When Helping *the* Victim Matters More Than Helping *a* Victim

SCOTT M. JAMES
University of North Carolina, Wilmington

Consequentialists insist there is no rational basis for distinguishing between determinate (or identifiable) victims and indeterminate (or statistical) victims. Whether it's a child drowning at our feet or needy communities abroad, our reason to help is the same. Experimental data indicate, however, that we regularly make such distinctions. In this article, I show that there are indeed persuasive normative grounds for preserving this distinction. When potential beneficiaries are determinate, they have a special claim on us grounded in fairness. I present several cases that demonstrate that treating determinate beneficiaries the same as indeterminate beneficiaries is unjust. I conclude with an analysis of the relevant social psychology data.

My aim in this article is to show that duties of beneficence are constrained by considerations of fairness. Such considerations have traditionally been overlooked in this context because they appear only after extended reflection on what I'll refer to as the *determinacy* of one's beneficiaries. Central cases of beneficence are such that one's assistance benefits a determinate beneficiary. There is, in other words, a fact of the matter as to whom one benefits when one acts even if, for contingent reasons, one cannot identify who that individual is. But there are other, rarely noticed, cases of beneficence in which one's assistance is not linked to any determinate beneficiary – even when evaluated *ex post*. Instead, one's assistance simply reduces the overall amount of suffering, without there being a determinate beneficiary. It's important to emphasize that in these latter cases, the indeterminacy is not merely epistemic. It's *metaphysical*: we cannot identify whom we've benefited since there is no fact of the matter as to whom one benefits. (That such cases can – let alone, do – obtain may seem implausible, but examples below demonstrate that this is hardly as obscure as it initially sounds.) The question that motivates much of what follows is this: does the determinacy of one's (potential) beneficiary make a moral difference? I argue that determinacy, to the extent that it uniquely introduces considerations of fairness, does indeed make a moral difference. It would thus be unfair – and, hence,

© Cambridge University Press 2013
Utilitas Vol. 25, No. 1, March 2013
doi:10.1017/S0953820812000246

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

wrong – to behave as if these two classes of beneficent acts are morally equivalent.[1]

Consequentialists, perforce, must reject these claims. So long as the number of beneficiaries is equivalent across situations, consequentialists have no grounds for distinguishing between determinate and indeterminate beneficiaries. This result, I maintain, represents a strike against consequentialism. If my arguments are successful, they would also thwart more contemporary efforts to undermine *non*-consequentialism by way of cognitive psychology. Joshua Greene, for example, argues that recent psychology results 'cast doubt' on non-consequentialism as a legitimate normative theory.[2] His argument, roughly stated, is that our non-consequentialist intuitions and judgements – unlike their consequentialistic counterparts – are produced by emotional 'centres' of the brain, which happen to be tuned to morally *irrelevant* features of our environment. We are, for example, responding to the 'up close and personal' aspect of situations in making stereotypically non-consequentialist judgments rather than (as non-consequentialists allege) a person's rights or inherent dignity or whatever.[3]

According to Greene, non-consequentialism amounts to little more than 'moral confabulation'. It is a sophisticated story we tell ourselves to make sense of our underlying feelings. In making his case, Greene cites the 'identifiable victim effect',[4] according to which people are typically more responsive to the needs of identifiable (or determinate) victims than non-identifiable (or statistical) victims. But this discrimination, claims Greene, is not grounded in any rational moral principle, but is, instead, just another manifestation of our emotional preoccupation with those up close and personal.

[1] To be clear, this is not an examination of fairness itself (I do not offer any substantive account of the nature of fairness), but rather the way in which the determinacy of potential beneficiaries can directly affect the fairness (and, hence, morality) of available actions.

[2] J. Greene, 'The Secret Joke of Kant's Soul', *Moral Psychology*, vol. 3, ed. W. Sinnott-Armstrong (Cambridge, MA, 2008), pp. 35–79.

