Playing Dice with Morality: Weighted Lotteries and the Number Problem

MATHIEU DOUCET
University of Waterloo

In this article I criticize the non-consequentialist Weighted Lottery (WL) solution to the choice between saving a smaller or a larger group of people. WL aims to avoid what non-consequentialists see as consequentialism's unfair aggregation by giving equal consideration to each individual's claim to be rescued. In so doing, I argue, WL runs into another common objection to consequentialism: it is excessively demanding. WL links the right action with the outcome of a fairly weighted lottery, which means that an agent can only act rightly if s/he has actually *run* the lottery. In many actual cases, this involves epistemic demands that can be almost impossible to meet. I argue that plausible moral principles cannot make such extreme epistemic demands.

Should the fact that there are more people in one group than in another have any bearing at all on our decision when we are choosing between them in distributing some benefit? Other things being equal, should we prefer 'the many' to 'the few'? Arguments about the 'Number Problem' have tended to revolve around cases such as the following: suppose you are in a boat equidistant from two small islands. There is one person (A) stranded on the first island, and two people (B & C) stranded on the second. All three are strangers to you. A storm is on the way that will kill those stranded on both islands, and you only have enough time and fuel to reach one of the islands and save its inhabitant(s). What should you do? Call this the *rescue case*.

For most consequentialists, of course, the Number Problem (and so the rescue case) is no problem at all: *ceteris paribus*, you should save B & C. To some, this conclusion might seem so obvious that it requires no argument. Surely two people dying is twice as bad as one person dying? In fact, the issue is not that simple. As John Taurek points out, there is no one *for whom* it is twice as bad to allow B & C to die (by saving A) than it is to allow only A to die (by saving B & C).[1]

Thinking that it *is* twice as bad to allow two people rather than one person to die requires aggregating claims in ways that non-consequentialists reject as unfair, since it ignores the interests of the minority and violates the separateness of persons. Hence the

[1] John Taurek, 'Should the Numbers Count?', *Philosophy and Public Affairs* 6 (1977), pp. 293–316, at 307.

© Cambridge University Press 2013
Utilitas Vol. 25, No. 2, June 2013
doi:10.1017/S0953820812000374

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

philosophical problem, for non-consequentialists: how can one take account of the relative number of people without aggregating the claims of individuals? This is a particular problem for non-consequentialist moral theory because common-sense morality seems to suggest that we should not be indifferent to the choice between 100 lives and a single life. If this is right, a convincing non-consequentialist theory needs to be able to explain it.

While the debate among non-consequentialists has yielded many different proposed solutions to the Number Problem, these can be classified into two broad categories:

1. Non-consequentialist justifications for saving the many.
2. Decision procedures that should be used to choose between the many and the few.

In this article, I evaluate one prominent decision procedure solution to the Number Problem: the weighted lottery (WL), which aims to take the interests of *each* person into account in deciding what to do. I will argue, however, that WL fails as a possible solution in many instances of the Number Problem. The objection can be stated quite simply: WL is useless in the absence of complete information about the numbers involved in a choice between the many and the few. Since ignorance of just this sort is quite possible in many specific instances of the Number Problem, WL is excessively demanding and can be, at best, a partial solution. Its failure, I suggest, stems from a misunderstanding of the 'problem' that the Number Problem represents.

I. THE OPTIONS

The simplest solution in the rescue case is to save the many, a solution many versions of consequentialism can easily explain. Agent-neutral maximizing act-utilitarianism (or 'pure' act-utilitarianism), which requires of all actions that they maximize overall utility, will clearly (absent any significant information about the identities of A, B and C) prefer saving the many, and most other versions of consequentialism can easily offer arguments to the same conclusion. Rule-consequentialists, for example, who assess the relative consequences of systems of rules, as opposed to individual acts, could easily argue that the adoption of a general rule preferring saving larger to smaller groups would be have better overall consequences.[2] Other versions of consequentialism modified in response to worries about pure

[2] See, for example, Brad Hooker, *Ideal Code, Real World: A Rule Consequentialist Theory of Morality* (Oxford, 2000).

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

act-utilitarianism's excessive motivational demands can also embrace a policy of straightforwardly saving the many in the rescue case. Those who reject pure act-utilitarianism's maximizing demands in favour of satisficing do so partly so as to deal with the potentially excessive demands of maximizing, but saving the many in the rescue case is no more motivationally or practically demanding than saving the few.[3] And those who reject pure act-utilitarianism's agent-neutrality can also easily embrace saving the many, given that saving the many is no more demanding and that all of the islanders are, by hypothesis, strangers to the rescuer, who therefore has no agent-relative interest in either group.[4] So while there will be many cases on which these competing consequentialist moral theories disagree, most can agree that saving the many is the morally correct choice in the rescue case.

If saving the many *is* the morally correct choice, non-consequentialists need to explain why without resorting to aggregation. Two prominent candidate explanations are the balancing solution and the egalitarian solution. According to the balancing solution, equal and competing claims are balanced against one another, and so the decision to save the larger group is determined by the unbalanced claims on that side.[5] According to the egalitarian solution, treating people fairly requires distributing goods between them equally, *ceteris paribus*. If the good is indivisible, as it is in the rescue case, this is not possible, and so one should distribute goods in the way that most closely approximates an equal division. Giving the good to the larger group accomplishes this.[6] On both solutions, the larger group are saved, but defenders of both argue that this can be accomplished without aggregating the claims of individual members of groups.

Not all non-consequentialists, however, are willing to agree that morality requires automatically rescuing the many.[7] They argue that saving the many gives the few no chance at all of being rescued, and

[3] See, for example, Jason Rogers, 'In Defense of a Version of Satisficing Consequentialism', *Utilitas* 22 (2010), pp. 198–221. The issue of demandingness will return in section V, below.

[4] For example, Douglas Portmore's act-consequentialism includes both agent-relative constraints and agent-relative options. See e.g. *Commonsense Consequentialism: Wherein Morality Meets Rationality* (New York, 2011).

