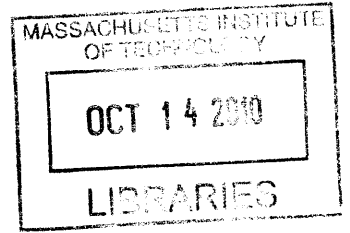# Help! Not Just Anybody.
Essays on Altruism and Conflicts of Interest
by
Tom Dougherty
B.A. Philosophy, Politics and Economics
Oxford University, 2004

Submitted to the Department of Linguistics and Philosophy in Partial Fulfillment of
the Requirements for the Degree of Doctor of Philosophy in Philosophy at the
Massachusetts Institute of Technology.

September 2010

Signature of author:...............................

Tom Dougherty
June 18th 2010

Certified by:..

Richard Holton
Professor of philosophy
Thesis Supervisor

Certified by:..

Rae Langton
Professor of philosophy
Thesis Supervisor

Accepted by:..............

Alex Byrne
Professor of philosophy
Chair of the Committee on Graduate Students

1

Help! Not Just Anybody
Essays on Altruism and Conflicts of Interest

Submitted to the Department of Linguistics and Philosophy on June 18th 2010 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Philosophy.

Thesis Abstract

If you can help someone without detriment to your interest, nor anyone else's interest, then it is clear that you ought to do so. But things are not always so easy. When there is a conflict of interest, you have to decide what to do. Whom should you help when you cannot help everyone? How much should you sacrifice to help others? May you expose people to risks when helping them? This dissertation addresses aspects of these questions.

You ought to save a larger group of people rather than a distinct smaller group of people, all else equal. Why? Chapter 1, "Rational Numbers," offers an explanation. Its two parts can be roughly summarized as follows. First, you are *morally* required to want each person's survival for its own sake. Second, you are *rationally* required to achieve as many of these ends as possible, if you have these ends.

Chapter 2, "Ambition and Altruism in the Dynamic Moral Life," poses a puzzle. We would like an account of beneficence to be moderately demanding, and yet still to require you to be ambitious with your altruism. How can these diverging desiderata be simultaneously met? Drawing on empirical work, the chapter defends the following solution: beneficence requires you to develop morally, and increase how much you give over time.

Chapter 3, "Chancy Charity and Aggregative Altruism," argues that two initially attractive claims are inconsistent. First, you must save someone's life rather than cure the headaches of many. Second, you may take a small risk of someone's death when curing this person's headache. Since we are unable to hold both these claims, we are in danger of lacking an explanation of some common intuitions about risk and the priority of serious needs. A candidate explanation is considered, but criticized.

Thesis supervisor: Rae Langton
Title: Professor of Philosophy

Thesis supervisor: Richard Holton
Title: Professor of Philosophy

Acknowledgments

Chapter One
Rational Numbers

1.   The numbers problem

When we cannot help everyone, we must make decisions about whom to help. Should we take into account the number of people we could help? Consider:

> Drug. You have 100ml of a drug that you own. Two strangers, Anna and Bert, each need 50ml to live. A third stranger, Clare, needs all 100ml to live. The strangers are in all relevant respects alike. You do not stand in any special relationship to these strangers. All else is equal.

What should you do? Most of us think that you ought to save Anna and Bert. In John Taurek's phrase, we think that "the numbers should count."[1] I will assume that the numbers should count. The question I will address is *why* they should.

Some of us are looking for an explanation of why the numbers should count because we reject what I will call the "Aggregative Explanation." This explanation has three parts. First, each person's survival bears some value. Second, the values corresponding to the survival of each member of a group sum to produce a greater total value corresponding to saving the whole group. Third, when it comes to helping strangers, "the sum of [this] value is to be maximized."[2] Why do people reject this explanation? There are two independent motivations worth considering.

The first motivation arises from denying that the reason why you ought to help people in a case like Drug is that doing so *produces value*. There are several reasons why people deny this. First, some are independently attracted to a moral theory that does not explain what you ought to do in terms of producing value. A virtue ethicist may explain this in terms of the actions a virtuous person would perform. A Kantian may do so in terms of whether you can rationally endorse a world in which your maxim for action is universal law. Second, some worry that a focus on producing value makes it hard to account for partiality and respect for rights.[3] Third, some deny that it makes sense to say that saving the larger group would produce more value. They reject the notion that there is a legitimate impersonal point of view from which we can evaluate saving them as "better."

---

[1] John Taurek, "Should the Numbers Count?" *Philosophy & Public Affairs*, 6 (1977): 293-216.

[2] Thomas M. Scanlon, *What We Owe to Each Other*, (Cambridge MA: Harvard University Press, 1998), p. 230.

[3] David Wasserman and Alan Strudler note that in order to seek an alternative rationale for saving the larger number, "one might seek to qualify or constrain consequentialist reasoning by adopting a pluralist moral theory that mixes or integrates consequentialist and nonconsequentialist elements." Presumably the type of nonconsequentialist elements they have in mind are permissions to be partial, and constraints to respect rights. However they "suspect that they are doomed attempts to breed species that are in essence incompatible." p. 72 of Wasserman and Strudler, "Can a Nonconsequentialist Count Lives?" *Philosophy & Public Affairs* 31 (2003): 71-94; In a similar vein, Judith Jarvis Thomson worries that if our metaethic allows that states of affairs are good *simpliciter*, then we will find inescapable a version of (agent-neutral) act-consequentialism, which cannot accommodate rights and partiality. Judith Jarvis Thomson and Gilbert Harman, *Moral Relativism and Moral Objectivity*, (Cambridge MA: Blackwell Publishers Inc., 1996).

Some find the property of "good, simpliciter" suspect across the board.[4] Alternatively, some take this position only with respect to cases where individuals' interests conflict. Taurek takes the latter position, and Weyma Lübbe neatly summarizes his view:

> Taurek does not even see how speaking of "goodness" and "betterness," as opposed to "goodness for" and "betterness for," is to be understood in such cases. Certainly it is better for the one if he is saved and better for each of the five if the five are saved, but what does it mean to say that it is better, period, if the five are saved?[5]

Although these theorists put their worries in terms of "goodness," I am sure they are no more comfortable with the notion of "value."

The second motivation stems from a worry that aggregation advocates "summing up ... small benefits to many to reach a sum that outweighs ... serious losses to a few."[6] Suppose you could save a stranger's life or cure the headaches of a different group of strangers. Consider the total value of curing the group's headaches, according to the Aggregative Explanation. This total value increases as the number of people in the group increases. Eventually this total value will exceed the value of saving the other person's life. So if the number of people with headaches is large enough, then the Aggregative Explanation will recommend curing their headaches rather than saving the other stranger's life. Some people find this result unacceptable. They hold that, no matter how many headaches you could cure, you ought to save the stranger's life.[7]

As a result, many have been searching for an alternative explanation of why the

---

[4] For example, Judith Jarvis Thomson denies that there is a property of "goodness *simpliciter.*" In support of this claim, she appeals Peter Geach's linguistic argument that "good" functions as an attributive adjective, like "big", and not a predicative adjective, like "black." If Mickey is a big, black mouse, we can infer that Mickey is black, but not that Mickey is big. The objection runs that similarly, if something is a good toaster, we cannot infer that it is good, simpliciter. Judith Jarvis Thomson, "Goodness and Utilitarianism," *Proceedings and Addresses of the American Philosophical Association* 67 (1993): 145-159; Peter Geach, "Good and Evil," *Analysis* 17 (1956): 32-42. Philippa Foot defends a similar position in her "Utilitarianism and the Virtues," *Mind New Series* 94 (1985): 196-209. Meanwhile, Taurek denies that people are valuable objects to be saved. Instead, Taurek claims that you should empathize with strangers, and care about their welfare because their welfare matters to them. Taurek, "Should the numbers count?" pp. 304 - 5 *et passim.*

[5] Weyma Lübbe, "Taurek's No Worse Claim," *Philosophy and Public Affairs* 36 (2008): 69-85. Lübbe refers to Taurek, "Should the numbers count?"

[6] Scanlon, *What We Owe to Each Other,* p. 237.

[7] This problem moves some consequentialists to reject the Aggregative Explanation. Consequentialism is logically distinct from the Aggregative Explanation because, as Scanlon puts it, "a purely teleological conception of value (whether impartial or not) need not be additive. It can hold that the value of a state of affairs is not the sum of the values of certain of its elements, but is arrived at in some more complicated way." Scanlon, *What We Owe to Each Other,* p. 81. Meanwhile, some non-consequentialists accept the Aggregative Explanation. Some hold that you ought to maximize value within deontological constraints, for example. So properly understood, the numbers problem cuts across the consequentialist / non-consequentialist divide. Some consequentialists are searching for an explanation of why the numbers should count. Some non-consequentialists are not.

numbers should count.[8] In the thirty years since Taurek's seminal paper sparked the numbers debate,[9] several non-aggregative explanations have been proposed,[10] and have often met with forceful criticism.[11] I will not enter into the debate over these proposals here, save merely to say that arguably the floor is currently open for new contenders.

In fact, some conclude that no explanation is satisfactory, and draw an extreme conclusion from this. They reject the very claim that the numbers should count. "Numbers skeptics", as Michael Otsuka calls them, deny that you ought to save the many in Drug.[12] Taurek himself says that he would toss a coin to decide which group to help. Tyler Doggett claims that you may save the few.[13] These skeptics acknowledge that their positions are counterintuitive. Taurek anticipates people responding to his view with "incredulous tones," while Doggett expresses some hesitancy about his view on the grounds that it strikes him as "lunatic."[14] They are moved to skepticism by their dissatisfaction with the explanations that they consider. Still, most of us are reluctant to join them. We would have little reason to do so if we had a satisfactory explanation of why the numbers should count—why adopt a counterintuitive view, if we have an adequate explanation of the intuitive one?

In this paper, I hope to offer an adequate explanation, which I will call the "Ends Explanation." It does not rely on the claim that value aggregates. It does not even presuppose that someone's survival bears value. Abstracting from some complications, the

---

[8] In addition, some people are moved to reject the Aggregative Explanation because by the contractualist idea that we should act in ways that can be justified to each person. I pass over this rationale, because I am moved by Derek Parfit's criticism of it. Parfit observes that a principle could be justified to each individual on the basis of collective benefits. Parfit, "Justifiability to Each Person", *Ratio (new series)* 16 (2003): 368 - 390.

[9] Taurek, "Should the numbers count?"

[10] For proposals, see Gregory Kavka, "The Numbers Should Count," *Philosophical Studies* 36 (1979): 285 - 294; Frances Kamm, *Morality, Mortality Vol. I: Death and Whom to Save From It*, (New York: Oxford University Press, 1993); Scanlon, *What We Owe to Each Other*; Iwao Hirose, "Saving the Greater Number without Combining Claims," *Analysis* 61 (2001): 341-2; Rahul Kumar, "Contractualism on Saving the Many," *Analysis* 61 (2001): 165 - 170; Joseph Raz, "Numbers, with and without Contractualism," *Ratio (new series)* 16 (2003): 346-367; Nien-Hê Hsieh, Alan Strudler and David Wasserman, "The Numbers Problem," *Philosophy and Public Affairs* 34 (2006): 352-372.

[11] For criticism of several of the aforementioned proposals, see David Wasserman and Alan Strudler, "Can a nonconsequentialist count lives?"; Michael Otsuka, "Saving Lives, Moral Theories and the Claims of Individuals," *Philosophy & Public Affairs* 34 (2006): 109-135. The latter also criticizes several explanations that are prima facie attractive, but have not been defended in print. I leave until footnote 28 for my criticism of the recent proposal of Hsieh, Strudler and Wasserman, "The Numbers Problem."

[12] Michael Otsuka, "Skepticism about saving the greater number," *Philosophy & Public Affairs*, 32 (2004): 413-26.

[13] Tyler Doggett, "Saving the Few," *Nous*, forthcoming. There is also some plausibility to attributing this position to Elizabeth Anscombe. Anscombe, "Who is Wronged?" *The Oxford Review* no. 5 (1967) 16-17.

[14] Taurek, "Should the numbers count?" p. 306. Otsuka claims that there is another problem for the numbers skeptic. He argues that a skeptic will irrational have intransitive preferences. Otsuka, "Skepticism about saving the greater number." For replies, see Frances Kamm, "Aggregation and Two Moral Methods," *Utilitas* 17 (2005): 1-23; Kirsten Meyer, "How to be Consistent without Saving the Greater Number," *Philosophy & Public Affairs* 34 (2006): 136-146; Nien-Hê Hsieh, Alan Strudler, and David Wasserman, "Pairwise Comparison and Numbers Skepticism," *Utilitas* 19 (2007): 487-504.

two main parts of this explanation of why you should save the many in Drug can be sketched roughly as follows. First, for each stranger, you are *morally* required to make his or her survival your end. Second, you are *rationally* required to save the larger group, if you have the survival of each stranger as one of your ends.

