the sweetened Beta option to its unsweetened alternative. If, however, there is a preferential gap so that L(.9, β2, β0)∥L(.9, α1, α0), we should expect that it's possible to sweeten the former without making it strictly superior to the latter and vice-versa (e.g., L(.95, α1, α0) might be superior to L(.9, α1, α0) without being superior to L(.9, β2, β0)). The simple prospectist welcomes agents with gaps and seeks to explain why such agents should choose L(.95, α1, α0) out of {L(.95, α1, α0), L(.9, α1, α0)} without denying that they might choose L(.9, β2, β0) out of sets that contain these other options.

When an agent's preferences display this kind of incompleteness, the prospectist doesn't think the agent needs to acquire new preferences by, say, reflecting harder. The prospectist notes that if the agent's preferences are otherwise coherent, there should be coherent completions that some hypothetical agent could have that would overlap with this agent's preferences. These sets of hypothetical preferences, in turn, tell us what this (non-hypothetical) agent would be rationally permitted to choose. To complete the agent's preferences, each preferential gap (e.g., α1∥α'1) is 'filled' (e.g., we replace α1∥α'1 with $\alpha 1 > \alpha' 1$, $\alpha 1 < \alpha' 1$, or $\alpha 1 \sim \alpha' 1$) and for each such completed set we'll have a new utility function (U', U'', etc.) that represents this agent's preferences. Our agent is permitted to choose an option if there is some such utility function, U', where this option maximises expected U'-utility. (She's required to choose an option if it maximises expected utility relative to every such utility function and forbidden from choosing those options that maximise expected utility relative to no such function.) Note that on this view, the permitted options will always be most preferred (in line with EVT). The key difference is that there are many preference relations to consider (i.e., the preferences of hypothetical agents suitably related to our agent), not just one (i.e., the preferences of our actual agent, which might include gaps).

If we replace EVT with the simple prospectist view, it seems the sceptic might avoid the objections from above and so say that these problems tell us more about the suitability of using EVT than it reveals anything about their moral outlook. In formulating each of the objections above, we assumed that the sceptical agent would have some single utility function, U, and tried to create trouble for the view by showing that there's no utility function that assigns the same expected value to L(.5, β2, β0) and α1 that also vindicates **No Heroes**, **No Daredevils**, etc. Freed from the commitment to explain all of our intuitions using some single utility function, the sceptical view might fare better.

It will be helpful to consider an alternative prospectist view:

> Ensemble Prospectism: It is permissible for an agent to choose an option iff it is selected by a non-stochastic ensemble-voting rule applied to the ensemble of rankings associated with coherent extensions of the agent's preferences (Kaivanto, 2017).

A non-stochastic rule will (non-randomly) choose a single option out of an option set by taking account of the recommendations of each coherent completion of the agent's preferences. For ease of exposition, we'll assume that it uses the Borda rule, but nothing here turns on which rule we choose.[13] Remember that EVT tells us to choose from among the most preferred options. The prospectists think that there might not be any options that are most preferred. In our cases, the closest thing we get might be options that are most preferred given some set of hypothetical preferences. The simple

[13] Unless the rule we choose is to select one completion to function as a dictator.

prospectist tells us we're free to choose an option that's most preferred by some coherent completion of our preferences. The ensemble prospectist view differs from this in two key respects. First, some voting rules (e.g., Borda) will identify losing options that are most preferred by some coherent completion and will identify some winning options that aren't most preferred on any coherent completions. Second, ensemble prospectism is *resolute* in that it identifies a single winner in each option set, but simple prospectism gives agents with incomplete preferences much more latitude. I can see the merits of both approaches, but I think this last difference shows that neither approach suits the needs of the sceptics.

If the sceptic is going to meet the explanatory challenge we've set out for them, they need to explain sets of claims that couldn't be explained using EVT:

1. Given {α1, β2}, an agent is permitted to choose α1, β2;
2. Given {α1, L(.5, Brenda, Bernhard)}, an agent is permitted to choose α1, L(.5, Brenda, Bernhard);
3. Given {β2, L(.5, Brenda, Bernhard)}, an agent is permitted to choose only β2.
4. Given {α1, β2, L(.5, Brenda, Bernhard)}, an agent is permitted to choose only α1, β2.