[3] Greene softens his position later by suggesting that the 'up close and personal' feature represents just a 'first-cut' hypothesis in need of refining. He is nevertheless committed to the idea that it is at root an emotional response that triggers stereotypical non-consequentialist judgements.

[4] T. C. Schelling, 'The Life You Save May Be Your Own', *Problems in Public Expenditure Analysis*, ed. S. B. Chase (Washington, D.C., 1968), pp. 127–76.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

My arguments would show that Greene is mistaken on this front.[5] I show that the determinacy of one's beneficiary can in fact make a moral difference – to the extent that it uncovers morally significant features like fairness. Again, this would expose a lacuna in consequentialism that has so far gone unnoticed. At any rate, these findings suggest that the empirical data do not show what Greene thinks they show. I conclude with a proposal for how to interpret subjects' behaviour under experimental settings.

DETERMINACY'S 'DISTORTING' INFLUENCE

The thought that determinacy might constitute a distinct normatively significant feature has arisen, among other places, in discussions of Peter Singer's now-familiar arguments for strict duties of beneficence.[6] On Singer's view, 'we have an obligation to help those in absolute poverty that is *no less strong* than our obligation to rescue [e.g.] a drowning child from a pond'.[7] Since people overwhelmingly resist Singer's conclusion, one might be tempted to offer – as an explanation of this resistance – this non-consequentialist response: the drowning child whom I could assist is *determinate*, whereas the suffering children abroad whom I could assist are *indeterminate* or statistical.[8] And my reason to respond to determinate individuals is more stringent than my reason to respond to indeterminate individuals.

Greene, however, maintains that there's a more plausible, *non-moral*, explanation of our responses. Because just the thought of the drowning child – being up close and personal – arouses my emotions in a way that those distant children do not, I find myself (presumably at a subpersonal level) with disparate affective states in need of interpretation – or, as Greene puts it, 'post-hoc rationalization'. Determinacy just happens to be one straw I might grasp at in concocting 'a rationally appealing story'. (I could have grasped at others.) But the point is that I do not, in the moment, grasp determinacy as a

[5] For criticism of Greene's general strategy, see F. M. Kamm, 'Neuroscience and Moral Reasoning: A Note on Recent Research', *Philosophy and Public Affairs* 37 (2009), pp. 330–45. Kamm cites instances of non-consequentialistic judgments in cases of *non*-'up close and personal' situations as evidence that non-consequentialism is indeed sensitive to factors beyond the 'up close and personal' aspects.

[6] P. Singer, 'Famine, Affluence, and Morality', *Philosophy and Public Affairs* 1 (1972), pp. 229–43; P. Singer, *Practical Ethics* (Cambridge, 1993).

[7] Singer, *Ethics*, p. 230.

[8] While my $50 contribution to UNICEF aids the agency's efforts, there is no fact of the matter as to which individual(s) were the direct beneficiaries of my contribution. See S. James, 'Good Samaritans, Good Humanitarians', *Journal of Applied Philosophy* 24 (2007), pp. 238–54, for a defence of this position and P. Unger, *Living High and Letting Die: Our Illusions of Innocence* (Oxford, 1996), ch. 3, for criticism.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

normatively significant feature of the situation. Instead, I fabricate a conceptually appealing story that makes sense of my feelings, perhaps 'with the help of some especially creative philosophers'.[9]

To strengthen his case, Greene draws attention to experimental settings wherein determinacy exerts an (apparently) irrational influence on subjects' decision-making. He cites the work of Small and Loewenstein.[10] In one experiment, subjects were given an 'endowment' of ten dollars. Half of the subjects, however, were forced to give up their endowment as a result of a randomized procedure; these subjects constituted the experiment's 'victims'. Non-victims – i.e. those who retained their endowments – were then asked, individually and in private, whether they would donate some of their endowment to a specific victim. But the experiment was designed to test the role determinacy plays in how much (if any) non-victims donated – in the following way. In one group, non-victims chose the number corresponding to a determinate victim *before* their decision to donate was made. Thus, in this setting a non-victim was deciding how much to give to a determinate victim.[11] In the second group, non-victims chose the number corresponding to a determinate victim *after* their decision to donate was made. Thus, in this setting a non-victim was deciding how much to give to a to-be-determined victim.