This explanation may seem disarmingly simple. In fact, I think that this is a good thing. Why the numbers should count *is* a simple question. So it is unlikely that it has a fiendishly complicated answer. Instead, it is more likely that the reason why the numbers problem remains on the books is that a simple solution has been overlooked. But do not let the simplicity of explanation hide some of its dialectical limitations—it may not convince hard-nosed numbers skeptics, theorists who take certain views about fairness. (Nor may it convince hard-nosed egoists, for that matter.) Still, I hope it may appeal to neutrals who are looking for a solution to the numbers problem, particularly as it suggests a broader strategy of explaining why the numbers should count across the board, and not only in simple cases like Drug.

2. The ends explanation

2.1. A moral requirement to adopt altruistic ends

One of the motivations for rejecting the Aggregative Explanation of why you should save the many in Drug was the denial that the reason why you ought to help is that doing so produces value. Since I aim to offer an alternative to this explanation, I want to begin by entering the frame of mind of someone who makes this denial. What alternative rationale should they give for why you ought to save people? Let us restrict ourselves to cases like Drug where you can help people at no cost to you. It is clear that you ought to save *somebody*. Why ought you do so?

I propose that an attractive answer is that you are morally required to make the survival of each person your *end*. By this, I mean that you ought to want[15] this person's survival for its own sake so that this goal guides your decision-making and action. I take it that morality requires us to have attitudes like this, as it places requirements on what sort of people we ought to be. For example, you are morally required not to take pleasure in others' misfortune. There can be these requirements even if these attitudes are not under your voluntary control, just as there can be epistemic requirements governing your beliefs, even if these attitudes are not under your voluntary control.

Moreover, it is not just the *survival* of someone that you encounter that should be your end. More generally, if you come across someone who is in need, whom you can help at no cost, then you are morally required to make the alleviation of their need an end of yours. (To avoid unnecessary controversies, I leave open here whether or not you are morally required to aim to benefit people, beyond satisfying their needs.) There could be many possible grounds for this claim, and we need not choose between them here. A virtue ethicist might say that the virtue of beneficence requires adopting altruistic ends. A

---

[15] I am framing my proposal in terms of you *wanting* certain ends. However, nothing I will say presupposes a particular moral psychology. A Kantian may hold that you *will* your ends. Others may hold that you *intend* these ends. The proposal could just as well be put in terms of these attitudes. I require only that the attitude be one that has differing degrees of strength. For example, I require that it is possible to *will* or *intend* some ends to a greater degree than other ends. This strikes me as a harmless requirement, since any view would need to be able to make sense of the possibility that someone's goal of growing a tomato plant is less important to him or her than his or her goal of raising a child.

7

Kantian might say that you could not rationally endorse a world in which everyone has the maxim of not aiming to alleviate each other's needs. Taurek himself might simply say that you ought to empathize with the needy person's point of view, and take on their concerns.[16]

Let us proceed to ask *how much* you are morally required to want the alleviation of someone's need for its own sake and not as a means to some further end. This depends on two factors. First, it depends on how serious their need is. For example, in Drug, the degree to which you want to alleviate Anna's need should depend on the fact that her need is serious. Second, it depends on how you are related to the person in need. For example, how much you want to help Anna should depend on the fact that she is a stranger to you. So I propose the following moral requirement:

> MORALLY ADOPT ENDS. When you come across someone in need, you are morally required to adopt as your end the alleviation of his or her need. How much you want this for its own sake should depend on
> (i)   how serious their need is; and
> (ii)   how you are related to them.

If you think we need to add further conditions besides these two, then I invite you to do so. For example, feel free to let it make a difference whether the person in need is a saint or a sinner. For our purposes, let us notice that in Drug the strangers are in all relevant respects alike. They are equally in need, equally strangers to you, equally virtuous and so forth. All that will matter for the explanation that I offer is that the following two claims are true of Drug. First, you are morally required to want each stranger's survival to an equal degree. Second, how much you want each stranger to survive does not depend on whether any other stranger does.

2.2. A rational requirement on how to pursue your ends

So in Drug you ought to have certain ends. This raises the question of how you should pursue these ends. Since we are after an alternative to the Aggregative Explanation, we should not think of these ends as having value, and these values as aggregating. But what alternative guidance can be given for how you pursue these ends? I propose that this particular question falls under the broader question of how you should rationally pursue your ends in general. When you are choosing between several similar ends, what general rational requirements govern how you should choose?

It will help to answer this question by focusing on an end that has nothing to do with helping people—let us say, the protection of bird nests. (Any example of something wanted for its own sake will do. The type of end in question can be stamp-collecting, planting flowers, painting portraits, a pleasurable experience like taking a warm bath, and so on. I invite you to alter the example, as you like.)

Suppose Jones wants to protect bird nests. He has a peculiar interest in these works of avian craftsmanship. He is not protecting them for the sake of birds. He simply wants the safety of each nest for its own sake. Suppose that there are three bird nests, but Jones cannot protect all three at the same time. Instead, he has to choose between protecting one of them or the other two. Imagine he says, "Well, then. If I cannot protect all three,

---

[16] Taurek, "Should the numbers count?"

8

then I do not mind whether I protect two or just one." That would be *irrational* of Jones. The reason why is that Jones is *rationally required* to maximize the number of bird nests he protects.

Or at least he is rationally required, if we make the following three assumptions. First, Jones wants to protect each bird nest to an equal degree.[17] (For example, it is not the case that one of the nests is a weaver bird nest, which is Jones's favorite type of nest.) Second, how much Jones wants to protect each bird nest does not depend on whether he protects any other nest. (For example, it is not the case that Jones wants to protect one, but only one nest, and he does not care which one this is.) Third, Jones's choice of which nest to protect, or whether to protect any at all, does not affect his pursuit of his other ends. (For example, protecting the nests does not prevent Jones from saving his sister's life.)

These three conditions are jointly sufficient for Jones to be rationally required to maximize the number of ends he achieves. There is nothing special about bird nests, however. The preceding considerations would apply for the rational pursuit of any type of end. So I propose the following general rational requirement:

> RATIONALLY PURSUE ENDS. You are rationally required to achieve as many of a group of ends as possible if
> (i)   you want each end to the same degree;
> (ii)   how much you want each end does not depend on whether you achieve another of these ends; and
> (iii)   your doing so does not affect whether you achieve other ends that are not members of this group.[18]

You could satisfy this requirement *either* by achieving as many of these ends as possible, *or* by not having a group of ends such that conditions (i) to (iii) obtain.[19]

This general rational requirement bears on what you should do in Drug. Suppose you

---

[17] We could weaken this to the weaker condition that you want them to *roughly* the same degree. For a discussion of mild incommensurability when helping people, see Caspar Hare, "Perfectly balanced interests," *Philosophical Perspectives* 23 (2009): 165 - 176.

[18] Note that RATIONALLY PURSUE ENDS is a purely procedural principle that makes no metaethical assumptions about value. Importantly, it is not a value-maximization principle of rationality. It could be accepted by someone who denies that there are values. The principle simply holds that sometimes when you have several similar ends that you want for their own sake, you ought to achieve more of these rather than fewer. Thus, a value skeptic could appeal to it to explain why Jones should save two bird nests rather than a third, for example.

[19] That is to say that the requirement takes "wide-scope" over the entire material conditional— you are required to make the conditional true. This means you can satisfy the requirement by making the conditional's antecedent false or its consequent true. Let us denote a requirement with "O" and mark its scope with the use of brackets. We should distinguish
> (i) O (If p, then q); from
> (ii) If p, then O (q).
In (i), the requirement takes "wide-scope" over the entire conditional. Therefore, you can satisfy (i) by making p false or q true. By contrast, in (ii) the requirement take "narrow-scope" over only the consequent. Therefore, if p is in fact true, then your only way of satisfying (ii) is by making q true. I owe the general idea of a wide-scope requirement to John Broome. I doubt that Broome would accept the claims that I make here, as he holds that rational requirements only govern mental states like intentions, and not actions. See, for example, Broome, "Does rationality give us reasons?" *Philosophical Issues* 15 (2005): 321–37; Broome, *Rationality Through Reasoning*, manuscript.

meet the requirement MORALLY ADOPT ENDS: you make each stranger's survival your end, and how much you want this end depends only on the severity of their need and their relationship to you. Then you would have a group of ends such that conditions (i) to (iii) of Rationally Pursue Ends obtain. First, because the strangers are in all relevant respects alike, you want the survival of each to an equal degree. Second, because how much you want the survival of each depends on their need and your relationship to them, it does not depend on whether any of the others survive. Third, because "all else is equal" in Drug, your helping any of the strangers will not affect your pursuit of any other ends that you have (or ought to have), besides your ends of helping these strangers. Therefore, if you meet the moral requirement, MORALLY ADOPT ENDS, then you will have ends, the pursuit of which is governed by the rational requirement, RATIONALLY PURSUE ENDS.

2.3. Putting the moral and rational requirements together

So far, we have seen what ends you morally ought to have, and how you rationally ought to pursue these ends in Drug. To complete the explanation, we need to put these two parts together to form an explanation of what you ought to do.

In Drug, you face a moral requirement and a rational requirement. At the very least, you defeasibly ought all things considered to satisfy them. So when you face no other requirements you ought all things considered to satisfy both. In Drug, these are the only two requirements that bear on you. The only way to satisfy both would be to save the larger group. Therefore, the only way to satisfy all the requirements that you face is to save the larger group. Therefore, you ought all things considered to do so.

Having an explanation for why you ought *all things considered* to save the larger number in Drug would constitute a respectable solution of the numbers problem. But we could be more ambitious, and appeal to these requirements to offer an explanation of the stronger claim that you *morally* ought to save the many.

Morality does not merely require you to have the strangers' welfare as your end. Morality also requires you to perform certain actions in light of these ends. But it does not require you to perform any old actions in light of these ends. It requires you to perform actions that are rationally required in light of these ends (at least when there are no other requirements bearing on what you ought to do.)[20] So I propose:

> MORAL-RATIONAL LINK. Moral requirements to have ends are accompanied by moral requirements to perform actions that are rationally required in light of these ends.

This is how moral requirements on the inputs of your deliberation (your ends) transfer to the outputs (your choices about which action to perform).

This connection between morality and rationality is particularly well illustrated by instrumental rationality. Suppose you are morally required to have your child's welfare as your end. Instrumental rationality would require you to act in certain ways, if you had

---

[20] Let me stress that I am not claiming that morality requires you to act rationally in general, nor that all of morality is to be explained in terms of rationality. I only make the much weaker claim that a rational requirement *in conjunction with* an independent moral requirement can lead to a further moral requirement.

   Also let me stress that a wicked person who has strangers' deaths as her ends would not be rationally required to kill as many as possible. Since RATIONALLY PURSUE ENDS is a wide-scope requirement (see footnote 19), they could also meet it by giving up their nefarious ends of the strangers' deaths.

10

this end. You would be rationally required to feed and clothe your child, for example, if you had the end of your child's welfare. This is because doing so would be a necessary means to your end. The conjunction of these moral and rational requirements can explain why you are morally required to provide these basic essentials. I do not claim that this is the only possible explanation of why you are required to do so, but only that it is an attractive one.[21]

This completes my proposed explanation of why you ought to save the larger number in Drug. As I will refer to the explanation later, it will help to give it a name. I will call it the "Ends Explanation."

3. Why fairness does not require giving equal chances

I presented the Ends Explanation as assuming that in Drug you face no other requirements except a moral requirement to adopt ends and a rational requirement governing your pursuit of these ends. Not every participant to the numbers debate would accept this assumption. Some hold that you also face a requirement of fairness. For example, some hold that fairness requires tossing a coin to decide which group to save in Drug. Even if fairness did require this, this perhaps need not be fatal to the Ends Explanation—it could be that the relevant moral and rational requirements are sufficiently stringent that you ought to conform to them, rather than the requirement of fairness. Still, a case can be made for the position that there is no requirement of fairness in play.

Let us consider the ground of your duty to help someone in Drug:

> Drug. You have 100ml of a drug that you own. Two strangers, Anna and Bert, each need 50ml to live. A third stranger, Clare, needs all 100ml to live. The strangers are in all relevant respects alike. You do not stand in any special relationship to these strangers. All else is equal.

In Drug, you are acting as a private individual, distributing a resource that you own, and have to choose whom to help. This is a case where you are called on to act out of a duty of *beneficence*. I suggest that across the board, when you act beneficently, you do not have to worry about treating recipients of your beneficence fairly. To support this view, compare other cases of beneficence. When you write a cheque to Oxfam, you need not worry that doing so is unfair to the people you would help by writing a cheque to Action Against Hunger instead.