[5] This solution was first proposed in by Frances Kamm in *Morality, Mortality*, vol. 1 (Oxford, 1993), pp. 129–41. It was given a contractualist interpretation in T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass., 1998), pp. 114–21, and refined in Rahul Kumar, 'Contractualism on Saving the Many', *Analysis* 61 (2001), pp. 165–70.

[6] Nien-hê Hsieh, Alan Strudler and David Wasserman, 'The Numbers Problem', *Philosophy and Public Affairs* 34 (2006), pp. 352–72.

[7] One subtext to this debate is the extent to which non-consequentialist moral theory should accommodate the judgments of what Kumar calls 'common-sense morality' ('Contractualism on Saving the Many', p. 165).

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

that this is unfair to minorities.[8] Minority claims may be included in the aggregate, but they are systematically outweighed, and this can strike the non-consequentialist as an unfair result of aggregation. In fact, it is in part the perceived unfairness of the potential tyranny of the majority that motivates some to reject consequentialism in the first place. Some of those who reject aggregation also argue that the balancing solution, in particular, seems to involve the kind of aggregation that should be ruled out by its own lights.[9] Instead, such non-consequentialists argue that those faced with a choice between the many and the few should use a decision procedure, since this is the only way adequately to recognize each person's equally strong claim to be saved. Automatically saving the many, they argue, is inconsistent with such equal recognition, and is therefore unfair.

Taurek's solution is to flip a coin, which would have the advantage of giving each person with a claim to be saved an equal chance at having that claim satisfied (50%).[10] This proposal, however, *also* fails to take into account each person's claim to be saved. Flipping a coin means treating the choice between A and B in the same way as the choice between A and B & C, which means that the addition of C makes no difference at all. C could complain that his presence was being ignored – things would be exactly the same if he were not there, which does not seem consistent with taking his claim into account.[11] We appear to have gone from not allowing A's presence to matter to not allowing C's presence to matter. If what matters is that each person's claim be taken into account, then both automatically saving the many *and* flipping a coin fail to do so.

More recently, some proponents of a decision procedure solution to the number problem have argued that the right thing to do is not to flip a coin, but to run a weighted lottery.[12] This proposal walks a middle path between the two existing options – it is a way of taking

[8] As Jens Timmerman puts it, 'it seems unfair simply to leave the minority without any chance to be saved' ('The Individualist Lottery: How People Count, but Not Their Numbers', *Analysis* 64 (2004), pp. 106–12, at 106).

[9] This objection is raised in Michael Otsuka, 'Scanlon and the Claims of the Many versus the One', *Analysis* 60 (2001), pp. 288–93; and in David Wasserman and Alan Strudler, 'Can a Non-consequentialist Count Lives?', *Philosophy and Public Affairs* 31 (2003), pp. 71–94, though Wasserman and Strudler are also in favour of automatically saving the many.

[10] John Broome agrees with Taurek that fairness requires tossing a coin, although he argues that the rescuer nevertheless ought to save the many, since in this case fairness is outweighed by the value of saving extra lives. 'Kamm on Fairness', *Philosophy and Phenomenological Research* 58 (1998), pp. 955–61, at 957.

[11] This objection is raised by Scanlon, *What We Owe to Each Other*, p. 232, and Iwao Hirose, 'Weighted Lotteries in Life and Death Cases', *Ratio* 20 (2007) pp. 45–56, at 47.

[12] Kamm calls this procedure 'proportional chances' (*Morality, Mortality*, pp. 128–42). She favours it in some cases, though not in those like the rescue case. Jens Timmermann

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

A's claim into account (unlike simply saving the many), while allowing each of B & C's claims to make a difference (unlike flipping a coin). The proposal works as follows: each person is assigned a number in the lottery. We then draw a single number, and rescue the person whose number we drew, along with everyone else on his/her island. In effect, this means that members of the larger group are more likely to be rescued than members of the smaller group: A has a 1/3 chance of being rescued and B & C both have a 2/3 chance of being rescued. Larger groups have a greater chance of being saved than smaller groups, and the greater the difference in relative size, the greater the chance that the larger groups will be saved. But this does not involve illegitimate consequentialist aggregating, or a diminishing of the strength of the claims of the few, since the decision is based entirely on the strength of the claim of the individual who wins the lottery, and each individual's claim is given equal weight. As Timmerman puts it, 'groups are *de facto* given a greater chance of survival, but not *de jure*, which would be illegitimate'.[13]

II. RUNNING THE LOTTERY

One consequence of WL is that the only way for a would-be rescuer to do the right thing is for her to *actually run the lottery*. As Timmerman puts it, the desire to treat each claim equally leads to a stalemate that can only be resolved by a fair method to decide: 'so grave a decision should not depend on a whim of the agent'.[14] In the rescue case, the right thing to do is to rescue the person selected by the lottery, along with everyone else on his/her island. This means that there is no way to do the right act without running the lottery, since the right act is *determined* by the lottery. Any choice the rescuer makes without first running the lottery would be, necessarily, unfair. For John Broome, who favours lotteries in some cases, this unfairness is not decisive: though he admits that not running a lottery is unfair, he believes that the demands of fairness 'should be weighed against the demands of overall satisfaction' and that in cases like *rescue* it is best (though still unfair) to save the many.[15] Such a view incorporates the fairness of a lottery as one element of an overall solution to the Number Problem. For a defender of WL such as Timmerman, however, the fairness of the lottery is not one consideration to be weighed against the demands of

('The Individualist Lottery') calls it the "individualist lottery'," and defends its use in the rescue case.

[13] Timmermann, 'The Individualist Lottery', p. 111.

[14] Timmerman, 'The Individualist Lottery', p. 108.