We cannot explain (1)-(4) using EVT because we get a violation of Sen's (2017) principle beta.[14] There is no utility function that rationalises these claims about permissible choice, but the sceptic might have some luck using prospectist decision theory to explain (1)-(4). They can now try to explain the data using multiple utility functions as inputs.

Let's see how the sceptic might use the simple prospectist view to try to justify **No Monsters**. We begin by noting the sceptical agent's attitudes with respect to these pairs:

α1 ‖ β2

α1 ‖ L(.5, Brenda, Bernhard)

β2 > β1 (e.g., save Bernhard but sacrifice Brenda)

Note that while every completion of the agent's preferences will include β2 > β1, some will include α1 < β2 and some will include α1 > β2. Let's imagine that one completion (U′) says that it is three times more desirable to save the Alpha than it is to save a Beta and another (U′′) that says that the inequality runs the other way. Given that the set of possible completions includes U′ and U′′, our agent with preferential gaps would be permitted to choose α1 out of any option set that also included β2 and/or L(.5, Brenda, Bernhard). This vindicates (1), (2), and part of (4). Given that there are some completions on which the reverse inequality holds, this gives (1), (2), and the remaining part of (4). We get (3) because every extension includes β2 > β1. In the setting where we're assuming EVT, we get to use only *one* utility function to determine what's permitted, but in this setting, we get to use sets of utility functions without any requirement that we use the same one for each choice. This gives us the flexibility needed to vindicate (1)-(4) and shows that the sceptic can justify No Monsters using a suitable decision theory.

This seems promising. Initially.

4.1 Prospectism and Scepticism Evaluated

---

[14] To illustrate this, consider Sen's Principle Beta. It says that if two options, X and Y, can permissibly be chosen from some set of options, it will be permissible to choose X iff it's permissible to choose Y from any set of options that's an expansion of this first set. In our discussion of No Monsters, we've seen how this might be violated by the sceptic. Whilst it should, on the sceptical view, be permissible to choose α1 or L(.5, Brenda, Bernhard) from {α1, L(.5, Brenda, Bernhard)}, the sceptic must deny that it's permissible to choose L(.5, Brenda, Bernhard) from {α1, β2, L(.5, Brenda, Bernhard)} even though she must insist that it's permissible to choose α1 from this set. If adding β2 makes it impermissible to choose α1, it seems we get something like the denial of the sceptical view. If adding β2 doesn't make it impermissible to choose L(.5, Brenda, Bernhard), we get the denial of No Monsters.

There's a familiar problem with simple prospectism. It faces a sure loss objection. Ensemble prospectism doesn't face this problem. The ensemble prospectist's fix makes it unsuitable for the sceptic. My hunch is that the first view gives the sceptic the flexibility they need to try to explain (1)-(4) but then its problem with the sure loss objection shows that it's incapable of explaining **No Daredevils**, **No Heroes**, etc. The second view avoid some of these problems because it doesn't give agents the same latitude to choose between missions that might benefit different groups of people, but then we see conflicts between the implications of ensemble prospectism and the sceptical view.

Peterson (2015) observes that the simple prospectist view has trouble with money pumps. Here's a version of the objection. Let's assume that our agent has the expected preferential gap (i.e., $\alpha 1 \parallel \beta 2$) and that when such gaps exist, there will be further gaps when it comes to sweetened or soured versions of these options (i.e., if $\alpha 1 \parallel \beta 2$, then $\alpha 1^{-} \parallel \beta 2$, $\alpha 1 \parallel \beta 2^{-}$, $\alpha 1 \parallel \beta 2^{--}$, $\alpha 1^{--} \parallel \beta 2$).[15] We will sweeten or sour our options using changes to the chance of success instead of cash. Let's assume that some slight differences in the chance of success in Alpha or Beta missions doesn't create any requirement that the agent prioritise missions to Alpha over Beta (or vice-versa). With this in mind, we can now state the objection. The agent might be offered an initial choice between $\alpha 1$ and $\beta 2^{-}$ and permissibly choose the latter. She might be offered a choice between $\beta 2^{-}$ and $\alpha 1^{--}$ and permissibly choose the latter. Through a series of permissible choices, she can end up with a soured option $\alpha 1^{--}$ when she could have kept $\alpha 1$ initially. Since simple prospectism sanctions both trades, it seems to sanction a series of choices that result in foreseeable sure loss where the loss in question is an increased risk of a failed mission. If the simple prospectist cannot block this objection, they cannot explain **No Daredevils**, but this should hold in the case of sequential choice.[16]