---

[21] In addition, some broader moral theories have independent reasons for positing a connection between rationality and morality. A virtue ethicist may explain what you morally ought to do in terms of what a virtuous person would do. If they hold that practical rationality is a virtue, then they will be explaining what you morally ought to do in terms of rational requirements. A Kantian may explain what you morally ought to do in terms of which worlds you could rationally endorse. You would have a duty to have altruistic ends if you could not rationally endorse a world in which people do not have altruistic ends. Further, you would have a moral duty to pursue these ends rationally, if you could not endorse a world in which people do not pursue these ends rationally.

By contrast, fairness would constrain how you distribute aid only if *justice* required you to distribute this aid.[22] Justice would be the ground of your duty to aid if the potential recipients have *claims* to this aid. By saying someone has a claim to aid, I mean that they are entitled to it, and can demand it as their due. Clearly, sometimes recipients do have claims to aid. If you had promised to give your drug to the strangers, or they jointly owned it, then they would be entitled to your helping them. It is a matter of justice that you keep your promises, and respect property rights. Similarly, if you stand in a special relationship to the recipient, such as being their doctor or parent, then they may have claims against you. In these cases, the potential recipients of your help would be entitled to your help, and it is plausible that fairness places some constraints on how you deal with people's entitlements. Of course, it would be a further step to claim that fairness requires giving people with equally strong claims equal chances, but perhaps that step could be motivated.[23]

Now, in fact, some people do say that the strangers in Drug have claims, and the strangers' claims explain why you ought to help them.[24] If we understand a claim along the lines of an entitlement, then I suggest that we should not follow them in saying this. Compare this variant case:

> Friend. You have 100ml of a drug. Two strangers each need 50ml to live.
> Your friend David needs all 100ml to live. You do not stand in any special
> relationship to the strangers. All else is equal.

---

[22] For more detailed accounts of this distinction between justice and beneficence, put in terms of whether you ought to help someone because they have a "right" to this help, see, for example, Judith Jarvis Thomson, "A Defense of Abortion," *Philosophy and Public Affairs* 1 (1971): 47-66; John Stuart Mill, *Utilitarianism*, ed. Roger Crisp, (New York: Oxford University Press, 1998, originally published 1861), chapter 5.

[23] The most detailed proposal I know of is John Broome's. On Broome's view, when people have equal claims to a good, fairness requires leaving people with an equal share of this good. If the good in question is indivisible, then fairness requires giving it to no one. But satisfying no one's claims is unacceptable on other grounds. Now, consider the option of satisfying some people's claims, and the option of holding a lottery that gives everyone an equal chance of having their claim satisfied. Broome claims that fairness requires that you take the option of the lottery rather than the option of satisfying some people's claims. He claims that the chance would be a "surrogate" for having one's claim satisfied, and that fairness would require that these chances be equally distributed. That said, Broome allows that in some cases requirements of fairness will be outweighed by other considerations: sometimes it will be more important to benefit many people, rather than benefit people fairly. Broome, "Fairness," *Proceedings of the Aristotelian Society* New Series 91 (1990 - 1991): 87-101.

Nien-Hê Hsieh, Alan Strudler and David Wasserman appear to agree that fairness requires that you equally satisfy claims of an equal strength. However, they draw different conclusions from this. They claim that you ought to satisfy as many claims as possible, because this is the best that you can do towards satisfying everyone's claims equally. Hsieh, Strudler and Wasserman, "The Numbers Problem."

[24] Conceiving of the needy as having "claims" is common in the literature on the numbers problem. See, for example, Otsuka, "Saving Lives, Moral Theories and the Claims of Individuals,"; Wasserman and Strudler, "Can a Nonconsequentialist Count Lives?"; and Hirose, "Saving the Greater Number without Combining Claims."

It may be that some people use the term innocuously, without intending that claimants of aid are entitled to this aid. However, the view that fairness plays a special role governing how you satisfy "claims" is only plausible when claims are understood as something like entitlements. Broome himself uses the term in a more robust sense, conceiving of a "claim" as a "duty owed to" the claimant. Broome, "Fairness," p. 92. He only hesitantly suggests that needs give rise to claims, p. 97.

I hope that you agree that in Friend you may help David. Furthermore, it would not be unjust of you to do so. However it is unfair to show partiality when satisfying people's claims. So if the strangers had claims on you, then it would be unfair of you to show partiality towards your friend. For example, if you had promised to each stranger that you would give them your drug, then they would have claims against you, and it would be unfair to give the drug to David. The fact that you may favor David shows that strangers in need are not entitled to your help, and that the ground of your duty to help is not justice. Instead, the ground is beneficence, as you are allowed to show partiality when practicing beneficence.

Let me rather briefly sketch my reasons for being dissatisfied with two obvious responses to this argument. The first response is that strangers in need have claims, but friends in need have stronger claims. This is unsatisfactory because if that were the case, then you would presumably be required to help David. But intuitively you are merely permitted to help him, although not required to do so. The second response is that strangers in need have claims to your help, but only when you do not have a friend nearby whom you could also help. This view of claims is unattractive, as it entails that strangers' claims would pop into and out of existence depending on whether a friend of yours is around.[25] But people's entitlements are not so fickle.

Now, my view of fairness is somewhat controversial.[26] The task of fully defending the view is too great for me to manage here. I am afraid that I am not able to satisfy champions of alternate views that these views are wrong. Here I only offer my general reasons for denying there is a requirement of fairness in a case like Drug, and hope that these might persuade neutrals.

Similarly, I do not expect that the Ends Explanation will persuade every possible numbers skeptic. A sufficiently committed numbers skeptic may well find some part of the explanation unacceptable. Still, it is worth recalling that numbers skepticism is a counterintuitive view. From a neutral's point of view, I hope the Ends Explanation is more attractive.

So the Ends Explanation may have some dialectical limitations. If a solution to the numbers problem should convince the skeptic or someone who holds that fairness requires tossing a coin, then the Ends Explanation may well fail in this regard. But to me this sets the bar too high for a solution. (Similarly, I judge that it would set the bar too high to require that a solution convinces an egoist who holds that you ought only pursue your own self-interest.) Instead, I suggest that it is enough to present an attractive explanation of why you should save the many in Drug that would persuade neutrals. This is what I hope I may have defended.

4. Why the numbers should count in more complex cases

[25] Adam Hosein makes this point in his doctoral thesis, *The Significance of Fairness*, Massachusetts Institute of Technology. In a similar vein, when considering the suggestion that people have rights to help when the help is easy, but not when it is hard, Judith Jarvis Thomson claims that "it's rather a shocking idea that anyone's rights should fade away and disappear as it gets harder and harder to accord them to him." Thomson, "A Defense of Abortion," p. 61.

[26] Some people think that fairness plays a wider role in morality, and they think that more equal distributions of goods are *ipso facto* fairer distributions. See, for example, Larry Temkin, "Egalitarianism Defended," *Ethics* 113 (2003): 764-782.

Thus far I have defended an explanation of why you should save the many in a case like Drug where strangers have the same type of needs. The numbers problem is traditionally conceived as that of finding an explanation of why the numbers should count in these cases. Finding a solution to the problem thus conceived would be an achievement in itself. But it would be an additional boon if we could extend this explanation to cases where individuals have differing needs.[27] I will make an attempt at this extension now. Consider this tidy, if somewhat gruesome, case:

> Legs (50 vs 1). You can save the right leg of each of fifty strangers or both legs of a stranger, Smith. You stand in no special relationship to the strangers. All else is equal.

My intuition is that here you ought to save the right legs of the fifty (and I invite you to increase this number, if necessary, until you share the intuition.) Presumably, this is because of the *number* of people whose right legs you would save. Why should this make a moral difference?

We cannot answer this question by appealing only to Ends Explanation, since this only addresses cases of like for like. This is because RATIONALLY PURSUE ENDS governs how you pursue ends you want to an equal degree. Still, the spirit of the Ends Explanation suggests a more programmatic proposal that would cover these more complicated cases where people have differing needs. In these cases, we should follow the same two-pronged strategy as before. First, we should consider what ends you are morally required to have. Second we should consider how you should rationally pursue these ends. I am not optimistic that in these more complicated cases we will always be able to find explicit principles like RATIONALLY PURSUE ENDS that state what rationality requires of you. Instead, the requirements governing how you pursue ends of different types may not be codifiable. But even if they are uncodifiable, we can still say something about the shape we should expect them to take. We should expect two general trends. First, these requirements will tend to take into account how much you want each end. Second, they will tend to take into account the *number* of ends. It is the latter feature that explains why the numbers should count in these more complicated cases.

Let me illustrate this point with reference to Legs (50 vs 1). You would be morally required to want to save both Smith's legs, and the right leg of each other stranger. Suppose you do. You have ends of two types.[28] Let us call the type of end of saving the right leg of a stranger "Single", and the type of end of saving two legs of a stranger, "Double." What you are rationally required to do depends on how much you want ends

---

[27] I see no way of extending some proposed explanations in the literature of why the numbers should count. For example, Nien-Hê Hsieh, Alan Strudler and David Wasserman argue that when people have claims to a good, you ought to distribute the good equally. However, in a case like Drug, you cannot distribute this good to everyone. They argue that you should give it to the many on the grounds that this is the closest that you can come to an equal division. I cannot see how this proposal would extend to cases where strangers have differing needs. If so, Hsieh et al. would need to offer a different sort of explanation of why the numbers should count in these cases. It would seem a drawback if they cannot offer a unified explanation of why the numbers should count, across the board. Pre-theoretically, we might hope for, and expect there to be, such a unified explanation. Hsieh, Strudler and Wasserman, "The Numbers Problem."

[28] I distinguish these types of ends, because I think lacking two legs is a very different way of being badly off than lacking one leg. If you do not agree, then feel free to replace "lacking one leg" and "lacking two legs", with lesser and worse harms.

of each type, and how many there are of each type. These are ends you should want to differing degrees: you should want a token of Double more than you want a token of Single. However, this should only be a *moderately* strong preference. Meanwhile, there are *many* more tokens of Single. In fact, here we would expect the sheer number of single legs to outweigh a moderate preference for ends of the type Double over those of the type Single. So you would be rationally required to save the legs of the fifty.

To buttress this claim, recall Jones, our bird nest enthusiast. Suppose he can save one weaver bird nest or fifty sparrow nests. As you may remember, the weaver bird nest is his favorite. Suppose that he only has a moderate preference for protecting these nests over sparrow nests. In light of this, we would expect the sheer number of sparrow nests to outweigh his penchant for the weaver bird nest. Rationality would require Jones to save the fifty sparrow nests.

In this manner, an explanation in terms of rationality can be sensitive to the number of ends you can achieve. Moreover, it lets the *absolute* number of ends make a difference. This is because sometimes the absolute number of ends you can achieve by performing an action will influence whether rationality requires you to perform it. For example, whether rationality requires Jones to save a weaver bird nest (his favorite!) or several sparrow nests may depend on the absolute number of sparrow nests. If Jones has a moderate preference for weaver bird nests over sparrow nests, then it may be rational for him to save fifty sparrow nests rather than one weaver bird nest, and yet rational for him to save one weaver bird nest rather than two sparrow nests. The rational requirement governing his pursuit of these ends is sensitive to the absolute number of sparrow nests: it is sensitive to whether this number is two or fifty.

This is significant for normative ethics because the absolute number of people in a group sometimes affects whether you ought to help this group. Consider a similar case to the one that we just encountered:

> Legs (2 vs 1). You can save the right leg of each of two strangers or both legs of another stranger Smith. You stand in no special relationship to the strangers. All else is equal.

I hope that you agree that here you ought to save Smith's two legs, rather than the right leg of each of the two strangers. But in Legs (50 vs 1), you ought to save the fifty strangers' right legs rather than Smith's two legs. The difference between these two cases is the absolute number of the strangers—in one case the absolute number is two, and in the other, this number is fifty. If the absolute size of the group reaches a certain threshold size (somewhere between two and fifty, and possibly a vague threshold), then you ought to

help that group rather than Smith.[29]

5. An alternative to the aggregative explanation

So rational requirements to pursue ends of different types can be sensitive to the number of ends in question. This is good news for the two-pronged strategy that appeals to moral requirements to adopt ends and rational requirements to pursue these ends. It means that this strategy can let the numbers count in cases where strangers have differing needs. But this raises a worry. One example of strangers having differing needs would be when strangers' need their headaches cured and another stranger needs his or her life saved. Does the two-pronged strategy share the Aggregative Explanation's commitment to saving some large number of headaches over a life?