[15] John Broome, 'Fairness', *Proceedings of the Aristotelian Society* 91 (1990), pp. 87–101, at 96.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

overall satisfaction in arriving at a solution: WL itself is the solution, and so proceeding without running the lottery would be wrong. Such a choice would be better than saving *no-one*, but it would nevertheless represent a violation of the obligation to be fair. As such it could not count as the right action, even though there are degrees of failure to live up to an obligation.

By itself, this is not an objection to WL, but it is important to note that it *is* a consequence of the view. Without running the lottery, there is no way of doing the right thing. When faced with a choice between 99 and 1, we cannot simply choose the 99 because the procedure is overwhelmingly *likely* to generate that result: saving the 99 without first running the lottery would be unfair. And if (as will happen 1 per cent of the time), the lottery directs us to save the 1, we should not baulk on the grounds that such a course of action is extremely unlikely (nor, of course, on the grounds that allowing 99 to die is much, much worse than allowing 1 to die, or that it is unfair).

Saving the many is the act that WL is likely to generate (and the likelihood increases the greater the difference between the many and the few), but in order to be fair, the decision must have been brought about *by* the lottery, and not simply match the decision that probably would have been brought about in that way. Nor can we evaluate the person's actions third-personally by saying that she did what the lottery would have probably told her to do, or even by running the lottery for her without her knowledge and praising her if her action happens to match the results of our lottery. There is no way, in other words, for a would-be rescuer to do the right thing by acting *as if* she had run the lottery, and no way for us to assess her actions as right *unless* she runs the lottery. Someone faced with the rescue case, or any other instance of the Number Problem, must have at her disposal some means of running a fairly weighted lottery – a large collection of different-sided dice, an adjustable wheel of fortune, a watch with a second hand,[16] a computer program specially designed for the purpose, or, at the very least, a hat, a pen, and pieces of paper. Without such tools, she simply cannot do what morality requires. She still has an obligation to do *something*, of course: she should still save someone. But if WL is the correct solution to the Number Problem, then she will be acting unfairly whomever she saves and however she decides whom to save. Again, this is not necessarily an objection to WL – some version of 'be prepared' may be a reasonable moral requirement – but it does emphasize the importance of actually being able to run the lottery. As we shall see, this requirement is potentially fatal to WL.

[16] Thanks to an anonymous referee for *Utilitas* for suggesting this method.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

III. LOTTERIES AND INFORMATION

At the origins of the debate on the Number Problem is Taurek's quite general question: other things being equal, is there any reason to prefer the many to the few? One of his examples is a Coast Guard captain's choice between directing his ship to the northern tip of an island, where 'a large number are gathered', and the southern tip, where 'a handful find themselves'.[17] The idea of someone on a boat having to choose to save one of two groups has been taken up repeatedly in subsequent discussions, which are generally put in terms of quite specific numbers – as a choice between two and one, or five and one.[18] Such specific choices are theoretical simplifications that help us to model choices between larger and smaller groups: if the numbers count, then a single additional person counts, and so it should be possible to express clearly the argument in terms of a choice between two and one. But we should note that Taurek himself speaks in vague terms: he describes a choice between 'a large number' and a 'handful'.

There is nothing wrong with such theoretical simplifications, of course. They are intended to abstract away morally irrelevant differences in order to concentrate on the precise issue at hand: whether (and if so, how) the numbers count. Once we have a grasp of the basic structure of justification in the simple case, we can extend it to more complicated cases. So, for example, the non-consequentialist arguments for saving the many are often put in terms of saving two over saving one, but the basic structure of the argument extends to any comparison between 'the many' and 'the few': so long as we know that there are more people in one group than in the other group, we know what is morally required. Arriving at a solution to the simple case shows us how to proceed in the complicated cases, and so provides an answer to the quite general question of how to treat choices between the many and the few. Provided we accept any of these arguments, we know which course of action is right once we know which of the groups is larger. We do not need to know *how much* larger the larger group is, or to know how many people are in each group.

The same, however, cannot be said for WL. Since it is necessary actually to run the lottery in order to take each person's claim to being

[17] Taurek, 'Should the Numbers Count?', p. 310.

[18] Boats reappear in Kumar, 'Contractualism on Saving the Many'; Wasserman and Strudler, 'Can a Non-consequentialist Count Lives?'; Hsieh et al., 'The Numbers Problem'; Timmerman, 'The Individualist Lottery'; Hirose, 'Weighted Lotteries in Life and Death Cases'; Matthew Liao, 'Who is Afraid of Numbers?', *Utilitas* 20 (2007), pp. 447–61; Thom Brooks, 'Saving the Greatest Number', *Logique et Analyse* 45 (2002), pp. 55–9; and Martin Peterson, 'The Mixed Solution to the Number Problem', *The Journal of Moral Philosophy* 6 (2009), pp. 166–77.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

saved into account in deciding what to do, specifying the numbers is not merely a theoretical simplification: it is a necessary requirement of the application of the theory. The lottery is not an idealization that helps us to grasp the structure of the solution: the lottery *is* the solution. This is because the way that the WL takes each person's claim to be saved into account is by giving each person a number in the lottery (or side of the coin, if there are only two people; or side of the die, if there is a die with either as many sides as there are people, or a number of sides equal to a multiple of the number of people). But in order to take into account each person's equal claim to be saved, each person must be assigned his or her own number. It is therefore crucial that we know the precise number of people with a claim to be saved. It is not enough to know that one group is larger than another, or even how much larger it is. If we are to decide in the way required by WL, we need to know the precise number of people with a claim to be rescued.