We know that some philosophers think we can avoid sure loss by, say, being resolute or by taking account of our past choices, but remember that the sceptic presumably thinks that (1)-(4) will continue to hold *even if* some agent has already decided, say, to help the Alphas. Just imagine, for example, that the ship's captain was momentarily confused and decided to head to Beta but was actually on course to Alpha. I take it that the sceptic would say that such a captain could permissibly decide to change course or remain on their present course. We don't want the captain's set of permissible options to contract just because they've made some past choice provided that their remaining options include Alpha missions and Beta missions that were, initially, permissible objects of choice.

Once we see why the combination of the sceptical and simple prospectist views runs into trouble, it should be clear that ensemble prospectism isn't the cure. Ensemble prospectism gives an agent the same directives at each choice point in cases of sequential choice. It thus seems to rule out in advance sanctioning the swaps that result in foreseeable loss. Using the Borda count, for example, if the agent is presented with {$\alpha 1$, $\alpha 1+$, $\beta 2$}, the agent will not be permitted to choose $\alpha 1$ at any point because any non-stochastic voting rule that selects $\alpha 1+$ or $\beta 2$ out of some set *thereby* makes it impermissible to choose the other when a switch is offered. Let's suppose that our agent is like this:

---

[15] This seems harmless enough. Remember that we've abandoned EVT, in part, because we agreed that the sceptic would want to say that she doesn't prefer $\beta 2$ to $\alpha 1$ *and* doesn't prefer $\beta 2$ to some lottery L(p, $\alpha 1$, $\alpha 0$) where p is sufficiently close to 1 and knows enough math to know that there must be some p' that sits between p such that (a) L(p', $\alpha 1$, $\alpha 0$) > L(p, $\alpha 1$, $\alpha 0$) and (b) $\beta 2$ is not strictly preferred to either. This gives us a pair of Alpha options such that one is strictly preferred to the other where the agent doesn't weakly prefer the Alpha options to $\beta 2$ or vice-versa.

[16] Given space limitations and my inability to improve upon extant work, I will not discuss the various moves and countermoves that someone might make to address the sure loss argument. I think that using the pumps designed by Gustafsson (forthcoming) defeat attempts to use inductive techniques to avoid sure loss. There are further objections to simple prospectism in the literature worth considering (e.g., those by (Schoenfield, 2014)) but they don't speak all that directly to the issues discussed in this paper and my objection isn't to prospectism, *per se*, but to the idea that the sceptic can meet our explanatory challenge using some form of prospectist view.

$\alpha 1 \parallel \beta 2 \parallel \alpha 1^{+} > \alpha 1$. We can complete this agent's preferences in three ways: $\alpha 1^{+} > \beta 2 > \alpha 1$, $\beta 2 > \alpha 1^{+} > \alpha 1$, and $\alpha 1^{+} > \alpha 1 > \beta 2$ and $\alpha 1^{+}$ has the highest Borda score. The view doesn't sanction a series of choices (including switches) that could lead to sure loss precisely because it doesn't let the agent choose a Beta option when a sweetened Alpha option is offered. In other words, the view was *designed* to explain why it's not rational to make the choices that, according to (1)-(4), an agent should be free to make.

There are two further problems with the way that ensemble prospectism handles the sure loss problems. These are more 'philosophical' than technical. There seems to be at least some superficial similarity between our prospectist views that try to explain intuitions about rational choice without recourse to some single utility function and views that try to explain intuitions about rational choice without assuming that an agent has a set of precise credences. One way to approach the latter problem is to use a kind of supervaluationist approach (Mahtani, 2019). We could try something similar in our case. We can treat the possible completions of the agent's preferences as ways of removing a kind of indeterminacy in the agent's evaluative outlook. When the supervaluationist does this for terms like 'bald' or 'pink', however, they focus on *admissible* ways of making things precise. We should have admissible precisifications that classify some single person as 'bald' and 'non-bald' if this person is a borderline case but not otherwise. Penumbral connections should be preserved. Should we see every coherent completion of the agent's preferences as admissible or should some completions be struck out as inadmissible?