It does not. This is because it appeals to the ends that you are morally required to have, and how you are rationally required to pursue these ends. We have independent motivation to hold that there are cases where someone is rationally required to achieve ends of one kind rather than any number of ends of a different kind. Recall Jones, our bird nest aficionado. Suppose Jones also has a sister, who he loves dearly. He has ends of two types—"saving his sister's life" and "protecting a bird nest." Naturally enough, Jones has a markedly stronger preference for a token of the type "saving his sister's life" over a token of the kind "protecting a bird nest." In fact, if this preference were great enough, then Jones would be rationally required to save his sister's life over protecting any number of bird nests. In this case, no one would stop for a moment to think that there must be some very great number of bird nests, such that Jones is rationally required to protect this many nests rather than save his sister's life. In general, the number of ends we can achieve can affect what is rationally required of us, but this does not entail that we will be rationally required to achieve some large number of trivial ends at the expense of one important end.[30]

Similarly, we could argue that you ought always save someone's life, rather than cure any number of headaches by making the following two claims. First, you ought to have a preference for a token of the kind "preventing a death" over a token of the kind "preventing a headache." Second, this preference should be so strong that you would be rationally required to prefer achieving a token of the former type over any number of

---

[29] This is a problem for a solution to the numbers problem that conforms to what Parfit calls an "individualist restriction" on your justification for helping people. Parfit, "Justifiability to Each Person, p. 388. Roughly, this restriction is that you may only justify your action on the basis of how your action affects individuals taken one by one. In presenting his contractualist theory, Scanlon glosses the restriction thus, "the justifiability of a moral principle depends only on *individuals'* reasons for objecting to that principle and alternatives to it." Scanlon, *What We Owe to Each Other*, p. 229. A paradigm example of an individualist explanation would be one in terms of a "pairwise comparison" of individuals' needs. By making pairwise comparisons, you could find which person(s) stands to lose the most. Some hold that if there is one unique such individual, then that you ought to help this person along with anyone else you can also help at the same time. Cases like Legs (2 vs 1) and (Legs 50 vs 1) pose a problem for an individualist explanation. By considering each individual's need in isolation, an individualist account cannot explain why the absolute size of a group would make a moral difference. It is only by considering individuals together, and not in isolation, that we can see that it makes a difference whether you can help two strangers or fifty.

[30] Indeed, this shows that the rationale for a principle like RATIONALLY PURSUE ENDS is not that each end bears some personal value for the agent, and that these values aggregate to produce a larger sum. Otherwise, the values to Jones of each bird nest would sum together to outweigh the value to him of his sister's life.

tokens of the latter type. The two-pronged strategy leaves room for making these claims. Because it does so, it does not share the Aggregative Explanation's commitment to holding that you ought to cure some large number of strangers' headaches at the expense of saving someone's life.

Dissatisfaction with this commitment of the Aggregative Explanation was one of the two motivations I considered for rejecting it. Let us also notice that the Ends Explanation is acceptable to people who reject the Aggregative Explanation because of the other motivation—the denial that the reason why you should help people in a case like Drug is that doing so *produces value*. Recall that there were three grounds for making this denial. Let me say why none are grounds for rejecting the Ends Explanation.

The first ground was an independent attraction to a theory like Virtue Ethics or Kantianism. The Ends Explanation could be seamlessly integrated into these theories. As I suggested earlier, a virtue ethicist should say that the virtue of generosity requires adopting altruistic ends, and the virtue of practical rationality requires pursuing them in a way consonant with RATIONALLY PURSUE ENDS. Alternatively, a Kantian should hold that you cannot rationally endorse a world in which people do not adopt altruistic ends, nor a world in which they do not pursue them rationally.[31]

The second ground was a desire to leave room for partiality and respect for rights. The Ends Explanation can leave room for both. A partiality theorist can hold that you are permitted to care more about an end of saving your friend rather than your end of saving a stranger. A rights theorist can hold that rights place constraints within which you pursue your ends. Of course, these theorists would still have to make a case for partiality and rights. For our purposes, let us note that the Ends Explanation leaves room for them to do so.

The third ground was a suspicion of the very property of value or "goodness, simpliciter." This worry can be dealt with quickly. The Ends Explanation simply does not posit such a property at all.

Thus, the Ends Explanation should be acceptable to people with any of the motivations for rejecting the Aggregative Explanation. It is an alternative explanation of why the numbers should count that is an alternative to thinking that you should bring about a state of affairs with the greatest possible value, where this value is an aggregative sum of the values of the state's parts.

6. Conclusion

So the Ends Explanation is an alternative to the Aggregative Explanation. Indeed, it is an alternative that can appeal to moral theorists of many stripes. That is not to say it will find favor with absolutely everyone—as I noted earlier, it may not convince certain numbers skeptics, certain theorists with strong views about fairness, or an egoist for that

---

[31] It is not just nonconsequentialists who could adopt the Ends Explanation. Suppose someone holds that each person's good bears some value, by itself. However, they deny that value is born by states of affairs that are composed of more than one individual's good. They say that such a state of affairs has parts that bear value, but the state of affairs itself does not bear value. In Drug, they would say that Anna's survival bears value, and Bert's survival bears value, but the state of affairs in which both survive does not contain value. Their position is plausibly a consequentialist position, and it could adopt the Ends Explanation. Because the survival of each bears value, they could say you ought to adopt this as your end. They could then appeal to RATIONALLY PURSUE ENDS in order to explain how you should pursue these ends.

matter. Even so, I hope that it is an attractive proposal that is worthy of the consideration of people who are looking for an explanation of why you should save the many in Drug, an explanation which would count as a solution to the numbers problem as traditionally conceived. Moreover, it suggests a more programmatic strategy for explaining why the numbers should count when people have differing needs—the two-pronged strategy of appealing to moral requirements to have ends, and rational requirements to pursue these ends. This is all to the good. After all, we do not merely want to understand why the numbers should count in simple cases like Drug; we want to understand their moral importance in general.

Chapter Two

Altruism and Ambition
in the Dynamic Moral Life

1. An ambition for altruism

Some people are *ambitious* about helping people in need. Even though they already give an impressive amount of their time, money and effort to help them, they are not content to give just this much. They strive to give more. I think that typically they are right to be ambitious. And, I think that this bears on what beneficence requires of us.

1.1. A conflict for altruists

Maybe you know someone like Amy. Amy now works in a non-profit organization that aims to persuade pharmaceutical companies to waive intellectual property rights for new vaccines in the developing world. She used to work as a corporate lawyer, but grew disillusioned about the contribution her career was making to the world. Each year, she sets aside a generous portion of her salary towards providing healthcare in developing countries. When a specific disaster, like an earthquake, strikes, she makes a further donation. There is more to her life than curing the world's ills, though. She is a supportive and loving mother and wife. She has her own interests besides. She loves to practice jazz on the piano, and to eat out. She expresses her aesthetic sense in the way she dresses and looks forward to holidays in interesting places as much as the rest of us do.

Amy is not a saint, but she is still doing pretty well when it comes to altruism. Even so, she tries to do better: she aims to increase how much time, effort and money that she spends helping people in need. Amy is an *ambitious* altruist. Just as some of her former colleagues strive to make partner and beyond, Amy strives to help others more. But her ambition is different from theirs in key respects. Amy's ambition is not for her own personal success, nor is it grounded in her personal preferences. She thinks that regardless of what she wants, she should do more to help others. What drives her ambition is her conscience. She feels morally pressured to help others more. Her experience is that she feels morally obligated to do so.

Interestingly, her ambition does not expire when she makes adjustments to her lifestyle in light of it. When she changed jobs, she still set herself the goal of doing more in her career. When she increased her yearly donation, she felt pleased with this effort, but still aimed for an even larger amount in the future. Amy's previous improvements do not take away the need for future improvements. As a result, her ambition is practically insatiable.

While Amy's ambition motivates her, it can also lead her to self-reproach. Sometimes, when she reflects on the fact that she should be giving more, she infers that she is not giving enough. From this, she concludes that she is presently failing to do what she is required to. This makes her feel guilty. Not always, though. At other times, she thinks that she *is* allowed to have a life of her own, she *is* doing an impressive amount to help others, and she *is* trying to do even more. How much does the world want of her? So Amy's ambition makes her feel conflicted. Sometimes, she thinks she is doing all that is required of her; sometimes, she thinks she is required to give more.

This conflict will be the topic of this paper. I feel moved by both sides to Amy's conflict when I consider the question of how much of our resources we are required to give up to help others in need. (By "resources," I mean our time, money, effort, and whatever else we value but could sacrifice to help others.) That is, I am inclined to accept both of these claims:

AMBITION   Beneficence requires an impressive altruist like
                      Amy to be ambitious, and aim to give more than
                      she is currently giving.

MODERATION   Beneficence does not require so much that an
                      impressive altruist like Amy is failing to give
                      enough.

Although each is attractive, they seem to contradict each other. My project will be to explore how far it is possible to accommodate both.

This project will be of some interest, (I hope), because I am far from alone in feeling the force of Amy's conflict. But given the philosophical disagreement over the issue of the requirements of beneficence, I predict that the assumptions behind the project will seem misguided to some. The desire to accommodate AMBITION will be lost on those who are attracted to a minimal view of beneficence, according to which we are permitted to concentrate predominantly on making the most of our own short lives, even while others are in great need. Meanwhile, those who are attracted to a more demanding view of beneficence, according to which we must concentrate mainly on alleviating others' needs, may well not see any reason to accommodate MODERATION. I certainly cannot resolve this broader debate here. Still, in light of the controversy, I feel that I ought to say a little about why it is that some of us are drawn to these assumptions. I do so in order that I may illuminate to opponents why we take the position we do, even if I cannot persuade them that the position is correct.

I will start with AMBITION, about which I have less to say. Some of us are inclined to share Amy's conviction that she is required to be ambitious, and not settle merely for an impressive amount of altruism. We think that Amy is required to make sacrifices herself to help people who lack the essentials for a minimally tolerable life—food, water, shelter, security from violence, education and employment opportunities. In addition, we think that she is right to see the severity of others' needs as a call to arms. She is right to be averse to resting content with her current sacrifices. The current state of much of the world is a perpetual emergency, and in the face of this urgency, a continual striving towards giving more to help others is requisite; it is incumbent on us to try to get by with less ourselves, so that we may help others more. Let me stress that AMBITION denies that Amy is going above and beyond in having her goals. That is, it is not the claim that this is supererogatory of Amy— good of her, but not required. Rather, the claim is that she is

morally required to do so.[32] She is required to engage with the full extent of others' suffering, and properly commit herself to alleviating them. Doing so involves an ambition for altruism. If someone asks for deeper reasons that explain why this is so, I doubt there is much more I can say besides reiterating the seriousness of others' needs, and the sheer numbers of people who have these needs. I do not know how to advance the debate with people who do not agree that these needs place these constraints on how we live our lives. But I hope that enough people are of a like mind to make the ensuing discussion worthwhile.

So much for AMBITION. I can offer more in way of an explanation of MODERATION, and in the following sub-section, I want to spend longer discussing why we should accept this claim. Indeed, I have a particular interest in getting this explanation straight because it will bear on much of my subsequent discussion. I should reiterate that these remarks will be too brief to count as anything like decisive arguments against alternative positions, and there is much a proponent of these alternative accounts may say in response. I mention them merely in order to frame my project in this paper.

1.2. Two limits to beneficence: not too difficult, and not suffocating

Why are some of us drawn to MODERATION? We think that there are limits to the demands of beneficence because we are allowed a robust concern with how our own lives go. At the core of this thought is the intuition that (at least typically) a good life is a moral one. That is, we could live good lives, and meet the requirements of morality. We think that this consideration places a limit on how much beneficence requires of us, on the grounds that a life dominated by altruism would not be a good life.

To flesh this point out, we should ask what a good life involves. For a start, it seems false that a good life is one in which we are slaving away constantly in pursuit of worthy ends. In this vein, consider Barbara Herman's attractive claim that

> a reasonable morality is well integrated into ordinary living, not something we are endlessly at war with (like a diet) or a distant goal toward which we direct substantial amounts of our energy.[33]

Here, Herman directly asserts that morality is not excessively onerous. I suggest that an attractive explanation of this claim is that a "reasonable morality" allows us to live good lives, and that good lives are "well integrated into ordinary living."