The fact that a would-be rescuer needs to know the exact numbers involved is, at the very least, a serious practical problem for WL. That is because, in many actual instances of the Number Problem, the person faced with the choice between the many and the few will not have the luxury of complete information about the exact numbers involved. Instead, he or she will have only a vague sense that one of the groups is likely to be larger than the other. This is true of Taurek's Coast Guard captain, who is faced with a choice between 'a large number' and 'a handful'. It is also true of pilots of planes likely to crash into more or less crowded areas,[19] military commanders who can send reinforcements to the rescue of larger or smaller groups of soldiers,[20] paramedics who can rush to the scene of a multiple car accident or a single car accident, fire-fighters who have to choose between fighting a fire in an apartment complex or a single-family home, and aid workers with food and medicine to distribute to one of two remote communities. Each of these are plausible, real-world examples of people faced with a choice between saving what they have reason to believe is a larger group of people and a smaller group of people (or, in the case of the pilot, between killing a smaller number of people and a larger number of people). In none of these examples is it reasonable to suppose that those who have to make the decisions have anything other than a rough estimate of the number of people involved, even if they have a justified belief that one of the groups is larger. The pilot has only the sense that one area is more heavily populated than another, the military

[19] An example from Philippa Foot, 'The Problem of Abortion and the Doctrine of Double Effect', *Virtues and Vices and Other Essays in Moral Philosophy* (Oxford, 1977), pp. 19–32, at 23.

[20] A similar example appears in Hsieh et al., 'The Numbers Problem', p. 357.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

commander might have no idea how many soldiers in each group have already been killed, the paramedic has no idea how many people were originally in each car and how many of those are already dead, and the same goes for the fire-fighter. The aid worker might know only that one of the communities was larger than the other the last time reliable information was gathered.

Such lack of precise information is no problem for those who, whether on consequentialist grounds or not, argue for saving the many. On such accounts, an agent can proceed once she knows (or has good reason to believe) that one group is larger than the other. Nor is such ignorance a problem for Taurek's coin-flip solution, which does not even require knowing which group is larger. Lack of precise information is, however, a problem that can often be, for practical purposes, insoluble for the proponent of WL. If the only fair way to proceed in the rescue case is by doing what the lottery tells one to do, then there is no way of proceeding fairly without knowing precisely how many people have a claim on one's aid. A would-be rescuer might have at her disposal a complete set of dice, or even a computer program specifically designed to run weighted lotteries in such situations, but all of that equipment will be useless to her if she does not know which die to select or which values to plug into the program. No one faced with a choice between the few and the many who lacks complete knowledge of the numbers involved will be able to do what morality requires, and so someone who, through no fault of her own, does not have that information has no way of discharging what the lottery solution insists is her moral obligation. This sort of blameless ignorance is certainly possible, and it is perhaps even common among those who, like our Coast Guard captain, pilot, soldier, paramedic, fire-fighter and aid worker, are faced with choices between the many and the few.

At this point, a defender of WL might object that the argument advanced above uncharitably saddles WL with an unreasonable information requirement. Precise information about the exact number of people in each group is not necessary, the defender could claim, since all that is required is the *ratio* of the two groups, and it is possible to know the ratio without knowing the exact number of people. An agent equipped with knowledge of the ratio can therefore run a fairly weighted lottery even without knowledge of the precise number of people in each group.[21]

There are two points to make in reply to this objection. First, such a move would represent a significant weakening of WL. For those concerned about the unfairness of aggregating, the point of a decision

[21] Thanks to two anonymous referees for *Utilitas* for urging me to consider this objection.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

procedure is to ensure that the Number Problem is solved in a way that respects the rights of each individual and treats each person's claim to be saved equally. Timmerman's view is that, in rejecting aggregation, 'it must be our duty first and foremost to save an *individual* person in need'.[22] On this view, a fair lottery must be an *individualist* lottery. Giving every individual a number in the lottery is not merely a way of ensuring that the probabilities line up properly; rather, it is the only way genuinely to give each individual an equal claim to be saved.

The procedure therefore has two steps. First, select the single individual whose claim is going to be satisfied. If each person gets a number in the lottery, then each person has an equal chance of having his or her claim satisfied: for any group of N, everyone has a 1/N chance of winning the lottery. Second, having selected an individual to be saved, upon reaching that person's island the rescuer incurs an *additional* obligation to save the rest of the people on the island, since this can be accomplished without loss.[23] As Timmerman puts it, '[e]veryone has the same chance of winning the lottery. In addition, some people stand a good chance of benefitting from someone else's good luck.'[24]

This two-stage procedure has the *effect* of making the outcomes of the individualist lottery equivalent to a lottery run on ratios, but for Timmerman this is not enough – this is why he distinguishes between giving the larger group a greater chance of being saved *de facto* as opposed to *de jure*. Only a *de facto* greater chance is legitimate, and this is what an individualist lottery does. A ratio lottery, however, gives larger groups a *de jure* greater chance, since it does not have two distinct stages, but directly gives the larger group a greater chance at the outset. While the outcomes of both procedures might be the same, what matters for certain committed non-consequentialists is not merely that the outcomes match a certain probabilistic distribution, but that those outcomes are brought about in a certain way – by treating the claims of each individual equally. For Timmerman, at least, this requires giving each person a unique number in the lottery, and so requires knowledge of the precise number of people involved.

Not all defenders of WL, however, need to be as dedicated to this individualist approach. Some may be satisfied with a ratio-driven lottery, since it represents the best approximation of the aim of giving everyone an equal chance to be saved. Nevertheless, the move from a lottery that requires knowledge of precise numbers to a lottery that

[22] Timmerman, 'The Individualist Lottery', p. 110.

[23] This two-stage process – first, select an individual to be saved, second, save everyone else on his/her island – also appears in Kumar's 'Contractualism on Saving the Many.'

[24] Timmerman, 'The Individualist Lottery', p. 111.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

requires only knowledge of the ratio of the groups does not significantly improve the ability of agents with incomplete information to run a fairly weighted lottery.

It is true that there will be more cases in which the agent knows the ratio of the two groups than cases in which the agent knows the precise numbers in each group, since one can always determine the ratio if one knows the precise numbers but not vice versa. Perhaps a would-be rescuer can know that all the men have gathered on the south side of the island, all the women have gathered on the north side, and that there are approximately as many men as women on the island. In such a case, the rescuer can know that the size of the two groups is close to equal, and can perhaps flip a coin to choose between them without knowing how many people are in each group.