So how did this difference affect subjects' level of charity? For those non-victims responding to determinate victims, the mean donation was *60 per cent higher* than the mean donation of non-victims responding to indeterminate victims. The median donation for non-victims responding to determinate victims was *twice* that of the other group. In a separate experiment, Small and Loewenstein derived similar results when they gave subjects a chance to donate to Habitat for Humanity. The first group was asked to donate to a needy family that had already been selected (though the identity of the family was never revealed). The second group was asked to donate to an as-yet-to-be-selected family. As in the earlier study, the mean donation in the former group was higher (by 25 per cent) than the mean donation of the latter group. And the median donation of the former group was twice the median donation of the latter group.

As a matter of psychological fact, then, determinacy is a potent force. But does it deserve to be potent *normative* force? Greene

[9] Greene, 'Secret Joke', p. 36.

[10] D. A. Small and G. Loewenstein, 'Helping a Victim or Helping the Victim', *Journal of Risk and Uncertainty* 26 (2003), pp. 5–16.

[11] The actual identities of victims were never disclosed, however. So while a non-victim might know that this number directly corresponds to some determinate individual somewhere, she would never learn any identifying features of that individual.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

thinks this is absurd. According to Greene, 'there's no rational basis' for giving more to a determinate individual than to an as-yet-to-be-determined individual.[12] Whatever is going on when subjects discriminate between the two has no normative foundation. Morally speaking, it doesn't matter whether our beneficiaries are determinate or not. What the Small and Loewenstein experiments demonstrate is that subjects indeed discriminate on the basis of determinacy. But such discriminations, according to Greene, make no normatively relevant difference.

Now why should all this constitute a threat to non-consequentialism? Greene does not explicitly say. After all, he rehearses these experimental findings to support only the hypothesis that our judgements (in the relevant cases) are produced by an emotional reaction to the determinacy of one's beneficiary. Nevertheless, it's fairly clear what Greene is up to. Non-consequentialism is threatened by these specific results when we conjoin his empirical hypothesis with the following two philosophical claims: (1) determinacy does not carry any independent moral value and (2) non-consequentialism assigns independent moral value to determinacy.[13] Greene does not offer any arguments in defence of (2), but the claim is not entirely unmotivated. For non-consequentialism is inherently *non-aggregative*, which entails that our moral duties are not individuated by reference to the total amount of good our actions produce. Rather, our moral duties are individuated by reference to how we treat determinate individuals. I cannot respect the rights of the overall state of world; I can only respect the rights of determinate individuals. For the purposes of this article, I'm prepared to accept (2). What I deny is (1). In the next section I mount a defence of the moral significance of determinacy.

BENEFITS WITHOUT BENEFICIARIES

The first point to insist on follows as a matter of logic: from the (alleged) fact that determinacy carries no normative weight in the Small and Loewenstein experiments, it does not follow that determinacy does not *ever* matter. It may be that, under the conditions of the experiment, whatever normative weight is assigned to determinacy is overridden by other normative factors. It could be the case that there is a *pro tanto* reason to attend, first and foremost, to determinate individuals. But under the specific circumstances, that reason is overridden by other

[12] Greene, 'Secret Joke', p. 19.

[13] True, even without Greene's hypothesis, non-consequentialism is threatened by the conjunction of (1) and (2). What Greene's hypothesis adds is an (allegedly plausible) explanatory alternative to the non-consequentialist account, thus rendering non-consequentialism superfluous.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

reasons, just as one's reason to keep a promise can be overridden by other reasons under specific circumstances. Be that as it may, what we need to investigate is whether or not the determinacy of our beneficiaries *ever* matters morally.