This general point about morality is relevant to the requirements of beneficence. For, even an impressive altruist like Amy would be "endlessly at war with" a life fashioned only out of altruism. She has to find a balance between living her own life and improving the

---

[32] By saying that she is morally required to do so, I mean that she would be failing, morally, if she did not. Moreover, I also mean that she is all things considered required to do so. To my mind, there is no conflict between morality and other spheres of our lives, so that whenever someone is morally required to do so, they are also all things required to do so. Admittedly, this point is controversial, and nothing I will say turns on it. I invite people who do not agree with me on this point to understand me as meaning "all things required to do so in virtue of moral reasons" by "morally required to do so." (In contrast to an action which is all things required, but not in virtue of moral reasons.) For stylistic reasons, I sometimes put things in terms of what she "ought" to do, or "should" do, but I am always interested in terms of what she is all things required to do in virtue of moral reasons.

[33] Barbara Herman, *Moral Literacy*, (Cambridge MA: Harvard University Press, 2007), p. 108.

lives of others, and this balance should not be so far in altruism's favor that she is miserable herself. Now, it may be that there are some exceptionally saintly individuals for whom a life of pure altruism is not excessively difficult. But for most of us, a life of pure altruism would not be a good life. In light of this, it seems counterintuitive that beneficence requires a life of pure altruism.

> TOLERABLE   Beneficence does not require so much of us that
> DIFFICULTY   meeting its requirements is excessively difficult
>   for most of us.[34]

Here, I leave it up for grabs what counts as *excessively* difficult. Where to set this limit is a matter for further debate that I cannot, and we need not, settle here.

So this first limit to beneficence allows us to live tolerable lives. But more can be said for our self-interest. For most of us, a good life involves pursuits beyond helping people. Susan Wolf points out that

> if the moral saint is devoting all his time to feeding the hungry or healing
> the sick or raising money for Oxfam, then necessarily he is not reading
> Victorian novels, playing the oboe, or improving his backhand.[35]

Wolf argues that a good life would include some pursuits like these, and she concludes that we ought not be moral saints. Many of Wolf's examples of worthwhile pursuits include enjoying and producing culture. But her view is far from elitist: Wolf values an appreciation of the Boston Celtics and the ice cream sundae, as much as that of the piano concerto, and the canard en croûte. Her aim is to defend "the enjoyment of material comforts," engaging in "the intellectual and physical activities of our choice" and "the love, respect and companionship" of others.[36] As I will put it, these are all examples of worthwhile "interests" that we may have beyond altruism.

---

[34] I put this thesis in terms of what is excessively difficult for "most of us" in order to avoid taking an unnecessarily controversial stance on who should count as the standard against which "excessive difficulty" is measured.

    One option would be to relativize what beneficence requires of each individual to what each individual finds difficult. Thus, Jones would not be required to give so much that this is excessively difficult for Jones. The drawback of this option is that it requires little of defective people, such as particularly selfish people, or people who find everything difficult. Suppose Jones is such a delicate flower that he finds even the tiniest amount of altruism very difficult. On the current proposal, then it will not require that he gives very much at all. But it seems strange that Jones's obligations are so limited in virtue of his being defective.

    Another option would be for this constraint to be sensitive to what would be excessively difficult for a non-defective person. This would create a baseline for everyone—delicate Jones would only be excused what would be excessively difficult for a normal person. However, this proposal has the result that particularly virtuous people are only constrained by what would be difficult for a non-defective person. Suppose Smith is lucky enough that in a certain situation giving $1000 would no trouble for him. It would be strange that Smith is not required to do so simply because a less virtuous person, (the merely non-defective person), would find it excessively difficult to give more than $500—why should the difficulty that this presents for others limit what Smith must give?

    A third option combines elements of each. On this hybrid view, beneficence does not require so much of someone that this amount would *both* be excessively difficult for them *and* excessively difficult for a non-defective person. This hybrid approach would offer attractive extensional results. There remains, however, the question of whether it can find enough independent motivation that it is not objectionably ad hoc.

[35] Susan Wolf, "Moral Saints," *The Journal of Philosophy*, 79 (1982): 419–439, at p. 419.

[36] Wolf, "Moral Saints," p. 420.

Wolf's position is attractive to many of us. We do think that a good life will have more to it than altruism.[37] As I will put it, beneficence should not be so *suffocating* that it leaves us little room for other worthwhile interests:

> NO SUFFOCATION   Beneficence does not require so much of us that our altruism is so suffocating that it leaves us insufficient room for other worthwhile interests.

Again, there are further controversies about this claim—how to balance these interests against others' needs?[38] How much weight to give to the consideration that someone identifies with a particular consideration?[39] I do not have much helpful to say on these matters, and fortunately, how we resolve these issues will not matter for our topic in hand.

1.3. The tension between moderation and ambition

So we have come across two constraints on how much beneficence can demand of us. In light of Herman's criticism of a diet-like morality, many of us hold that altruism should not be excessively *difficult*. In light of Wolf's defense of the worth of other projects, many of us hold that our altruism should not be excessively *suffocating*. Each of these constraints is attractive to many of us in its own right. Together, they illuminate how it could be the case that Amy is giving enough to satisfy the requirements of beneficence, even though she is not a moral saint:

---

[37] Wolf seems to think that no flourishing life would be chiefly dominated by altruism. For example, she begins her paper by stating that "I am glad that neither I nor those about whom I care most are [moral saints]," "Moral Saints," p. 419. For a response that defends the lives of some actual saints, see Robert Adams, "Saints," *The Journal of Philosophy* 81 (1984): 392 – 401.

[38] Wolf's own view seems to be that no weighting is possible because these values are incommensurable with moral values. This is a controversial position to take, and challenged by Sarah Buss in her "Needs (Someone Else's), Projects (My Own) and Reasons," *The Journal of Philosophy* 103 (2006): 373 – 402. Buss argues that even if two values are irreducibly distinct, it can still be the case that we are unjustified in promoting one value over another in a particular situation.

[39] Many of us think that beneficence should not crowd out other worthwhile interests. But *whose* worthwhile interests we might ask. The agent's own interests? Or, other worthwhile interests that appeal to others? Say that Marks has a particular love of playing golf competitively. He is so passionate that he would feel dejected not pursuing it (and let us suppose that he would also feel dejected playing it at a recreational level.) Now, the fact that the worthwhile activity is central to Marks's identity does not mean that it is above ethical scrutiny. As Buss has argued in her "Needs, Projects, Reasons', we can still ask whether we can justify our most central interests. But still, we can ask how accommodating beneficence is of Marks's particular passion.

Moreover, suppose that to enjoy his passion, Marks needs to practice assiduously, so his passion is particularly time-consuming. Let us suppose his sport is particularly costly (he needs the right golf clubs, membership to the right greens and so forth.) As a result, his sport leaves him few disposable resources to divert to helping those in need. Now, in their own right, playing golf, and developing one's skill at it, are worthwhile activities. But there are cheaper, and less time-consuming alternatives, which other people find just as fulfilling. Moreover, other activities do more good for people in need. Others derive as much satisfaction form them as Marks does from golf. How ethically relevant is the fact that Marks has the interests that he does, and could not easily come to have less expensive ones?

This question is difficult because two extremes are unattractive (at least to many of us). First, it seems that individuality should get some due, so that beneficence does not require us all to aspire to the most beneficial, and least expensive, worthwhile activities. Second, it seems that personal preference does not reign supreme—some tastes are too expensive, and do too little good, to be justified. For example, people like to charter private helicopters at the cost of several thousand dollars, in order to reach the peaks of mountains that they then snowboard down. It is hard to navigate between these two extremes.

MODERATION   Beneficence does not require so much that an
                              impressive altruist like Amy is failing to give
                              enough.

However, adopting a moderate view of beneficence makes it hard to accommodate the claim that Amy ought to be ambitious. For if Amy is doing enough, then how could it be that she is required to aim for more?

This difficulty besets, for example, a moderate view according to which we are only required to give away some quota of our resources. This could be a universal quota—J. O. Urmson holds that the beneficence that is required of each of us is that which can be demanded of everyone.[40] Or, it could be a quota that varies with individuals' circumstances—Liam Murphy holds that what is required of each of us is our fair share of a collective altruistic effort, and that each of our fair shares may vary.[41] For our purposes, it is enough that the quota not be excessively difficult, and that it leaves room for other worthwhile interests besides altruism; this would be enough to make it a moderate view. Can a quota account accommodate a requirement for Amy to be ambitious? Not straightforwardly. Amy would only be required to give more than she is if she was currently giving below the minimum. But if Amy is giving below the minimum, then she is failing to give as much as she is required to give. By contrast, suppose she is already giving more than this minimum. Then Amy would be mistaken to feel obligated to give any more. Instead, what this account will say about Amy is that it would be "supererogatory" for her to strive to give more than the minimum—it would be good for her to do so, but optional.

My reason for mentioning quota views is not to dismiss them. Indeed, the positive proposal that I defend is arguably consistent with a certain type of quota view. Rather I mention them as examples of a more general danger we face if we adopt a moderate view: we are under pressure to deny that we are required to be ambitious.

Moreover, there is pressure in the other direction too. It seems that if we accept that we ought to be ambitious, then we are under pressure to embrace a demanding view of beneficence. A demanding view could be formulated in many different ways and held by

---

[40] J. O. Urmson, "Saints and Heroes," p. 64, in *Essays in Moral Philosophy*, ed. A. I. Melden, (University of Washington Press, 1958). I read Urmson as implying that our duty to act beneficently goes no further than actions that have trivial costs for ourselves. Urmson holds this view because his paradigm of a duty is something we can reasonably *expect* and *demand* from each other. On his view, what we can expect of each other is adherence to an "ideal moral code" composed of "simple rules... binding on all." (p. 70). These are rules that we can enforce by applying "pressure" on each other, by means of "censure in the event of non-performance" (p. 71). It is in terms of these rules that Urmson marks off a category of actions that are morally good, but not required.

I think that Urmson's conception of a "duty" may be useful for the question of how we ought to *sanction* wrongdoers. (Should we punish them? Express our disapproval? Privately think poorly of them?). However, I do not think it bears on what we are morally required to do. For example, we are required to be ambitious, even if others may do little to censure us if we fail to do so.

[41] Liam Murphy, *Moral Demands in Nonideal Theory*, (Oxford: Oxford University Press, 2003). To be clear, Murphy does not motivate his view on the grounds that it avoids making beneficence too demanding. But some may find it a virtue of his approach that when there are sufficiently many other people that could help, someone's fair share would not be an overly demanding amount.

a rainbow of moral theorists.[42] I will focus on a justly famous example—Peter Singer's view. Singer holds that we are required to give away anything that does not satisfy an important need of our own to help people in need.[43] Singer's view easily accommodates the claim that Amy is required to be ambitious about how much she is giving—she is required to give more than she is presently managing, because she is spending resources to enrich her own life without satisfying an important need of hers. Indeed, if we accept that Amy ought to be ambitious, then we may easily be led to such a view. For, if she is giving a moderate amount, and this is not enough, then presumably she should be giving a more demanding amount.

What Singer's view cannot accommodate is MODERATION: a demanding view will hold that Amy is correct when she reproaches herself for not giving enough. But many of us think she is making a mistake here. Overall, we think that she is doing pretty well when it comes to helping others. In light of this, a demanding is view unattractive. For, if she really were failing to do what she ought, then she would be failing morally. If she were failing morally, then our overall evaluation of her should be a negative one. It could still be true that she is *less* blameworthy than most, and if we only focus on what she gives, then we praise her for this. But when we focus on her entire lifestyle, then we should blame her. This is what seems wrong to many of us.

1.4. A summary of the project

To sum, up, I am assuming that it would be desirable for an account of beneficence to accommodate the following claims:

---

[42] Common versions of consequentialism entails a demanding view: according to many such theories, a great deal of the time you would bring about a better outcome by giving away resources to help others rather than spending them on yourself.

According to some versions of virtue ethics, you ought to act as the virtuous person would. If a saint like Mother Theresa is a paradigmatic virtuous person, then this might amount to giving away rather a lot.

Lastly, some interpret Kant as holding a demanding view. This is a somewhat controversial interpretation, but it is suggested by at least three textual points. First, Kant holds that someone is to try, "as far as he can, to further the ends of others." [4:430] of Immanuel Kant, *Groundwork of the Metaphysics of Morals*, translated by Mary Gregor, (New York: Cambridge University Press, 1998), p. 39. See also [402] Immanuel Kant, *The Metaphysics of Morals*, translated by Mary Gregor, (New York: Cambridge University Press, 1991), p.203, and [453] of the same, p. 247. Second, he omits to defend a *permission* to pursue our own happiness, beyond our "true needs" [393], *The Metaphysics of Morals*, p.197. Third, he claims that the only thing that limits our duty of beneficence is another duty, such as a duty of gratitude to our parents: [390] of *The Metaphysics of Morals*, p. 194. David Cummiskey defends an interpretation of Kant as holding a demanding view in his *Kantian Consequentialism*, (New York: Oxford University Press, 1996). Thomas E. Hill Jr argues for a more moderate interpretation in his *Human Welfare and Moral Worth*, (New York: Oxford University Press, 2002), Chapter 6, "Meeting Needs and Doing Favors."