There will therefore be a number of cases that WL can handle even in the absence of complete information about the number in each group. Such cases are likely to be relatively rare, however. In none of the examples discussed above – the Coast Guard captain, pilot, soldier, paramedic, fire fighter, or aid worker – does the rescuer have even approximate knowledge of the ratio of the two groups. In some of them, the rescuer does not even know which group is larger. In none of these cases, then, is a ratio-driven lottery any better than an individualist lottery.

After all, in general the best way to establish the ratio of two groups is by knowing the precise number in each. Ignorance of the numbers is very often ignorance of the ratio as well. The only cases in which one can be confident that one has knowledge of the ratio in the absence of knowledge of the precise numbers are cases in which one has some *independent* reason to be confident that a given ratio holds; for example, knowledge that the men and women are separated into two groups, and that there are approximately equal numbers of each. But this imposes an additional information requirement on the ratio version of WL – in the absence of knowledge of the precise numbers involved, one can only run the lottery on the basis of the ratio if one has independent evidence for that ratio. In the many cases in which this information is as inaccessible as information about the precise number of people in each group, the move from an individualist lottery to a ratio-driven lottery will be no help.[25] So while allowing for a ratio-driven lottery might make it possible to arrive at a decision in some instances of the number problem, these cases are so rare that the information requirement still represents a serious barrier to running a fairly weighted lottery. And since running the lottery is necessary, the information requirement,

[25] Thanks to an anonymous reviewer for *Utilitas* for suggesting this line of argument.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

even in ratio-driven lotteries, makes it effectively impossible for many moral agents to act as WL requires.

IV. CHOICE IN DYNAMIC CONTEXTS

The argument just presented assumes an agent faced with a choice between the many and the few in which each group had a set membership, but the agent has no way of knowing the precise size or ratio of each group. The possibility of such ignorance presents a practical problem for defenders of WL, but it is not the only kind of ignorance that can confront an agent faced with such a choice. That is because not all such choices are made in static contexts with unchanging numbers of people. Again, it may be helpful to think of them in this way for the purposes of theoretical clarity – hence the popularity of islands and boats – but we can also face such choices between larger and smaller groups in dynamic contexts. In these cases, an agent must choose between directing aid or resources to one of two groups where the exact make-up of the two groups will be partly determined by the very choice that the agent makes. In dynamic contexts running a WL is impossible, not simply through lack of knowledge of the exact number of people in each of the relevant groups, but because the precise size of the relevant groups is partially *constituted* by the agent's choice, and so cannot serve as part of the *grounds* for that choice.

For example, imagine an aid worker in a position to dig a well in one of two remote villages. The residents of the village that gets the well will fare significantly better than the residents of the village that does not – their water will be safer to drink and they will have to spend much less time getting it, so they will be healthier, more prosperous, happier, less dependent on outside aid, and so on. Village A currently has more people than village B. But knowing the precise number of people currently in each village might not be enough properly to take into account all of those likely to be affected by the aid worker's decision, and so all those who might have a legitimate claim to her assistance. This is because where she decides to dig the well will have a significant influence on, not just the current residents of the villages, but also the future population size of the two villages. The village that gets the well will be healthier than the one that does not, and so will enjoy more natural population growth: more babies will be born, more of them will live past infancy, fewer adults will die of preventable disease, and so on. In addition, the village that gets the well is likely to attract immigration from other villages that also have a legitimate claim to have wells dug but which, for logistical reasons, are not candidates for receiving the well. The village that does not get the well, by contrast, is likely to experience population decline, not just through preventable deaths,

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

but also through emigration, as its residents are forced to leave the village (something, we can assume for the sake of the argument, they very strongly do not want to do) in search of opportunities elsewhere.

This is a dynamic situation, since the number of people in each of the relevant groups (village A and B) will be partly determined – through births, deaths, immigration and emigration – by the choice that the aid worker makes. (An additional complication: it could be that only some of these potentially affected people have a legitimate claim on the resource in question.) This dynamic example is like the standard static examples of the Number Problem in that the good to be distributed is indivisible, and either a larger or smaller group must get it. In dynamic cases, however, running a WL is impracticable, not simply because the agent faced with the choice does not know the precise number of people in each group, or even the ratio of the groups (as might be the case in static cases), but because the number of people in each group is, in some sense, *indeterminate* prior to the agent's choice. And if this is so, then it is not possible to assign each person a number and make the choice by running a WL – the number, identity and group membership of the people who qualify for inclusion in the lottery cannot be determined in advance. If this is the case, then a WL cannot be what is required in order to choose fairly in all instances of the Number Problem.

V. TWO SENSES OF 'DEMANDING'

The problem with WL is that, as a general solution to the choice between the many and the few, it can require moral agents to possess knowledge that it is difficult or even impossible for them to acquire. This criticism raises the issue of just what sorts of demands a moral principle – and, more broadly, a moral theory – is entitled to make without counting as unreasonable.

In order to see how the demands of WL are unreasonable, it will help to distinguish two broad senses in which a moral principle can be demanding. It can be, first, *motivationally* demanding. Pure act-utilitarianism is often seen to be demanding in this sense: in requiring that all our actions aim to maximize overall welfare, it seems to reach deep into our lives, and to require impartiality, altruism and systematic self-sacrifice on a level that few are willing to undertake.[26] The idea that the correct moral theory should not demand such extreme and potentially character-undermining self-sacrifice is often referred to as the demandingness objection. While this objection was first raised with

[26] Bernard Williams frequently argues that this kind of demandingness counts against utilitarianism. See, for example, 'A Critique of Utilitarianism', in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge, 1973).