The second point I wish to argue for, then, is that determinism can indeed matter. Although my argument will turn critically on the details of a particular case, it is essential to emphasize that the success of my argument does not depend on any case-specific moral intuition. Indeed, most respondents (myself included) do not report any robust moral intuitions when confronted with the forthcoming case. But this is just as well since my argument for determinism can be made without any appeal to case-specific moral intuitions. My argument relies, instead, on reflection on general moral principles, which are quite compelling in their own right. This approach, I contend, exposes an additional problem with Greene's general methodology: if the case for non-consequentialism can be built without appeal to case-specific moral intuitions, then subjects' case-specific moral intuitions – and the means by which they're produced – do not vitiate non-consequentialism.

Consider, at long last, the following case, a case in which one's decision hinges exclusively on the putative significance of determinacy.

A cruise-liner, with 1,000 passengers aboard, is taken hostage by a terrorist group. For undisclosed reasons, they intend to choose 500 passengers at random to kill unless their demands are met. During the course of negotiations, the terrorists (again, for undisclosed reasons) offer to free just *one more* hostage for ransom. The government, perhaps to buy time, agrees. But *how* this additional individual will be chosen remains undecided. The terrorists offer two options (but express complete indifference as to which option is chosen). On Option A, the terrorists will order their men to choose at random only 499 passengers to kill instead of 500. Option B is more complicated. As it turns out, when the ship was taken over in the middle of the night, every passenger on the ship was in his room (on the sleeping deck), save *one*. This passenger (call this passenger no. 1,000) happened to be strolling the main deck when the ship was commandeered. No. 1,000 was thus locked in a storage room apart from the rest of the passengers. On Option B, the terrorists will kill no. 1,000 along with 498 randomly chosen passengers from the sleeping deck. To be clear, on Option A, no. 1,000's life will be spared; the 499 passengers to be chosen randomly will be chosen from the sleeping deck only.

Now assuming that the decision has been made to spare one more person from dying, are there any moral reasons to prefer one option over the other? If so, what are those reasons?

One might immediately doubt whether there is any such reason to prefer one option over the other. After all, whether you choose Option A or Option B the outcome is identical: 501 lives saved and 499 lives lost. The only difference between the two options is one of *determinacy*: on Option B there is no determinate beneficiary; on Option A there

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

is. To see this, imagine a world wherein you choose Option B. And suppose that in that world the passenger in room no. 358, say, is *not* selected to be killed. There nevertheless remain many near *possible* worlds wherein passenger no. 358 *is* selected to be killed even though you chose Option B in that world. *Mutatis mutandis* for worlds where Option A is taken, no. 358 cannot claim that his life was lost as a result of your choosing Option A, since, for all that anyone knows, he would have lost his life even if Option B was taken – indeed, even if *no* option was taken. It must be pointed out though that on Option B, while there is no determinate beneficiary, there *is* a determinate victim – viz., no. 1,000. Option A reverses this: while there is a determinate beneficiary (i.e. no. 1,000), there is no determinate victim. What we know, by choosing Option B, is that 499 individuals disembark safely as opposed to 500. What we do not know, nor could know, is who that additional individual is.[14] Not so in the case of Option A.

Now if determinacy did not make a moral difference, if, that is, the determinacy of a potential beneficiary (or victim) should not be taken into account in deliberating what you ought morally to do, then you should be, at most, indifferent as to which option you take. Again, apart from determinacy, the options are equivalent. But I submit that you should *not* be indifferent as to which option you ought to take. And the reason (I shall defend below) is this: by choosing Option B you bestow a minuscule benefit on any one individual on the sleeping deck *at the cost of no. 1,000's life*, and this is inconsistent with moral values we already adhere to – fairness, first and foremost.