[43] Peter Singer, "Famine, Affluence and Morality," *Philosophy and Public Affairs*, 1 (1972): 229 - 243. Singer first states his principle in terms of sacrificing anything that is not "morally significant" (p. 231). He later elaborates that something does not count as morally significant if it does not address an "important need" (p. 235). Here, Singer aims to construct an argument that does not presuppose consequentialism, and is based on common-sense principles. In fact, Singer endorses the stronger claim that you are required to distribute your resources wherever they will have the greatest marginal benefit, whether this means your giving up something "morally significant" or not.

| AMBITION | Beneficence requires an impressive altruist like Amy to be ambitious, and aim to give more than she is currently giving. |
| MODERATION | Beneficence does not require so much that an impressive altruist like Amy is failing to give enough. |

Moreover, as a rationale for MODERATION, I am assuming:

| NO SUFFOCATION | Beneficence does not require so much of us that our altruism is so suffocating that it leaves us insufficient room for other worthwhile interests. |
| TOLERABLE DIFFICULTY | Beneficence does not require so much of us that meeting its requirements is excessively difficult for most of us. |

My project is to explore how an account of beneficence could accommodate all of these claims. The difficulty in doing so arises because the conjunction of MODERATION and AMBITION seems paradoxical, if not contradictory. For, if Amy is only required to manage a moderate amount of ambition, how could she be required to be ambitious? How does a practically insatiable ambition for altruism leave room for other interests, for example? We would need to find a way of reining in such ambition, without it leading to an excessively demanding amount of altruism. The puzzle is to work out how an account of beneficence can accommodate all of the above claims at once. That is, my project in this paper is to frame a moderate account that respects both limits to beneficence, but still manages makes an ambition for altruism mandatory.

2. Dynamic beneficence as a solution to the puzzle

Here is my diagnosis of why it can be hard to find a view of beneficence that accommodates both moderation and ambition. This difficulty arises from thinking about the issue *statically*. If we fix on a particular time, and judge that someone is required to give more, then it must follow that they are not giving enough. However, there is a different way we can think about beneficence. Instead of fixing on a particular time, we can think about beneficence *dynamically*. We can consider our altruism in the context of our lives, and ask what pattern this altruism is required to take over time.[44]

How could this help with the matter in hand? We could distinguish what we are required to do *now* from what we are required to do in the *future*. When someone like Amy asks herself the common question, "am I giving enough?", we would press her to make clear whether she is asking if she is giving enough at the moment, or whether she should be aiming to still give this amount in the future. By separating these questions, we could claim that someone like Amy is required to give more in the future, even though she

---

[44] Note that I am not claiming that thinking about things dynamically motivates a moderate view. A demanding view of beneficence could also hold that the amount we are required to give varies over time—it could simply hold that it is always a demanding account. Rather, my point is to show how considering the dynamism of beneficence frees up some theoretical space for a moderate account to incorporate a requirement to be ambitious.

is giving all that is required of her now. Moreover, we could hold that Amy is required to increase how much she gives *gradually*: she is not required to make a large increase all at once. If we say this, then we could have it both ways: we could have a moderate view of beneficence that also includes a requirement to be ambitious. Of course, all this depends on our having a principled way of making these claims. Now, I will try to offer one.

2.1. Dynamic beneficence: fueled by a requirement to develop morally

When we look at things dynamically, we need not take our characters nor our interests as fixed. Over time, we can change them. I will argue that we are required to develop morally into more altruistic people, in order to bring our altruism and self-interest closer to harmony. This will tend to increase how much we should give. I propose the following two-part thesis:

DEVELOPMENT   (i)   Beneficence requires us to develop morally;

                         (ii)   The more morally developed we are, the more beneficence requires us to give.

The requirement in the second part of the thesis should be understood as having an implicit all things equal clause. In other words, it claims that when we develop, beneficence requires us to give more than we would be required to give if we had not developed. So suppose that Amy develops morally but loses her job. Beneficence requires unemployed Amy to give more than she would have been required to give if she was unemployed and less morally developed. But it need not be the case that beneficence requires unemployed and more mature Amy to give more than she was required to give when she was employed and less morally developed. DEVELOPMENT simply claims that holding fixed other relevant factors, such as employment, the more morally developed we are, the more beneficence requires us to give.

If you like, you could think of these two requirements in terms of the following analogy.[45] Consider someone who is learning the high jump. At any particular jump, they should aim to clear the bar. But they should also aim to improve over time. As they do improve, they should raise the bar higher. Similarly, at any particular time, beneficence requires us to give a certain amount to help others. It also requires us to develop morally. As we become better people, we should increase what this amount is.

DEVELOPMENT makes an ambition for altruism appropriate. Let us return to Amy. According to this thesis, when all else is equal, Amy is required to develop over time, and then increase how much she gives. So looking forward into the future, she ought to be ambitious about how much she gives—she ought to be aiming to give more then. Importantly, this need not mean that Amy ought to be giving more now.

Furthermore, I will argue that we can get this result, without encountering any of the problems we saw earlier. Let us take seriously the idea that beneficence cannot require Amy to make her altruism too difficult nor suffocating. So she is only required to develop gradually, and at a pace that is not excessively difficult nor suffocating. Moreover, the end result of her moral development would be an ideal that is not too difficult, nor suffocating.

---

[45] Thanks to Julia Markovits and Caspar Hare for suggesting and refining this analogy, respectively.

In this manner, we can hold a moderate view of beneficence, while also finding a role for ambition.

Now, I will offer an argument for why the thesis DEVELOPMENT is true.

2.2. The argument for the requirement to develop morally

My argument for DEVELOPMENT is simple. Earlier, I noted that there are two constraints on how much beneficence can require us to give up to help others. First, our altruism should not be too *difficult*. Second, we also have legitimate interests besides altruism—our altruism should not be *suffocating*. Suppose we could help others more without this increase making our altruism too difficult nor suffocating. Further, suppose there are no further reasons against doing so. For example, doing so does not affect how we meet our other obligations. In such a case, others' needs count in favor of us being required to help them, and there are no counterbalancing considerations. So beneficence requires us to do so. In light of this, the first premise in my argument is that beneficence requires us to pursue a course of action if it allows us to help others more without this being excessively difficult or suffocating for us.

> INCREASE If you can give more to help others without this being too difficult or suffocating, (and there are no other reasons not to), then beneficence requires you to do so.

The clause specifying that "there are no other reasons not to" is in place because beneficence clearly does not require you, for example, to steal from others in order to give more.

Shortly, I will argue that moral development allows us to increase how much we give to help others:

> MEANS Moral development is a means to giving more to help others without this being too difficult or suffocating.

This will be the second premise in my argument. I take it that it is uncontroversial to assume that there are no further reasons not to develop morally and give more than one would have given as a less mature person. On this assumption, the first premise, INCREASE, and the second premise, MEANS, together entail DEVELOPMENT:

> DEVELOPMENT (i) Beneficence requires us to develop morally;
>
> (ii) The more morally developed we are, the more beneficence requires us to give.

Soon, I will turn to some key forms of "moral development," and thereby defend MEANS. Before I do so, I would like to make a brief point about the type of rationale for moral development that this is.

The argument concludes in a moral requirement to develop morally: beneficence requires us to do so. Thus, moral development is required as a *means* of helping others. I think of this rationale as having a Kantian flavor. On Kant's view, only an act performed from a sense of duty has moral worth. He claims that we ought to develop morally in

28

order that we are better able to act from a motive of duty. As Marcia Baron puts it, Kant posits "a very central, even foundational, duty: the imperfect duty we each have to strive to perfect ourselves morally."[46] So on Kant's view we develop morally in order to have the right motivation when we act. The current proposal is similar to the Kantian view in that it is also a purely *instrumentalist* rationale for moral development. While Kant espouses moral development for the goal of having the right motivations, my argument espouses it for the goal of aiding others.

By contrast, there is another possible rationale for moral development: we ought to develop for our own sakes. This non-instrumentalist rationale has an Aristotelian flavor. On Aristotle's view, we ought to develop because virtuous states of character are ends that we ought to want for their own sake, as well as being means to our ultimate end of flourishing (eudaimonia).[47] Our reasons to develop ourselves, on Aristotle's view, are ultimately grounded in our reasons to value our self-perfection for its own sake, and as constitutive of our flourishing.

These types of rationale for moral development are different. The first proposes developing for the sake of doing our duty; the second proposes developing for our own sakes. Although these rationales are different, they need not be competitors. It could be that we have multiple reasons to develop. Here, I only insist that a requirement of beneficence is one of these. To make good on this claim, I need to defend MEANS, which is the second premise of the main argument for DEVELOPMENT. I turn to this task next.

3. How moral development allows for more altruism within the two limits to beneficence

To show how moral development can enable altruism, I will focus on some particular kinds of moral development. When possible, I will aim to support what I say with reference to the relevant empirical literature. First, I will look at how *habituation* and *strengthening our capacities for will-power* lets us give more to help others, without this being excessively difficult. Second, I will look at how *shaping our interests* can let us be more altruistic without this being so suffocating that we do not have other interests. It is in light of these points that I think we should accept the claim of the preceding section:

MEANS   Moral development is a means to giving more to help others without this being too difficult or suffocating.

This would complete my defense of the argument for DEVELOPMENT, which illuminates why we are required to be ambitious.

3.1. Giving more within the first limit: habituation and building will-power to make altruism less difficult

---

[46] Marcia Baron, "Kantian Ethics and Supererogation," *The Journal of Philosophy* 84 (1987): 237-262 at p. 249. The position that Baron attributes to Kant is an excellent example of a dynamic view that takes moral development seriously, and the proposal of this paper has profited a great deal from her paper.

[47] Aristotle, *Nichomachean Ethics*, translated by Terence Irwin, (Indianapolis: Hackett Publishing Company, Inc., 1999)

I will begin by showing how moral development can allow us to give more while there is the first limit to beneficence encountered earlier:

> TOLERABLE DIFFICULTY   Beneficence does not require so much of us that meeting its requirements is excessively difficult for most of us.

I will do so by looking at two forms of moral development that make altruism less difficult. First, I will look at habituation. Second, I will look at building will-power.

The first way in which we can make altruism less difficult is by forming appropriate habits. Habituation's role in building character is famously an insight of Aristotle, who claims that

> what we do in our dealings with other people makes some of us just, some unjust; what we do in terrifying situations, and the habits of fear or confidence that we acquire, make some of us brave and others cowardly. The same is true of situations involving appetites and anger... To sum it up in a single account: a state of character results from the repetition of similar activities.[48]

On Aristotle's view, by repeating actions we engrain dispositions to perform certain actions and feel certain emotions in certain situations.[49]

In support of this view, there is a picture of habituation emerging from the contemporary psychology literature about how we come to automaticize actions.[50] By repeating certain actions in response to certain cues (often environmental cues), we come to make these responses *automatic*. When we do so, we no longer have to engage in conscious, effortful decision-making in order to choose to perform these actions. Instead, the responses are adopted as the result of through-put mechanisms that are triggered directly by these cues.

What dispositions would it be helpful to form? The most obvious are habits of helping others. Equally important are habits of paying attention to the needs of others.[51] Psychologists such as Daniel Batson defend the "empathy-altruism hypothesis": our

---

[48] Book 2 Chapter 1, [1103b14 – 1103b22], of Aristotle, *Nicomachean Ethics*, p. 19.

[49] Aristotle's view of character as a stable and general disposition is controversial, and disputed by some philosophers in light of claims by situationist psychologists. Situationists claim that some situations trigger the same behavior in different individuals. Some philosophers object that this shows people do not have such general and stable dispositions. (For example, see John Doris, *Lack of Character: Personality and Moral Behavior*, (New York: Cambridge University Press, 2002)). However, we do not need to adopt Aristotle's view of character to profit from his insight about habit. It is enough to note that we do form behavioral and emotional dispositions to a considerable extent.

[50] For recent work on automaticity, see Ran Hassin, James Uleman and John Bargh (eds.), *The New Unconscious*, (New York: Oxford University Press, 2005).

[51] The role for empathy may be nuanced, however. It seems that empathy could add to the difficulty of altruism. If someone is constantly paying attention to others' suffering, and sharing in it, then they will be failing to have a reasonably flourishing life themselves. It is not clear to me where to strike the appropriate balance.

empathizing with others motivates our altruistic behavior.[52] Our ability to empathize is increased the more we know about others' situations, and the more we pay attention to their plights. Forming the habits of inquiring into others' plights and of paying attention to them would be a form of moral development.