Let's say that the coherent completions of the agent's preferences are *egalitarian* iff they present the rescue of each stranger as equally important and *inegalitarian* otherwise. For the prospectist to explain why it might be permissible to save one person on Alpha instead of sixty-four on Delta to the Delta, we need some admissible completions to be radically inegalitarian (e.g., ones that 'say' that it's vastly more important to rescue this stranger than any other stranger we might help). I think it's rather odd, however, that such preference relations could have any bearing on what the agent ought to do precisely because they are so radically inegalitarian. Isn't any reasonable agent *certain* that it cannot be so much more important to rescue this one person on Alpha than anyone else? If we say, on the other hand, that only the egalitarian preference relations are admissible, we can still allow some room for vagueness. It might be clear that when choosing between saving a person from death and a person from a papercut, we should prefer to save the person from death. If the harms that a person faces might cause enough suffering, it might be borderline whether it's more important to save this person than another person from death. If, however, we're stuck with just the egalitarian preference relations, we're saddled with a view that's indistinguishable from the simple view when it comes to the cases that we're interested in (i.e., cases where many strangers face death and we have to choose between saving the many or the few). At any rate, if the sceptic thinks, as I suspect they do, that we should ignore all the promptings from inegalitarian preferences, I suspect they wouldn't want to use the prospectist machinery to explain why their moral intuitions are correct. But if they ignore this advice and say that some radically inegalitarian completions are admissible, they'll fail to explain **No Heroes** since it will be just as easy to cook up preference relations relative to which $L(1/63, \alpha 1, \alpha 0) > \delta^{\delta}64$ as it is to find ones relative to which $\alpha 1 > \delta^{\delta}64$.

Let's consider one final problem for the prospectist. I don't see why any sceptical agent would feel any rational pressure to follow the guidance of the ensemble prospectist's choice rule. Remember that they needed a choice rule to deal with problems involving sure loss that would forbid choosing an option that's not less desirable than the option it requires us to choose (e.g., it has to say that we cannot be permitted to choose $\alpha 1^{--}$ or $\beta 2$ out of $\{\alpha 1, \beta 2\}$ and $\{\alpha 1^{--}, \beta 2\}$ respectively even though the agent is indifferent between the options presented at each choice point *and* might be permitted to choose, say, $\beta 2$ if it's certain that no swaps will be offered after an initial choice). When offered the chance to swap things when indifferent between them, the ensemble prospectist doesn't want to say that the sole reason that an agent would turn this down is that they must turn this down to avoid choice-sequences that result in sure loss. One good answer might be this. If we thought that the agent was *uncertain* about the desirability of different options, some choice rules might handle this uncertainty better than others. Ensemble prospectism can be defended on these grounds. We, however, are interested in agents who are *certain* that neither option they're presented with is more desirable than the other, so this justification is off the table. I cannot imagine what alternative justification might be given.[17]

To sum up, the simple prospectist views gives agent's too much latitude to try to justify **No Daredevils** and **No Monsters**. The ensemble prospectist view fares better, but it does so by removing the latitude that the sceptic wants (i.e., the freedom to try to save someone on Alpha or the people on Beta when there aren't dramatic differences in the chances of our missions succeeding). Neither view vindicates **No Heroes**.

At this point, we shouldn't be surprised if the sceptic were to say something more radical in response. They might say that the above just shows that the prospectist view remains too close to EVT because it still tries to use comparisons about the importance of rescuing distinct individuals.[18] They might say that if we really grasp the consequences of seeing the reasons to rescue distinct individuals as incomparable, we wouldn't try to appeal to these sets of hypothetical preferences to try to show (in sceptic-friendly terms) why claims like **No Heroes** are correct. We might just settle for a view on which some option is more desirable than another iff we can use Pareto principles to explain why that's so. This might give us **No Monsters** and **No Daredevils**, but the sceptic might decide that **No Heroes** isn't something they'd want to explain.