Why should we regard the benefit bestowed in Option B as 'minuscule?' You are, after all, reducing the number of deaths by one. The reason is: the most that any one individual (on the sleeping deck) can claim is that by choosing Option B you increase his chances of being spared by a tiny fraction. That fraction can in fact be measured. For any individual on the sleeping deck, the probability of his being spared, on Option A, is 0.4995.[15] By choosing Option B, you increase his chances of being spared to 0.5005, an increase of 0.001. Hence, for any individual on the sleeping deck, the difference between Option A

[14] This fact marks an important distinction between this case and, for example, another, oft-discussed, example introduced by Sophia Riebetanz, 'Contractualism and Aggregation', *Ethics* 108 (1998), pp. 296–311. In that example, a landmine is buried somewhere in a field in which 100 labourers are about to work. As with the cruise-liner case, there is no fact of the matter as to who will be harmed when the case is evaluated *ex ante*. But unlike the cruise-liner case, there *is* a fact of the matter as to who was harmed when the case is evaluated *ex post*. We can identify who was injured by the landmine. In the cruise-liner case, we cannot identify the additional life lost on Option A simply because there was no additional *determinate* life lost.

[15] Remember that, with no. 1,000 partitioned off from the rest of the passengers, 500 individuals will be chosen from *999* – not 1,000.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

and Option B is *one-tenth of 1 per cent*. But this is not the case at all for no. 1,000. For no. 1,000, the difference between Option A and Option B is the difference between life and death.

WHY DETERMINACY MATTERS

Let us concede that every individual on the ship has a reason to want to avoid being killed. The question concerns whether any entity has a *greater moral claim* to your assistance than another, where the 'entity' in question may, for all we know, consist of a group of individuals. If there is such an entity, then we should not be indifferent as to which option you take. Three possibilities present themselves. There is, first, the group of individuals on the sleeping deck, that is, the 999 passengers *taken as a single entity*. Second, there is any single individual on the sleeping deck. And, finally, there is no. 1,000. Let's consider these in turn.

It might be asserted that the *group* has a greater moral claim to your assistance than no. 1,000. To defend this view requires more than pointing out that Option B will guarantee the survival of one more individual, for this observation in fact supports the hypothesis that one should be indifferent as to which option you take – not the hypothesis that you should place greater moral weight on one entity over another – since no. 1,000 has the identical claim: Option A will also guarantee the survival of one more individual. The only difference lies in the identity of the surviving individual. More plausibly, one might argue that in choosing Option A you treat the group unfairly since you are putting one individual's claim for assistance ahead of the claims of 999 individuals. But if these individuals, taken together, have the same claim as no. 1,000 – viz., that the option favouring them will result in one less death – then this would again support the view that we should be indifferent as to which option we take. We should not attach special moral weight to the group's claim when no. 1,000 has the identical claim against you. And vice versa. Thus the view that the group has a moral claim to your assistance, if coherent, collapses into the view that we should in fact be indifferent as to which option we take.

On the other hand, if the assertion is that any individual on the sleeping deck has a greater claim on your assistance because, in general, one has an overriding moral reason to want to increase his chances of survival, then no. 1,000, *if anyone*, should have a greater claim on your assistance since Option A increases his chances to one. It is instructive to consider this possibility in more detail because I believe it brings into stark relief why no. 1,000 has indeed a greater moral claim on your assistance than any other entity under the circumstances.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

As we've already conceded, every passenger on the sleeping deck has a reason to want to be rescued when it is in your power as a potential rescuer to do so; therefore, in that sense each has a claim on your assistance. So there is no denying that by choosing Option A you will not be directly honouring any individual's claim except no. 1,000. But it is essential to bear in mind what any such individual's claim is a claim to – viz. to increase his chances of survival by 0.001. *Ex hypothesi*, this is the *most* you can do for him. It is thus false (let alone misleading) to suppose that his claim is a claim to anything like 'being saved'. So what would any one on the sleeping deck be demanding in demanding that you choose Option B? He would be demanding that you allow no. 1,000 to die *so that his chances of survival could be increased by 0.001*. Here the point is sharpest. This is not a justifiable demand. And it is not merely indecent. To demand such a thing is wilfully to ignore the most basic recognition of the importance to others (and, for that matter, to oneself) of not dying prematurely.