Both of these benefits are exemplified by the case of Mary. Mary was raised in a religious community, which had the following practice. At each celebration, such as a birthday, the children are all provided with a penny that they symbolically put into a collecting can for a charity. Mary has since cut her ties with the community, but has kept up the practice of giving symbolically. Whenever she passes by someone asking for money, whether it be someone collecting for charity or a homeless person, she always gives *something*, even if it is a small amount. Her boyfriend has suggested that her money would go further if she donated only to a particularly efficient charity. Mary defends her habit on the grounds that ensures that she gives regularly and that it keeps her open to the needs of others, stopping her from turning a blind eye when her help is called for. She aims for her habit to be a means to building her own character, as well as a means of helping others.

Acquiring behavioral and affective dispositions is one part of forming character. Another part is developing our capacities for will-power. Kant puts this at the heart of his account of moral development. He conceives of developing morally as strengthening our "capacity... to overcome all opposing sensible desires." In more prosaic terms, we might refer to this as our ability to resist temptations. Kant also singles out a key means to strengthening our will-power. He claims that a central way to strengthen our capacities for will-power is "by practicing virtue."[53]

Recent work in psychology lends support to Kant's insight. Roy Baumeister has argued that a faculty of self-control is similar to a muscle.[54] Like a muscle, it uses a limited energy resource. And, like a muscle, it is something that we can strengthen over time. Just as a gym-rat can become more physically powerful by using her muscles to lift weights, so someone can increase their will-power by using it to resist temptation.

Why would helpful habits and strong will-power facilitate giving more to help others? Partly, they make it more likely that we perform helpful actions, pay attention to others' needs, and shape our interests. In addition, and equally importantly for our purposes, they would make altruism easier. If decisions are removed from the realm of conscious decision-making and made habitual, then these decisions are not occurrently experienced

---

[52] C. Daniel Batson and Laura L. Shaw, "Evidence of Altruism: towards a pluralism of pro-social motives," *Psychological Inquiry* 2 (1991): 107-122.

[53] Section [397] of Kant, *The Metaphysics of Morals*, p. 200. In the same section, Kant mentions another way of strengthening will-power. This is "by contemplating the dignity of the pure rational law in us." Perhaps a non-Kantian way of putting this point would be "concentrating on the importance of doing what you are required to do." This is an interesting further possibility for moral development. However, I think commonsense folk psychology does not lend support one way or the other, and I do not know of any empirical studies that support this. (Indeed, some of the empirical literature on self-control suggests that a way to resist temptation is to distract yourself in order to avoid considering what you ought to do altogether. This would be in order to keep the thoughts of temptation at bay.)

[54] Mark Muraven, and Roy F. Baumeister, "Self-regulation and depletion of limited resources: Does self-control resemble a muscle?" *Psychological Bulletin* 126 (2000): 247-259.

as taxing. Similarly, if we have a strong faculty of will-power, then acting contrary to other desires will be less strenuous.

By analogy, if someone has to make a conscious decision not to eat cake at each meal, then this regimen may well be very taxing. On the other hand, if they have engrained the habits not to do so, then refraining will be much less of an ordeal. They would do even better if they had engrained the habit not to even think about whether to eat cake. Further, if the cake-eater has strong will-power, then when the thought of cake does occur, it will not be such a struggle for them to banish it, and stick to their resolution.

So to the extent that we can make altruistic behavior habitual, and strengthen our will-power, we make altruism less difficult. This frees up room to be more altruistic, without the sum total of all this behavior being excessively difficult.

3.2. Giving more within the second limit: shaping our interests to make altruism less suffocating

Earlier, I made the assumption that the requirements of beneficence are constrained by our freedom to pursue other worthwhile interests:

> NO SUFFOCATION Beneficence does not require so much of us that our altruism is so suffocating that it leaves us insufficient room for other worthwhile interests.

Now, typically, our interests do not stay fixed. We acquire new ones. Some interests fade away. We also develop our interests in various directions. These are all ways of *shaping* our interests, over time. Doing so in the right way is a form of moral development.

First, we can shape our interests towards more altruistic forms. Let me give you an example of what I have in mind. I used to play on a soccer team with a man called Adam. Adam was a beefy man, with a shaven head—quite an intimidating proposal for opponents. However, his appearance turned out to be entirely deceptive. Adam was a big softie, who volunteered his time to promote youth sport among disadvantaged children in his neighborhood. Adam had taken his passion for soccer, and developed it into a way of helping others. He was enabling children to enjoy the sport. While doing so, he was also giving them the opportunity to learn about cooperation and sportsmanship. He was giving them an extracurricular goal to motivate themselves towards, which he hoped would make them more motivated in other aspects of their lives.

Crucially, Adam had found a way to help others that was not at the expense of his having a flourishing life himself. By all accounts, he seems to be positively thriving in his work with kids. His altruism is built around a distinctive interest of his that marks him out as an individual – his love of soccer. His life is still colored by his appreciation of one of the many forms of "personal excellence" which Wolf advocates.[55]

Second, we can shape our interests towards less expensive forms. For one, we can aim to rely less on material comforts. We can also take cost into account when deciding which activities to engage in. These vary in terms of how much time, effort and money they consume. Typically, many different projects could appeal to us. Life is too short to take an interest in all the worthwhile things. If we can steer our interests towards less expensive

---

[55] Wolf, "Moral Saints," p. 426.

forms, then we free up resources to spend helping others. And, we could do so while still having reasonably flourishing lives ourselves.

Take the following example. Jess recently traveled to Thailand, where she considered whether to learn to scuba-dive or learn how to cook Thai food. Both activities appealed to her in their own right, but she chose the cooking course on the grounds that it was less expensive. While she had not developed a particular interest in diving, foregoing it was no great loss to her life, and she thought it better to keep it that way. In this way, Jess avoided acquiring an expensive interest. This did not leave her life devoid of interesting activities. She simply economized.

Now, suppose Jess was already a scuba-diver when she arrives in Thailand. This is arguably a worthwhile pursuit, all else equal. Still, it is a relatively expensive interest, and she might reflect how far the money saved would go towards helping others. If she did, then she might choose to take a cooking course instead. In doing so, she could let one interest fall by the wayside, and replace it with another. Indeed, it is not uncommon for altruistic people to gradually lose their taste for expensive pursuits, as they continue to reflect on the costs, and the good that they could otherwise be doing with the money.

So we can shape our interests into more altruistic and less expensive forms. In this way, we can increase how much we give to help others, without this increase being suffocating.

4. A conclusion, a consequence and a coda

4.1. Conclusion

So there are ways for us to develop morally without this being excessively difficult (habituation and building will-power), and ways for us to develop morally to increase how much we give without this being suffocating (shaping our interests). This completes my defense of a key premise of my main argument:

> MEANS   Moral development is a means to giving more to help others without this being too difficult or suffocating.

My other main premise was that if there is a course of action open to us that allows us to help others more without this being excessively difficult or suffocating for us, and there is no other reason not to, then beneficence requires us to take this course. On the assumption that there is no other reason not to develop morally and give more, these two premises constitute my argument for my central conclusion:

> DEVELOPMENT (i) Beneficence requires us to develop morally;
>
>   (ii) The more morally developed we are, the more beneficence requires us to give.

Thus concludes the instrumentalist rationale to develop morally on the grounds that doing so and giving more is a means to alleviating others' needs.

DEVELOPMENT is the key to unravelling the puzzle with which we began. For we can accept this thesis, while positing the two limits to beneficence encountered earlier:

No Suffocation	Beneficence does not require so much of us that our altruism is so suffocating that it leaves us insufficient room for other worthwhile interests.
Tolerable Difficulty	Beneficence does not require so much of us that meeting its requirements is excessively difficult for most of us.

Now, in order to accommodate these two limits, an account would have to take a stand on what *rate* of moral development is required, and where it is ultimately headed. In addition to any limits on what is humanly possible, the rate should not be so fast that moral development is excessively difficult or suffocating. Similarly, the final destination of development would be an ideal life which is not excessively difficult nor suffocating. I mentioned earlier the philosophical difficulty with determining what counts as suffocating. This turned, in part, on the difficulty in judging how much weight to give to the fact that someone identifies with the particular interests that he or she happens to have. This difficulty resurfaces with the requirement to shape our interests. Even without being able to address this difficulty, we can say at a minimum that the requirement to shape interests requires at least this much—that among the various forms that appeal to us, we shape our interests towards more altruistic and less expensive forms.

By respecting these two limits to the requirements of beneficence, DEVELOPMENT could be accepted by someone who holds a moderate view of beneficence:

MODERATION	Beneficence does not require so much that an impressive altruist like Amy is failing to give enough.

Note that here I am not arguing that DEVELOPMENT motivates MODERATION. Indeed, I suggest that it does not—a demanding view of beneficence may well also include DEVELOPMENT. Rather, my claim is merely that DEVELOPMENT is compatible with MODERATION. This is in keeping with my overarching project which was to assume, (rather than argue for), a moderate view of beneficence, and explore how it could accommodate a requirement to be ambitious.

I hope to have shown that it could accommodate this because DEVELOPMENT makes it appropriate for an altruist to be ambitious. Someone like Amy ought to be aiming to develop morally, and thereby increase how much she give:

AMBITION	Beneficence requires an impressive altruist like Amy to be ambitious, and aim to give more than she is currently giving.

Thus, a requirement of beneficence to develop morally provides a solution to the puzzle with which we began. The key to including both moderation and ambition in an account of beneficence is to turn our attention from momentary questions about what to give at a particular moment, and turn to the question of what pattern altruism should take over time. The answer defended here is that the dynamic moral life is one of moral development, and of increasing levels of altruism. An ambitious attitude like Amy's is the correct stance to take towards the future stages of such a life.

Does this resolve all of the conflict that an altruist like Amy may feel? Sadly, not. So far, all she knows is that she is required to develop morally, and increase how much she gives, when all else is equal. There remains for her the further question of how much she ought to be giving at a particular time. She may well feel conflicted about this. She may feel unsure whether her current level of altruism is enough *for now*. Sadly, resolving this conflict is work for another occasion.

4.2. A (controversial?) consequence: different targets for different ability levels

Now, DEVELOPMENT has a potentially controversial consequence. According to DEVELOPMENT, Amy ought to develop and increase how much she gives as she does. So as she develops into a more morally mature person, she will be required more to give than she was before. In this way, the amount she should give is indexed to stages of her moral development.

An upshot of this is that the requirements of beneficence will turn out to be indexed to persons. For, persons will vary in their stage of moral development; Amy may be more morally developed than Belinda, for example. Assuming all else is equal in their lives, then Amy would be required to give more than Belinda. So the level of altruism that beneficence requires may vary from individual to individual.

This personal variation may seem objectionable to some. For a start, it would seem more fair if everyone were required to give a more uniform amount. But things may seem even worse. This is because of the reason that Amy is required to give more: she is so required because she is a better person. Thus, in a certain sense, Amy would appear to be punished for being a better person, while Belinda is let off the hook for being a worse one.

Now, this may seem to you to be a bullet to bite, and if so, I hope the positive arguments in favor of DEVELOPMENT make up for this cost. Alternatively, you may think that it makes sense to hold people with greater moral potential to higher standards. Moreover, we should enquire into the reason why Belinda is less mature than Amy. If the explanation is that Belinda has failed to develop morally, then the account has it that Belinda is failing to meet a moral requirement. If this is right, then the unfairness would seem to lie not with the account of beneficence, but at the hands of Belinda herself.

4.3. Coda: an ambitious altruist

I will end by returning to a champion of a demanding view of beneficence–Peter Singer. Apparently, Singer is often asked whether he is living a simple life and donating his income to help the poor. He responds

> I'm not living as luxurious a life as I could afford to, but I admit that I indulge my own desires more than I should. I give about 25% of what I earn to NGO's, mostly to organizations helping the poor to live a better life. I don't claim that this is as much as I should give. Since I started giving, about thirty years ago, I've gradually increased the amount I give, and I'm continuing to do so.[56]

---

[56] http://www.princeton.edu/~psinger/faq.html as of 1/25/10.

In this response, Singer claims he is not living as he ought to, since he holds a demanding view of beneficence. To many of us, his own altruism seems a counterexample to this view—25% is a very impressive amount to be giving, and it would be counterintuitive to judge that he is failing morally. Instead, I have recommended we adopt a more moderate account of beneficence that includes a requirement to develop morally. But I want to highlight that Singer exhibits just the sort of ambition for altruism that I have argued is required. He is trying to do better, and increase how much he gives, year by year. Now, you might think that many moral philosophers would do better to practice what they preach. If my proposed account is correct, then, unusually, Singer would do better to preach what he practices.