My main concern with this response is that this saddles the sceptic with the commitments that I feared were theirs. It suggests that the sceptic might have a perfectly coherent view that's incapable of explaining some of our answers to easy questions. The explanatory shortcomings of this view won't be quite as dramatic as the views that tried to play the game, but they are, nevertheless, shortcomings. It seems we might culpably mismanage our limited resources if we use them to carry out nearly hopeless missions instead of those nearly certain to succeed.

5. Numbers without aggregation?

Let me try to identify the source of the sceptic's problems. If the sceptic takes on board any of the tools that we've developed for thinking about how we ought to choose in the face of uncertainty, they'll have to accept these two claims:

> **Necessary Compensation**: If L1 is strictly preferred to L2, either L2's prizes are less desirable or the probability of winning L2 is greater to a sufficient degree.
>
> **Discounting**: The strength of a reason to choose L1 decreases as the probability of winning L1's prize decreases.

Given **Necessary Compensation** and **Discounting**, the sceptical view runs into trouble because the sceptic accepts **Double Sweetening** in some cases while rejecting it in others. To vindicate **No Monsters**, they have to admit that the prospect of additional lives can sweeten a prize. Crucially, however, they have to say that it can do this iff it yields a prize that's Pareto superior to the original prize. (Similarly, they must say that increasing the probability of winning a lottery's prize can sweeten in some contexts (in order to justify **No Daredevils**) but then deny that it can sweeten an option in others. Everything turns on whether we're comparing the chance of success of saving the same individuals.) In a setting where all prizes can be compared (e.g., when we're trying to use EVT), this attitude towards sweetening is incoherent. With **Double Pareto Sweetening** (i.e., the idea that we can make a lottery more desirable if we either make someone better off without making anyone worse

[17] In personal communication, Kaivanto confirmed that he saw his proposal as offering a theory of rational choice that addresses the problems caused by preferential gaps and *not* those associated with incomparability. Thus, the justifications he might offer for following the guidance of the voting rule (e.g., that it helps us track what's valuable) might explain why we should adhere to the rule's recommendations. This justification isn't available to the sceptic.

[18] I can imagine a sceptic saying this if they endorse Muñoz's (n.d.) claims about comparability or Tenenbaum's (n.d.) Kantian claims about dignity and price.

off or increase the odds that some will be better off without decreasing them for any of these individuals) and the sceptical view, we'll get violations of EVT because we'll be forced to use an inconsistent set of instructions for representing the desirability of options (e.g., while holding the desirability of $\alpha 1$ fixed and insisting that β2 is equally desirable, we can sour some Beta option such that it's just as desirable as $\alpha 1$ and strictly less desirable than β2 (e.g., a lottery in which it's certain that some Beta will be rescued but uncertain which will be)).

The obvious 'fix' is to drop the assumption that we can compare prizes when they aren't seen as Pareto sweetenings or sourings of each other. This requires us to change our decision-theory to something like prospectism, but no such view can justify **No Heroes**. The prospectist must admit that a rational agent might be indifferent between saving one person or saving a much larger group that this person isn't part of iff this agent would be indifferent between a mission that's certain to save any number of people and a mission that's nearly certain to benefit nobody. It's obvious how the sceptic could avoid this. They could impose restrictions on admissible hypothetical preference relations so that the agent ignores all hypothetical preferences that prioritises the rescue of some strangers over others, but this leads us right back to EVT. We get the flexibility in prospectist decision theories that the sceptic needs iff we don't impose these restrictions on admissible hypothetical preference relations.

Where do we go from here? I'd like to find a justification the claim that we should save the greater number that doesn't assume that the sceptic is wrong about everything. In my view, the most interesting sceptical argument concerns aggregation. Taurek warns us against making a kind of metaphysical mistake, that of thinking that we might sum or add the suffering of many so that we might use that to explain why the interests of the many take precedence over the interests of the few:

> To me pain and suffering are magnitudes that cannot be added or summed across individuals. They are like physical beauty, boxing skill, or artistic talent ... If I call your attention to the great beauty in this woman's face, and having acknowledged it, you reply, 'But her beauty pales when compared to the awesome beauty we contemplate when we add together or sum the beauty found in each of a sea of ordinary faces,' I will not know what to make of this ... To me pain and suffering are magnitudes of this kind. I can compare the pain or suffering of this one person with the pain of the suffering of any of these many. But when I am asked to compare the magnitude of her pain or suffering with the magnitude of the pain or suffering obtained by adding or summing the pains of these many I am at a loss to know what it is I am to focus my attention on (2021: 313).