Moreover, such a demand is inconsistent with other aspects of our behaviour. For example, I willingly – indeed, cheerfully – increase my chances of dying prematurely by an even greater amount (i.e. greater than 0.001) every time I get into car or aeroplane, swim in the ocean, or eat ground beef. We regularly and consciously accept such risks even when the rewards are minor. But if increasing our chances of dying prematurely by only a fraction meant *that* much to us, if we were indeed prepared to allow another person to *die* in order to avoid such risks, it would be altogether mysterious why we would so willingly and so frequently engage in such activities when doing so was quite unnecessary. But it is not mysterious at all. There are innumerable reasons why we accept such trade-offs; life is too important – too rich in experience and experiment – to cut oneself off from any risk. The point, however, is that if our reasons for accepting an increase in our chances of dying prematurely can include such things as aesthetic and intellectual considerations (to name just two), then surely they could – and indeed should – include someone's losing his life as a relevant consideration.

No. 1,000's predicament is altogether different from the predicament of any other passenger, for you have the option of saving his life, not merely increasing his chances of survival. By choosing Option B you allow him to die. It is plain then what no. 1,000 is demanding when no. 1,000 demands that you choose Option A. It is a demand that each person on the sleeping deck allows his chances of dying prematurely to increase by 0.001 *so that one person may live*. That this difference has moral implications (whatever they might be) is an intuition that appears to be shared by many different audiences.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

The problem with indifference in this case is that you are assigning equal weight to the claims of every member on the ship, including no. 1,000, when the weights of the various demands are patently unequal. There is a critical difference between demanding that someone decrease his chances of survival by a fraction of 1 per cent so that someone else may live, on the one hand, and demanding that someone be allowed to die so that another can increase his chances by a fraction of 1 per cent, on the other. However one chooses to locate this datum in a moral theory, it should be clear that reason strongly favours drawing a principled distinction.

We can thus complete our argument. Since the only difference between Option A and Option B is one of determinacy, and since this difference yields important moral differences in this case, then determinacy makes a moral difference. It makes a moral difference by exposing the morally unique position a determinate person occupies. Beneficent acts disproportionately bear on that individual's well-being, such that treating that individual as just another (potential) beneficiary is unfair. To the extent that non-consequentialism assigns moral value to fairness in this way – and consequentialism doesn't – non-consequentialism holds a theoretical advantage.

CLARIFICATIONS AND OBJECTIONS

I've argued that since determinacy is the only feature that distinguishes Option A from Option B and since Option A is morally preferable to Option B, determinacy makes a moral difference. As stated, the conclusion may tempt one to infer that determinacy *in itself* matters morally. But that is not a claim I've tried to defend. I've defended only the weaker claim. Careful examination of the cruise-liner case reveals only that determinacy functions as a way of segregating individuals in such a way that some available acts (e.g. Option B) come to instantiate a morally salient property that other available acts (e.g. Option A) lack, namely, the property of being unfair. Option B treats each passenger, including no. 1,000, as if her predicament was the same as that of any other, an approach that would be fair if the predicament of each was the same. But no. 1,000's predicament is not like that of any other.[16] No. 1,000 occupies a different moral space, so to speak, than any other passenger. That Option A would save a determinate

[16] By the same token, of course, choosing Option A guarantees no. 1,000's survival at the cost of reducing (by a tiny fraction) each of the remaining passengers' chances of survival. Does this mean that *they* are treated unfairly? It's hard to see why. After all, as far as each of the rest of the passengers is concerned, it's only a statistician who would discern any difference at all between Option A and Option B. Life or death is, by almost any measure, a coin-flip – whichever option is chosen.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

individual matters morally only because, qua determinate individual, no. 1,000 can literally be saved at a negligible cost to others, and that is a normatively significant consideration that counts persuasively in favour of Option A. The moral matter thus recommending Option A does not reside in the nature of determinacy itself; instead, it draws on considerations that only determinacy exposes.