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

respect to pure act-utilitarianism, subsequent discussions have argued that a similar problem arises for other consequentialist moral theories, and even for *non*-consequentialist moral theories. Given the massive levels of global economic inequality and widespread suffering among hundreds of millions of the world's poor, it seems that any theory that requires us to save lives in emergencies is potentially susceptible to some version of the motivational demandingness objection.[27] Defenders of a variety of moral theories have therefore recently sought to respond to the demandingness objection in ways that aim to show that the theories in question are not, in fact, as demanding as it might initially appear.[28]

Motivational demands are not the only kind of demand that a moral principle can make, however – such principles can also be *practically* demanding in various ways, making it difficult to follow even for those who are suitably motivated to do so.[29] A principle can, for example, be *physically* demanding, by requiring strength, skill or endurance that few possess. A principle that required running sub-2:30 marathons might not be motivationally demanding for running enthusiasts, but the large majority of highly motivated runners are unable to run at such elite speeds. A principle can also be *logistically* demanding, by requiring, for example, the rapid and complex coordination of the actions of many far-flung strangers.

Most importantly for our purposes, a moral principle can be practically demanding by making a variety of *epistemic* demands. A principle is *informationally* demanding if acting in accordance with it requires an agent to possess a considerable degree of knowledge that he or she would be otherwise quite unlikely to have and that would be extremely difficult to acquire, such as knowledge of the exact number of people on some islands, the precise outcomes of uncertain future processes, or the exact welfare levels of various people or groups. And a principle is *cognitively* demanding if acting in accordance with it requires an agent to engage in extremely cognitively difficult tasks that are beyond the ability of most people, such as calculations of the precise welfare outcomes of a large number of alternative courses of

[27] Garrett Cullity argues that rights-based approaches are equally susceptible to such an objection. See *The Moral Demands of Affluence* (Oxford, 2004). Elizabeth Ashford has made a similar argument with respect to contractualism in 'The Demandingness of Scanlon's Contractualism', *Ethics* 113 (2003), pp. 273–304.

[28] Several such approaches are gathered in *The Problem of Moral Demandingness*, ed. Timothy Chappell (Basingstoke, 2009).

[29] Onora O'Neill distinguishes between motivational and practical demands in 'Demandingness and Rules', in Chappell, pp. 59–69, at 59. The account of practical demands she outlines is broader than the narrowly epistemic demands emphasized below. O'Neill does emphasize that such demands can be excessive even for agents who are motivated to follow them.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

action or counterfactual reasoning about the relative probabilities of each of these outcomes obtaining.

These two broad senses in which a moral principle can be demanding are quite different, in that a principle can be very demanding in either one of these senses without being at all demanding in the other. Giving away all of one's money is not at all demanding in the practical/epistemic sense: one needs only to write a cheque. But it is quite demanding in the motivational sense: few are inclined to do it, even if they suppose that morality requires it of them. WL is arguably no more motivationally demanding than other proposed solutions, all of which claim that the would-be rescuer has an obligation to help someone. But it can be *considerably* more demanding in the epistemic sense. According to all of the justifications for saving the many, a rescuer does the right thing so long as she rescues the many (provided, that is, that her decision is not wholly motivated by a desire to, for example, kill her enemy who happens to be among the few). She does not need to run through the precise justification in her head prior to making her choice. So even if *understanding* that justification is cognitively demanding – requiring, for example, the calculation of a Gini-index to measure inequality – acting as the principle demands is not demanding in this way.[30] Since WL requires actually running the lottery, and since doing so can be extremely practically demanding, WL is potentially much more demanding than the other solutions to the Number Problem.

There is some disagreement over the extent to which even extremely stringent *motivational* demands count as an objection to a moral theory. Such demands make us uncomfortable, but this might reveal a problem with us, and not with the principle. Those worried by the demandingness objection certainly take it to be a problem that any correct moral theory should overcome, but there are some for whom even very stringent motivational demands do not necessarily constitute a problem for a moral theory.[31] I take no stance on this

[30] Peterson considers, and ultimately rejects, the use of a Gini-index at 'The Mixed Solution to the Number Problem', pp. 172–3.

[31] David Sobel, for example, argues that the demandingness objection 'cannot itself provide good reason to break with Consequentialism since it must presuppose the truth of prior and independent breaks with Consequentialism' (David Sobel, 'The Impotence of the Demandingness Objection', *The Philosopher's Imprint* 7 (2007), pp. 1–17, at 1). Kant is even less persuaded: 'actions of which the world has perhaps so far given no example, and whose very practicability might be much doubted by one who bases everything on experience, are still inflexibly commanded by reason' (Immanuel Kant, *Groundwork of the Metaphysics of Morals*, trans. Mary Gregor (Cambridge, 1997), 4:408). And of course one common target of the demandingness objection, Peter Singer, is of the view that morality *does* in fact make extremely stringent demands, perhaps most famously in 'Famine, Affluence, and Morality', *Philosophy and Public Affairs* 1 (1972), pp. 229–43.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

question here. Rather, I raise the issue to highlight two points. First, a moral principle can be demanding in ways that extend beyond the sorts of motivational demands that have been the main focus in discussions of the demandingness objection. Second, while the problem of *motivational* demandingness is one that applies in much the same way to all possible solutions to the Number Problem, the problem of *epistemic* demandingness is uniquely acute for WL, since WL has much higher informational demands than other solutions.

A moral principle that is too epistemically – or, broadly speaking, too practically – demanding for an agent to implement successfully has a potentially fatal flaw, since it cannot be properly action-guiding even for those who are motivated to act in accordance with its requirements. What makes agents unable to live up to the demands of the principle is not a motivational failure that a broader moral theory can perhaps accommodate by assigning blame to such failures, but simply a practical inability either to identify what the principle demands, or to carry out those demands once identified. Such practical problems can beset agents who are morally motivated and practically wise: the best practical deliberator in the world will be unable to comply with the demands of morality if doing so requires her to possess information that she cannot obtain or engage in mathematical calculations that are well beyond her (or any non-mathematician's) ability. If a principle makes demands that an agent cannot carry out, then it (a) fails to be action-guiding and (b) blames agents for a failure to do what morality requires, even when such agents are properly motivated by moral considerations, and when their failures are brought about by a morally blameless *inability* to do what morality requires.