I don't disagree. I also don't always see how we're meant to derive the sceptical conclusion from his premises. Let's suppose we side with Taurek against some consequentialists in saying that the state of affairs in which the Betas are saved cannot be just better or better simpliciter than the one in which one Alpha is saved. What moral premise are we supposed to add to this metaphysical one to take us to the sceptic's conclusion? Parfit (1978: 291) imagined it was something to the effect that, special obligations aside, the only moral reason there could be to choose the Betas over the Alpha is that one outcome is better (period, simpliciter) than the other. I think it would be strange for the sceptic to appeal to this premise. It seems to assume, special obligations aside, that our moral explanations should be barely distinguishable from one a consequentialist would give.

It's one thing to say that the warning about the metaphysical mistake should serve as a warning against giving a consequentialist explanation as to why we should save the greater number and another to say that it establishes the sceptic's conclusion. I can't see how to connect the premise to the conclusion without attributing to the sceptic the view that our justification for preferring some rescue mission to another must ultimately appeal to claims about what's good or better simpliciter that, in turn, assume some problematic kind of aggregation. I shall sketch an explanation that doesn't assume the kind of aggregation that Taurek is rightly sceptical of. The starting point is the idea that we should look out for each other the best we can. If we choose the principles that govern rescue cases in a certain way, we'll get the conclusion I'm after without committing the metaphysical mistake Taurek warns us against.

Allan is heading to camp Alpha. His parents are trying to decide what to pack. They could pack the stuff for spider bites or the stuff for snake bites. We might imagine that the bite of a snake or spider would be fatal if untreated, that there's room enough only for one antidote, and that the possibility of such bites doesn't give Allan's parents decisive reason to keep him away from camp. Suppose everyone knows that it's much more likely that Allan would be bit by a spider than a snake. Suppose Allan discovers that his parents packed the stuff for snakes. Imagine him reflecting on this later in life. He might think the following things:

a. It was fine for them to do this so long as they first flipped a coin to decide what to pack;

b. It was fine for them to do this so long as they first rolled a die with the right number of sides to decide what to pack;

c. It was fine for them to do this without first using some prop;

d. They should have packed the stuff for spiders.

I can't imagine Allan thinking anything but (d). When it comes to his expectations about how his benefactors look out for his interests, he reasonably expects them to prefer the means more likely to be useful. I sort of doubt that the risks in question constituted harms, but even if we think harming is a way to do something wrong, the same holds for risking. The case is a reminder that the morality of rescue depends, in part, upon considerations about the potential outcomes but also upon the process by which we reach decisions about how to respond in light of these potential outcomes. Allan's complaints might sometimes be about harms, but they might also be about the ways potential benefactors seek to spare him of those harms.

Let's consider two more scenarios. Allan and his nine siblings are sent away to Camp Kappa. It's much more likely that Allan and his siblings would encounter spiders than snakes. Upon discovering that their parents had packed the stuff for snakes, each of the kids would think (d). Each would have reasonably expected better from their parents if the parents send the stuff for snakes. This should give you an idea of the kinds of principles that we might use to flesh out what we should do when looking out for others if we care about doing the best we can. This all seems to be in keeping with, say, **No Daredevils**.

The kids head off for another adventure. This time each is bitten by something or other. What they know is that eight have been bitten by spiders, two by snakes, and that it's only possible to get them the stuff for spiders or the stuff for snakes. Nobody knows what bit them. Here's my hypothesis. Since each kid has a credence of .8 that they need something for a spider bite, each will think (d). It's what they'll do, they think, if they survive long enough and have kids of their own. Parents who can anticipate this should, in my view, send the stuff for spiders. A good way to think about ways to benefit beneficiaries is thinking about what the beneficiaries want, after all.