One might legitimately wonder, however, whether these results speak more to the arbitrary features of the case than to determinacy itself. After all, the cost to others is 'negligible' only as a consequence of the arbitrary number of passengers aboard the ship. One's chances of not being selected to be killed rise or fall as a direct function of the number of passengers on the ship. With 500 passengers instead of a 1,000, one's chances of not being selected improve by 0.002 on choosing Option B. With 100 passengers, one's chances improve a hundredfold over the original case – to 0.01. At the extreme, with only six passengers, one's chances of being selected to be killed on Option B are 0.2 (no. 6 is automatically selected along with one randomly selected individual among the remaining five), but choosing Option A increases one's chances of being killed to 0.4, a 20 per cent increase. In such a case, no. 6's demand that you choose Option A over Option B represents a non-negligible demand on others, for now two (of the five) will be randomly selected to be killed instead of just one. And it's far from obvious to me that acquiescing to that demand is obligatory. This would appear to vitiate my claim that determinacy matters morally since, in the six-passenger case, the demands of a determinate individual are being held up against the demands of the remaining passengers. But here the latter demands are at least as normatively persuasive as the former.

This objection, however, misidentifies the morally critical matter. Were one committed to the claim that determinacy *itself* made a moral difference, then the objection might stick. But I'm committed only to the claim that determinacy, by virtue of the considerations determinacy exposes, matters morally. What we see in the six-passenger case is that determinacy exposes, but does not settle, how morally to proceed. Is it unfair or unreasonable to demand that individuals increase their chances of dying by 20 per cent so that another individual may be saved? To the extent that my intuitions tell me anything, they tend towards the affirmative. But that's a matter to be resolved in the course of working out a conception of fairness, something I've made no attempt to do. What I hope to have done is shown how attending to the determinacy of beneficiaries draws out normatively significant considerations (such as fairness) that are otherwise opaque to the consequentialist. For the consequentialist will be indifferent (*ceteris paribus*) to which option is chosen *no matter how many potential beneficiaries are in need.* The non-consequentialist can concede that the normative equation becomes

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

murky as the number of beneficiaries decreases (as we saw in the six-passenger case). But that's a theoretical price worth paying. For the consequentialist confronts the much steeper price of having to deny that no. 1,000 occupies a different moral space, so to speak, from the rest of the passengers. For the consequentialist, 501 saved are 501 saved.

This brings me back to more general implications for beneficence. Recall that Singer-style arguments for strict duties of beneficence attempt to undermine any distinction between preventing n determinate individuals from dying soon and preventing a group of individuals from being reduced by n, where there is no fact of the matter as to whose life is lost in the reduction. If successful, these arguments show that I have just as much reason to donate to Oxfam (where, presumably, my donation will prevent communities of needy individuals from being further reduced) as I do to rescue a child drowning in front of me. That there is no fact of the matter as to whom my donations to Oxfam have saved is supposed to be morally irrelevant.[17] I've argued that there are in fact normative grounds for this distinction. The fate of that child at my feet is 'at my mercy' in a way that the fate of no potential beneficiary of humanitarian relief is. No distant child in need, that is, can claim that my *failure* to donate to Oxfam last month is responsible for his imminent death – just as no such child can claim that my decision to donate last month is responsible for his survival. His claim is instead analogous to the claim of any one passenger on the cruise-liner: how I choose to respond to solicitations for help nudges his chances of survival up or down by some slight fraction. Conflating this claim with the claim of the child at my feet – as the consequentialist will wish to do – makes no moral sense.

Up to this point, my argument has not so much turned on any specific case as on general normative principles. But it's possible to marshal case-specific moral intuitions in strengthening my case. Here's an example adapted from Paul Gomberg:[18]

Vintage Boots. As the owner of a pair of vintage boots, you decide one day to sell the boots and donate the money to humanitarian relief efforts abroad. You don the boots and strike out in the direction of a local pawnshop. On the way there, however, you come across a child drowning in a shallow, deserted pool. Before rushing to the child's aid, it occurs to you that if you enter the pool in

[17] For a defence of the non-normative claim that there is no fact of the matter as to who (if anyone) is saved by our donations to large humanitarian organizations like UNICEF and Oxfam, see James, 'Good Samaritans, Good Humanitarians'.