A plausible moral principle, then, and indeed a plausible moral theory, should make demands that are constrained by the ability of properly motivated, practically wise agents to carry them out. In addition to not demanding the logically or physically impossible, such a theory should also take into account the practical limitations of even the most virtuous agents. We are physically and cognitively limited beings, and if a moral theory is to be a theory for beings like us, then what it demands must not only be possible in a thin logical sense, but more robustly achievable by beings constituted, as we are, with such cognitive and physical limitations. If a moral theory entails that we are obliged to have either knowledge or cognitive and logistical abilities that few if any possess or indeed could possess, then the theory is excessively demanding and should be modified or abandoned.

A common objection to pure act-consequentialism is that it is excessively demanding in the epistemic sense: it seems to require agents to engage in a cognitively draining, and potentially endless, calculation of the utility upshots of alternative courses of action. This

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

task is so complex that an agent often has no way of knowing which of the actions available to her is morally required.[32] In fact, such calculations can be self-defeating: we would often do better, from the perspective of overall utility, if we did not engage in such frustrating and time-wasting calculations.

The standard consequentialist reply to this objection is to distinguish between a criterion of right and a decision procedure. Though the right is defined in terms of maximizing the good, we need not insist that right action must be the upshot of moral deliberation that proceeds by calculating expected outcomes for every possible action.[33] On this view, there are many acceptable and even praiseworthy deliberative routes to right action, and not all of them involve calculating, or even considering, the way to maximize the good.

Defenders of non-consequentialist solutions to the Number Problem that favour saving the many can employ a similar strategy, distinguishing the justification for saving the many from the kinds of deliberative processes that might make someone praiseworthy for saving the many. Saving the many might be right because of considerations of balancing, or equality, but this does not mean that an agent faced with the choice must consciously and explicitly make her choice on those grounds: perhaps saving the many strikes her as so clearly right that she engages in *no* explicit deliberation. This strategy, however, does not seem open to those defending WL. There can be no distinction between the criterion of right and the decision procedure when the criterion of right *just is* (or at the very least, requires the application of) the decision procedure. And so the practical demands of the decision procedure are a legitimate area of concern for WL in a way that they need not be for consequentialism. Because WL requires agents to be able to run potentially complex lotteries, and to have precise knowledge of the numbers of people with a legitimate claim on the distribution of some resource, it can make demands that are unreasonable and beyond the capacity of almost any moral agent, regardless of his or her level of moral motivation or degree of practical wisdom. WL can therefore fail to guide properly the actions of those

[32] James Lenman presents a particularly forceful version of this objection in 'Consequentialism and Cluelessness', *Philosophy and Public Affairs* 29 (2000), pp. 342–70.

[33] This distinction can be traced back to J. S. Mill's claim that an agent's motive 'has nothing to do with the morality of the action' even if it bears on 'the worth of the agent' (*Utilitarianism* (Indianapolis, 2001), p. 18). For more recent defences of this strategy, see R. Eugene Bales, 'Act-utilitarianism: Account of Right Making Characteristics or Decision-making Procedure?', *American Philosophical Quarterly* 8 (1971), pp. 257–65; and Peter Railton, 'Alienation, Consequentialism, and the Demands of Morality', *Philosophy and Public Affairs* 13 (1984), pp. 134–71.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

moral agents. As such, it fails as a general solution to the Number Problem.

VI. DROPPING THE LOTTERY

I have argued that WL's requirement that we actually *run* the lottery makes it excessively epistemically demanding and so it fails to be action-guiding. How might defenders of WL respond? One possibility is to modify the view by denying that it is necessary actually to run the lottery – what matters is simply that the probability of saving the many is equivalent to the proportion of the many relative to the few. In the standard rescue case, that would mean that the agent acts rightly so long as the probability of his/her saving each group corresponds to what WL prescribes – that probability does not need to be actually *generated* by a lottery.

If this is right, then perhaps the proponent of WL can appeal to the distinction between a criterion of right and a decision procedure after all, and can therefore deny that WL is unreasonably epistemically demanding. So long as the probability of saving each group corresponds to what the WL prescribes, in other words, the agent does not do anything unfair.

Modifying WL in this way would represent a significant concession. Timmerman, for example, explicitly presents WL as decision procedure intended to solve a stalemate: since everyone's claim to be rescued is equal, 'we must turn to procedure ... anyone stuck in the rescuer's dilemma requires a *fair method* to decide'.[34] WL is that method. So on his view, WL is not merely a criterion of right. Nevertheless, the practical and epistemic challenges to WL might force its defenders to make this move in order to salvage the motivating insight of WL – that each person's claim to be rescued should be given equal consideration in a way that allows each person a genuine chance to be saved. It is this commitment to fairness that leads away from simply saving the many to a decision procedure in the first place. But is such a modification possible? Can WL be recast as a criterion of right?

If WL really is a criterion of right rather than a decision procedure, then the failure to be action-guiding is no longer an objection. But if it is a criterion of right, then agents faced with the Number Problem still need to implement *some* decision procedure. Moreover, that procedure will have to make their actions accord with the demands of WL, which is to say that it will have to ensure that their decisions match the probabilities of a fairly weighted lottery.

[34] Timmerman, 'The Individualist Lottery', p. 108 (emphasis in original).

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

The original objection here is that, as a decision procedure, WL is impractical. It is hard to see how this problem can be solved by substituting another decision procedure in its place. Surely the easiest way to ensure that your decisions match the probabilities generated by a lottery is actually to run one. The alternative would seem to be to implement (perhaps unconsciously) a mental process that somehow randomizes choices in a way that matches the desired probabilities. Perhaps some people have such mental processes or dispositions, but 'choose according to a disposition that matches the probabilities generated by a fairly weighted lottery' is hardly more practical as a decision procedure than 'run a weighted lottery'. The mechanism for such a disposition is entirely mysterious, and our confidence that any particular rescuer's choice actually exhibited such a disposition would be, at best, extremely weak. In fact, in the absence of a series of choices between groups with equivalent ratios, it is hard to see what could possibly count as evidence that any particular choice was generated by the appropriate disposition. There would simply be no fact of the matter whether someone faced with a single rescue case selected on the basis of the right disposition. Since it is selection on the basis of the disposition that ensures that the criterion of right is fulfilled, there is no way of knowing whether the action of someone faced with a single rescue case is right or not.