We reach the conclusion that we should save the greater number here by taking account of what each kid wants. Their wants reflect an understanding of what it takes to use the best means for trying to serve our interests.[19] Since none of them prefer the responses that they do because they see

[19] There is one and only one rule, that we might choose *ex ante*, that ranks options in terms of the expected lives and that's the rule that says that we should save the greater number. It's an appealing rule, from an *ex ante* perspective, because it's the rule that maximises each individual's chance of rescue if those chances are determined *ex ante*. The sceptic, of course, has to deny that we equate expected desirability with expected number lives. This was obvious from the start given their insistence that there's not more reasons to save two on Beta than one on Alpha. To formulate a view that deviates from the one defended here, the sceptic has to say that expected desirability (or its surrogate) is sensitive to something besides the expected number of lives and so I think that we would each reject this principle, *ex ante*, since there isn't some salient competing consideration that we would be happy to say should guide decision makers that they should trade off against lives. Now, to

saving the greater number as better period or best simpliciter, it seems we've found a justification for the view that we should save the greater number that doesn't presuppose the metaphysical mistake Taurek warns us against.

It's at just this point that I expect someone to object.[20] The cases I've discussed differ from some of the cases that matter in the debates with the sceptic in an important way. The sceptic asks us to consider a different case, a case where the personal reasons are a bit more personal. In the case I've described, we are considering the reasons of potential beneficiaries who don't yet know what their personal predicament is but know what it might be. In this way, I'm imagining that each can speak for each person by thinking about how, owing to the vagaries of fate, they might find themselves in the different groups of people who need assistance and deciding how we ought to respond without taking account of their wants, requests, demands, pleas, etc. once they see how things turn out for them in particular. What if Allan learns that he's been bitten by a snake? Before heading off to camp, Allan might have told his parents that if nobody knows what bit him or what will bite him, they should send the stuff for spiders, but now that he knows that it's a snake, he's going to ask his parents to send the stuff for snakes. Doesn't this change things? Once each kid knows what bit them, we won't have unanimity.

I could try to carve out a partial victory and say that we should save the greater number when we have limited information, but I'd like to aim for something more ambitious. I think we need to decide whether to reject or suspend the principles that would be chosen *ex ante* in light of these kinds of *ex post* complaints or demands. Let's contrast two ways of approaching this issue. On the first, we recognise that people like Allan will prefer the sorts of principles that I like *ex ante* but then ask that they be waived or suspended *ex post* if he sees their application won't serve his interests and we ignore what people like Allan have to say when it doesn't agree with what they'd agree to *ex ante*. On the second, we should listen to what people like Allan have to say about the application of principles we'd choose *ex ante* when they see how their application impacts them.

It's hard for me to think of any non-circular justification for preferring the first way of approaching the issue.[21] I think that everyone (Allan included) would say that we should ignore anyone (Allan included) who tries to get us to suspend or reject the application of the principles chosen *ex ante* because they can see that every conceivable alternative does worse in terms of what's expectably best for them and for others. However, we also know that everyone might also be disposed to say that such principles should be suspended or rejected when they find themselves in Allan's predicament and seeing that the principle that looked best *ex ante* looks unappealing in light of further information. I won't try to give this non-circular justification but conclude with the observation that any view that tells us to suspend or reject the principles we choose *ex ante* because of complaints like the ones voiced by Allan suffers from just the kinds of problems we've been discussing.

When Allan asks his parents to send the stuff for snakes and to ignore the requests to send the stuff for spiders, he puts his parents in a bind. What should they do? They can't just count the number of kids bitten by spiders and snakes and use that to decide. As they think through their position on

---

be fair, I haven't really argued that we should save the greater number or identify expected desirability (in the relevant range of cases) in terms of expected number of lives. I haven't, for example, said a single thing about risk aversion. I note that now to note that debates about the rationality of risk aversion have no bearing on debates about whether to save the greater number.

[20] Thanks to [] for pressing a version of this objection in conversation.