[18] P. Gomberg, 'The Fallacy of Philanthropy', *Canadian Journal of Philosophy* 32 (2002), pp. 29–66.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

the boots, they'll be ruined and, once ruined, worthless. What's more, you need the assistance of another person to remove the boots (you've nearly outgrown them), but no such assistance is available. You reason that the money you planned to send abroad will likely save a few more lives than the single life you could save here and now. On these grounds, you allow the child to drown and continue on to the pawnshop where you sell the boots for humanitarian ends.

While those already devoutly committed to utilitarianism will approve of the action, the intuitions of those with no moral theoretical commitments reflect strong disapproval. Some report being repulsed by the idea of turning one's back on the drowning child. It's impossible of course to locate the source of subjects' responses to this case, but I have offered a principle explanation that predicts exactly the pre-theoretic response observed.

Like no. 1,000 in the cruise-liner case, the child in Vintage Boots is 'at your mercy' in a way that no individual in needy communities abroad is. Treating the child's plight as no different from the plight of those abroad is, in correlative fashion, unfair – and morally unjust. Perhaps it is a coincidence that people report unease at the thought of turning one's back on the child in favour of needy individuals; however, it's also possible that that unease stems from an inchoate awareness that it matters morally that there is a fact of the matter as to whom who could save.

REVISITING SMALL AND LOEWENSTEIN

There is still the matter of the Small and Loewenstein experiments. If we're correct that determinacy can expose important moral considerations, does it make a moral difference in those experiments? If, as Greene argues, it 'has no rational basis', why do subjects appear to think it does? My position is that Small and Loewenstein's subjects' responses are, all things considered, irrational, though perfectly understandable. They're irrational because the victims are indeterminate *only* when assessed *ex ante. Ex post,* victims are determinate. Subjects' responses are understandable because we almost never confront situations that are structurally analogous to the Small and Loewenstein experiments – that is, situations in which potential beneficiaries are statistical (or indeterminate) when assessed *ex ante*, but determinate when assessed *ex post.* Either beneficiaries are determinate both *ex ante* and *ex post*, as when you rescue a drowning child or offer a few dollars to a homeless individual; or beneficiaries are indeterminate both *ex ante* or *ex post,* as when you donate to your church or contribute to humanitarian relief efforts or pay taxes. The Small and Loewenstein experiments present subjects with a quite unfamiliar situation. So is it any wonder, then, that undergraduates

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press

(who, incidentally, are given no incentive to conform their responses to 'correct' standards) fail to engage in the kind of deliberation necessary to recognize the implications of shifting determinacy?

For it might be that a non-victim in the 'non-determined' group (let's call her X) is overly influenced by the fact that, when evaluated *ex ante*, no victim is at the complete mercy of X. By contrast, a non-victim in the 'determined' group (let's call her Y) might recognize that, when evaluated *ex ante*, there is an individual at the complete mercy of Y. In the former case – but not in the latter case – any potential victim (A) has a chance of receiving more than X contributes since A has a chance of being linked to *other* non-victims. This highlights the crucial difference between the Small and Loewenstein study and the cruise-liner case. At the conclusion of the Small and Loewenstein experiment, we can identify a determinate victim whose compensation (if any) was directly linked to a non-victim's contribution. But as we saw in the cruise-liner case, we could never identify a determinate beneficiary on Option B (even evaluated *ex post*), nor could we ever identify a determinate victim on Option A. It's this difference that eliminates the morally significant features that arose in the cruise-liner case. So, on one charitable reading of the Small and Loewenstein study, subjects' deliberations are being influenced by appropriate *ex ante* evaluations, but fail to appreciate matters *ex post*.[19]

jamess@uncw.edu

[19] Several (determinate) individuals deserve thanks for helping me get these ideas into their final form: Brian Ballard, Jason Bowers, Peter Carruthers, Matt Eshleman, Mark Schroeder, Matt Waldschlagel and Heath White.

https://doi.org/10.1017/S0953820812000246 Published online by Cambridge University Press