Perhaps the defender of WL will reply that the rescuer need not actually run the lottery, or even act from a disposition that mimics the lottery – her choices must simply match the probabilities of the lottery. And, since the lottery is more likely to choose the larger group, she should simply save the many. Read in this way, WL does not require a mysterious mental process that can mimic lotteries. Instead, it is like the balancing solution and the egalitarian solution in requiring saving the many, except that it will hold that this is the right course of action because that is the action that a fairly weighted lottery is most likely to select.

Since part of the motivation for WL is a desire to avoid illegitimate aggregation, the proponent of WL would have to show that its automatically saving the many in this way avoids aggregating even if the balancing solution and the egalitarian solution does not. And if it does avoid aggregation, WL will have been transformed from an *alternative* to simply saving the many to a *justification* for doing so. If the motivation for WL was the sense that automatically saving the many fails to treat each person's claim fairly, it is hard to see how WL satisfies its own criterion for an adequate solution to the Number Problem.

Moreover, the proponent of WL cannot simply claim that we ought to save the larger group because that is what a fairly weighted lottery

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

is likely to require. It is true that, in each individual case, WL is more likely to select the larger group. But that does not mean that WL is likely to select the larger group in *every* case. In fact, this is relatively unlikely. In the choice between five and one, WL has a probability of 5/6 of selecting the larger group. But the probability of selecting the larger group six times in a row is relatively low: around 1/3. And the probability of selecting the larger group ten times in a row is much lower still – just under 1/6. To conclude that *every* choice is likely to select the larger group on the grounds that *each* choice is likely to do so is a mistake. So even if we were to grant that a person faced with the rescue case a single time might legitimately choose to save the many, it would be a mistake to conclude on this basis that WL licenses a general policy of saving the many for someone faced with repeated instances of the Number Problem.

The person who faces repeated instances of the Number Problem would therefore have to find some way of ensuring that her choices are distributed in a way that matches the probabilities of a fairly weighted lottery. And in order to do so, she would have, at least, the kind of information about ratios that I argued above is often difficult or even impossible to acquire. Moreover, if the motivation behind WL is to treat each claim to be rescued equally, people stranded on the islands might legitimately complain that the claims of those involved in *previous* rescue cases – who have either been rescued or have died – should not influence the decision on what to do in *this* case.

Moving away from casting WL as a decision procedure, then, would represent a significant concession on the part of its defenders, who see procedures as a way of avoiding the unfairness of automatically saving the many. Even if such a concession is made, however, WL fails to be compelling as a solution – it either requires an entirely mysterious disposition that mimics already demanding lotteries, is reduced to another justification for saving the many (thus undermining its original purpose), or cannot explain how to handle repeated choices between the many and the few. Dropping the requirement of running the lottery, then, would not make WL any more plausible as a solution to the Number Problem.

VII. WHAT KIND OF 'PROBLEM' IS THE NUMBER PROBLEM?

The source of WL's failure, I suspect, is that it misunderstands the problem described by the Number Problem, and so misidentifies the desiderata for a plausible solution to that problem. Since Taurek's essay, the philosophical debate about the Number Problem has largely treated it as a problem for non-consequentialist moral *theory*: should

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press

the correct moral theory take numbers into account, and if so, how can it do so without committing the sins of consequentialism? (In this theoretical sense, the 'Number Problem' is no problem at all for consequentialists.) This is a genuine philosophical problem: how is it that we can reconcile a commitment to fairness to minorities with a recognition that the number of people affected can make a difference? But the Number Problem is not only a theoretical problem for moral *philosophers*; it can *also* be a practical problem for moral *agents*. In fact, it is a theoretical problem in part *because* it is a practical problem. Actual moral agents have to face choices between the many and the few, and for such agents the 'problem' is not how to give a theoretically satisfying answer to the question of how to take account of the numbers of people with a stake in the outcome of the choice, but rather the problem of *what to do*. It is because the problem is pressing on moral agents that it is pressing on moral theory, and plausible solutions to the problem should be ones that take into account the practical nature of the problem. That is not to say that moral agents find the choice *difficult:* one of the aims of some non-consequentialists is to solve the problem in a way that recognizes the common-sense view that the choice is generally *easy*, and that saving the many is clearly the right thing to do. But even if the choice is easy for moral agents, it is still a choice that such agents have to make. Because WL single-mindedly directs itself at the philosophical problem, the solution it identifies is one that moral agents often cannot implement in practice. This is a problem for WL, since its requirements can prevent agents from being able to make a choice at all.

Our theorizing about moral problems should take the agents who face such problems into account, and this is something that WL fails to do: far too many moral agents faced with the often hard choice between the few and the many have neither the means with which to run a properly weighted lottery, nor the information required in order to do so. Perhaps a lottery is one justifiable way of solving choices between the few and the many in cases with full information and a wide selection of dice.[35] WL cannot, however, be an adequate solution to the general question of how we ought to choose between the many and the few.[36]

mdoucet@uwaterloo.ca

[35] Though perhaps not: see John Broome, 'Selecting People Randomly', *Ethics* 95 (1984), pp. 38–55; Kamm, *Morality, Mortality*, pp. 136–40; and Hirose, 'Weighted Lotteries in Life and Death Cases'.

[36] Thanks to Rahul Kumar, Gwen Bradford, and two anonymous reviewers at *Utilitas* for helpful comments on an earlier version of this article.

https://doi.org/10.1017/S0953820812000374 Published online by Cambridge University Press