[21] For reasons of space, I cannot address the fairness objections that someone like Timmerman (2004) might raise. I would note, however, that I don't see how it could be unfair to any person to choose and stick to principles that each would choose *ex ante*. I also think that complaints about not giving the snake bitten victims a chance ring hollow since we took account of such chances ahead of time and the distribution of bites seems to be a lottery-like device that we could use to decide which group to save.

choosing principles *ex ante* and sticking to them or waiving them as Allan hopes they'll do, they might wonder what it would take to make Allan happy the next time or happy in different circumstances. If Allan alone had been bitten by a snake, he'd complain that they were sending stuff for spiders. What if Allan were uncertain and only slightly more confident that he had been bitten by a snake? Presumably, he'd still complain if they sent the stuff for spider bites. Even if, say, it were a choice between a response that gave him a .51 chance of survival while giving his nine siblings no chance of survival, that's what he'd ask for. And what if the stuff for snake bites wasn't nearly as effective as the stuff for spiders, so we're looking at a choice between, say, 1/64 chance of saving Allan and the near certainty of saving his nine siblings? Knowing Allan, we might expect that he'd still want his parents to send the stuff for snakes rather than spiders. He'd want them to violate **No Heroes**. But **No Heroes** seemed so eminently plausible before. As much as we like Allan, maybe we should keep **No Heroes** and let him go.

As I see it, once we start waiving, suspending, or rejecting the principles chosen *ex ante* because of people like Allan, we might end up with a sceptical view, but we'll also end up with a view about the morality of rescue that cannot do justice to our answers to the easy questions. We'll be stuck with a view that cannot explain why we should dedicate our time, our efforts, and our limited resources to the rescue missions that are likely to succeed rather than those nearly certain to fail.

References

Anscombe, G.E.M. (1967). Who is Wronged? *Oxford Review*, *5*, 16–17.

Doggett, Tyler. (2013). Saving the Few. *Noûs*, *47*(2), 302–315.

Fischer, Jessica J. T. (2021). Counting People and Making People Count. *Philosophy*, *96*(2), 229–252.

Gustafsson, Johan E. (forthcoming). *Money-Pump Arguments*. Cambridge University Press.

Hare, Caspar. (2013). *The Limits of Kindness*. Oxford University Press.

Kaivanto, Kim. (2017). Ensemble prospectism. *Theory and Decision*, *83*(4), 535–546.

Kamm, F. M. (2005). Aggregation and Two Moral Methods. *Utilitas*, *17*(1), 1–23. https://doi.org/10.1017/s0953820804001372

Mahtani, Anna. (2019). Vagueness and Imprecise Credence. In Richard Dietz (Ed.), *Vagueness and Rationality in Language Use and Cognition* (7–30). https://doi.org/10.1007/978-3-030-15931-3_2

Muñoz, Daniel. (forthcoming). The Many, the Few, and the Nature of Value. *Ergo: An Open Access Journal of Philosophy*.

Muñoz, Daniel. (n.d.). The Many, the Few, and the Nature of Value. *Ergo: An Open Access Journal of Philosophy*.

Munoz-Darde, Veronique. (2005). The distribution of numbers and the comprehensiveness of reasons. *Proceedings of the Aristotelian Society*, *105*(2), 207–233.

Parfit, Derek. (1978). Innumerate Ethics. *Philosophy & Public Affairs*, *7*(4), 285–301.

Peterson, Martin. (2015). Prospectism and the weak money pump argument. *Theory and Decision*, *78*(3), 451–456. https://doi.org/10.1007/s11238-014-9435-2

Reibetanz Moreeau, Sophia. (1998). Contractualism and aggregation. *Ethics*, *108*(2), 296–311.

Resnik, Michael. (1987). *Choices: An Introduction to Decision Theory*. University of Minnesota Press.

Savage, Leonard J. (1954). *The Foundations of Statistics*. Wiley.

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

Schoenfield, Miriam. (2014). Decision Making in the Face of Parity. *Philosophical Perspectives*, *28*(1), 263–277. https://doi.org/10.1111/phpe.12044

Sen, A. K. (2017). *Collective Choice and Social Welfare: An Expanded Edition*. Harvard University Press.

Taurek, John. (1977). Should the Numbers Count? *Philosophy and Public Affairs*, *6*(4), 293–316.

Taurek, John. (2021). Reply to Parfit's 'Innumerate Ethics'. In Jeff McMahan, Tim Campbell, James Goodrich, & Ketan Ramakrishnan (Eds.), *Principles and Persons: The Legacy of Derek Parfit* (311–322). Oxford University Press.

Tenenbaum, Sergio. (n.d.). *Can't Kant Count? Innumerate Views on Saving the Many over Saving the Few*.

Timmermann, Jens. (2004). The individualist lottery: How people count, but not their numbers. *Analysis*, *64*(2), 106–112.