Chapter Three
Chancy Charity and Aggregative Altruism

1. Introduction

It is plausible to think that it is wrong to cure many people's headaches rather than save someone else's life. On the other hand, it is plausible to think that it is not wrong to expose someone to a tiny risk of death when curing this person's headache. I will argue that these claims are inconsistent. For if we keep taking this tiny risk, then it is likely that one person dies, while many others' headaches are cured. This conflict is perplexing. I will consider a possible way to reconcile the intuitions in question. But this proposal faces serious problems of its own.

2. Intuitive opposition to aggregation

I will begin by fleshing out the two claims that I will argue are inconsistent. The first claim concerns decisions about whom to help.[57] Consider the following case:

> Islands. If you go to the North island, then you can save Jones's life. If you go to the South island, then you can cure the headaches of a billion people. You only have enough fuel to go to one island.[58]

You have two options. First, you can provide a large benefit to Jones by saving Jones's life. Second, you can provide a small benefit to a billion people by curing their headaches. Intuitively, you ought to save Jones's life. Moreover, you ought to do so, no matter how many people's headaches you could otherwise cure.

What principle would explain the fact that you must save Jones? A natural candidate is the following:

> MANY-FEW   You may never provide small benefits to many people rather than save the life of someone else. (All else equal).

When I say you may not provide these benefits, "all else equal," I am holding fixed features of these people, such as their relationship to you. I am also assuming that your behavior affects only the people who could receive the small benefits and the person whose life you could save.

If MANY-FEW is correct, then we should reject explanations that endorse a utilitarian approach to beneficence. According to this approach, small benefits to sufficiently many

---

[57] I am assuming that the potential beneficiaries do not have claim-rights against you, and that you do not act in any institutional capacity.

[58] This case is inspired by Scanlon's classic case:

> World Cup. A worker has fallen onto electric cables in the television studio and is in agony. If you turn off the power, billions of people miss the climax of the World Cup final.

T. M. Scanlon, *What We Owe To Each Other*, (Cambridge MA: Harvard University Press, 1998).

people "aggregate" to form a total amount of welfare that is more important than Jones's survival. If you ought to save Jones in Islands, then the utilitarian approach to beneficence is misguided.[59]

3. Intuitive tolerance of risk-taking

The second view concerns risk-taking. Sometimes, we take minuscule risks that someone will die in order to provide this person with a small benefit. To spare Amanda the inconvenience of taking the subway, I might drive her to the airport. I do so, even though this exposes her to a slightly increased chance of death. (I will ignore the risks to pedestrians and other motorists.) For Ben's culinary pleasure, I might serve him raw fish. I do so, even though cooking the fish would slightly reduce his risk of dying from a fatal bacterium. Do we sometimes kiss people for their sake? If so, we should pay heed to James Lenman's observation that

> it seems morally pretty unobjectionable to kiss one's lover in circumstances where there is (and I guess there *always* is) a fantastically small chance that, for example, one has contracted some deadly transmissible disease which one thereby passes on to them.[60]

It seems a kiss is worth the coyest flirtation with death.

What would explain the fact that we may take these risks? A natural candidate is the following principle:

> RISK TOLERANCE You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else equal.)

Again, the "all else equal" clause is in place to restrict the scope of the principle—the clause ensures that the principle only has bite in a case where the only morally significant feature of your action is that it has these chances of harms and benefits. I leave for an appendix a discussion of what difference it makes whether this risk is metaphysical or doxastic in nature. (Quick answer: little difference.)

So we have seen two claims that are initially plausible. First, you may never provide small benefits to many rather than save someone else's life. Second, you may expose someone to a tiny risk of death when providing this person with a small benefit.

4. The Repetition Argument

I will offer an argument that turns on the likely effects of repeating risky actions. It is helpful to consider the repetition of risks when these risks are small. This allows us to put these small risks "under the moral microscope," in terms of Derek Parfit's apt phrase.[61] Parfit's preferred microscope is universalization: he magnifies a small effect of an action

---

[59] This is the point that Scanlon makes about his World Cup example described in the previous footnote. He offers this case as a counterexample to aggregative theories. Scanlon, *What We Owe To Each Other*.

[60] James Lenman, "Contractualism and risk imposition," *Politics Philosophy Economics* 7 (2008): 99 - 122, at p. 112.

[61] Derek Parfit, *On What Matters*, manuscript.

by considering *everyone* performing the action.[62] In addition, we can magnify a small effect of an action by considering *you* performing the action many times: another way of putting a risk under the moral microscope would be to consider your repeating the risk.

The main argument of this paper turns on the likely effects of repeating the risk.[63] I will call it the "Repetition Argument. It aims to show that RISK TOLERANCE and MANY-FEW are inconsistent. It does so by starting with a premise entailed by RISK TOLERANCE, and reaching a conclusion that contradicts MANY-FEW.

To run my argument, I will need an example of a risky action that you would be permitted to perform, according to RISK TOLERANCE:

> RISK TOLERANCE You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else equal.)

I will choose the following case:

> Poison. In a nearby room, Victim has swallowed poison. If you do nothing, then Victim will have a headache and then die. You can bring Victim only one of two antidotes:
>
> > The Reliable Antidote is certain to save Victim, but will do nothing for Victim's headache.
> >
> > The Chancy Antidote is most likely to cure Victim's headache and save Victim's life, apart from a one in a billion chance that it does nothing.

If you choose the Chancy Antidote, then you are most likely to relieve Victim's headache, but take a tiny risk that you fail to save Victim's life. I assume that if the situation described in Poison occurs in isolation, then you may bring the chancy antidote to Victim, according to RISK TOLERANCE. I invite you to alter the size of the risk as you see fit so that you are satisfied that RISK TOLERANCE entails that you may bring the chancy antidote. For example, if you feel that a one in a billion chance is too large, then I invite you to make it smaller. The numbers that I use will have no substantial effect on the argument that I offer. So I will assume that if we accept RISK TOLERANCE, then we should accept the following:

> 1. You may bring the chancy antidote in Poison once. (From RISK TOLERANCE)

This will be the first premise in the argument.

Now, suppose the situation described in Poison occurs again and again. On the first day, Smith is in the position of Victim, then the next day Roberts is in this position, and so on. If you may bring the chancy antidote to Smith, then you may also bring the

---

[62] Some moral theories stress the importance of universalizability. For example, Kantianism holds that you ought only act on a maxim that you could rationally endorse being followed as universal law. Meanwhile, some rule-consequentialist theories hold that you ought to follow the rules that would lead to the highest expected value if everyone followed them.

[63] For a similar argument, see Alaistair Norcross, "Great harms from small benefits grow: how death can be outweighed by headaches," *Analysis* 58 (1998): 152 - 158.

chancy antidote to Smith and then bring it to Roberts on the following day. If the risk of bringing the chancy antidote is such that you may bring it to one person, then the risk is such that you may bring it to one person and then another. Indeed, we could suppose that you face a very long sequence of people in exactly the same situation. Assuming you may bring the chancy antidote to each person alone, you may bring it to each of them in turn:

2. If you may bring the chancy antidote in Poison once, then you may bring the chancy antidote to any sequence of people in Poison.

From these two premises, we can infer:

3. *Therefore*, you may bring the chancy antidote to any sequence of people in Poison. (From 1 and 2)

Now, we should ask what will happen if you do bring the antidote to many people.

As you increase the number of times that you repeat a risk that there is someone who dies, it becomes increasingly likely that one of these repetitions results in there being someone who dies. Let us bring this point to bear on Poison. Suppose this situation repeats with different people playing the role of Victim. Indeed, suppose it occurs one billion times, and you bring the chancy antidote each time. Among these one billion repetitions, we would expect that one time you will fail to save someone's life. Moreover, the "Law of Large Numbers" entails that as the number of times that you bring the chancy antidote increases, it becomes increasingly likely that you fail to save someone's life one in every billion times.[64] The law entails that it is near certain that this distribution will result with a sufficiently large, finite sequence. Note that I really do mean *near certain*. If we make this sequence of people large enough, then we can bring the chance as close to certainty as we like.[65] If we made the sequence long enough, then the probability could be 0.99999999, for example. Thus:

4. If you bring the chancy antidote to a sufficiently large sequence of people, it is near certain that you fail to save the life of one in every billion people to whom you bring the chancy antidote and cure the headaches of the rest. (From the Law of Large Numbers)

Notice that these repetitions do not increase the chances of your failing to save any particular person—this chance remains one in a billion. To think otherwise would be to commit the gambler's fallacy.

Continuing the argument, we can infer:

5. *Therefore*, you may make it near certain that you fail to save the life of one in every billion people to whom you bring the chancy antidote in Poison and cure the headaches of the rest. (From 3 and 4)

---

[64] Roughly, the law says that as the number of repetitions of the risky action tends to infinity, the probability that the resulting distribution is one death in every billion tends to one (certainty). For our purposes, the differences between the weak and strong versions of the law do not matter.

[65] Indeed, the Law of Large Numbers entails that if you repeated this action *infinitely* many times, then statistically you would be guaranteed to fail to save someone's life one in every billion times—the probability of this pattern emerging is 1.

That is, repeating these risks would be likely to result in a particular distribution of benefits. The distribution is that for every billion people to whom you bring the chancy antidote, one person dies, while all the rest have their lives saved and have their headaches cured. By always bringing the reliable antidote, another distribution would result. This distribution is that for every billion people to whom you bring the chancy antidote, all have their lives saved but suffer headaches. The comparative difference between these distributions is that in the former 999,999,999 people enjoy a small benefit of having their headaches cured, while one person suffers the large loss of death:

6. *Therefore*, you may make it near certain that you cure the headaches of 999,999,999 people rather than save the life of someone else. (From 6)

This is the consequence of your being permitted to repeat the risky action.

This brings us to another substantive premise in the Repetition Argument:

7. If it is wrong to bring about a particular distribution, then it is wrong to make it near certain that this distribution obtains.

By contraposition, if it is permissible to make it near certain that the distribution obtains, then it is permissible to bring about the distribution. This premise is an important step in the argument as I offer it here.[66] But it strikes me as very plausible, and I cannot see any reason for denying it. Any reason not to bring about a distribution would equally seem to be a reason not to make the distribution near certain. Recall that I really mean *near certain* here. The likelihood of this distribution arising can be as close to certainty as we like. (For reasons of space, I will not illustrate this point by giving an example of a probability that would take several pages of 9s to write down.) I suggest that we accept this premise unless we are given some reason not to do so. I cannot see any such reason. So I conclude that we ought to accept it.

Accepting this premise puts us in a position to infer:

---

[66] In addition, there are two ways of running a similar argument that appeal to weaker premises.

First, Richard Holton has pointed out to me that an analogous argument could be run that relies on this premise instead:

7* If it is wrong to bring about a particular distribution, then there is a probability, p, such that it is wrong to make the probability of this distribution obtaining p.

This would require some tweaks to the other premises in the argument that strike me as benign, but would be complicated to write out. For example, we would have to replace references to "near certainty," with references to this probability p. Stating this would require some complexity to write down, as it would get complicated specifying the quantifiers that bind each premise, which is why I avoid doing so here.

The second way of altering the argument begins by strengthening premise 2 to 2*:

2* If you may bring the chancy antidote in Poison once, then you may bring the chancy antidote to any sequence of people in Poison.

By the Law of Large Numbers:

4* If you bring the chancy antidote to an infinite sequence of people, there is a probability of 1 that you fail to save the life of one in every billion people to whom you bring the chancy antidote and cure the headaches of the rest.

Then we would only need to appeal to the weaker premise 7**:

7** If it is wrong to bring about a particular distribution, then it is wrong to make it the case that there is a probability of 1 that this distribution obtains.

This premise strikes me as undeniable.

8. *Therefore*, you may cure the headaches of 999,999,999 people rather than save the life of someone else. (From 6 and 7)

But this conclusion in line 8 contradicts the thesis that we encountered earlier:

> MANY-FEW   You may never provide small benefits to many people rather than save the life of someone else. (All else equal).

So by starting with an assumption entailed by RISK TOLERANCE, we have ended with a conclusion that contradicts MANY-FEW.

Let me summarize the Repetition Argument:

1. You may bring the chancy antidote in Poison once. (From RISK TOLERANCE)

2. If you may bring the chancy antidote in Poison once, then you may bring the chancy antidote to any sequence of people in Poison.

3. *Therefore*, you may bring the chancy antidote to any sequence of people in Poison. (From 1 and 2)

4. If you bring the chancy antidote to a sufficiently large sequence of people, it is near certain that you fail to save the life of one in every billion people to whom you bring the chancy antidote and cure the headaches of the rest. (From the Law of Large Numbers)

5. *Therefore*, you may make it near certain that you fail to save the life of one in every billion people to whom you bring the chancy antidote in Poison and cure the headaches of the rest. (From 3 and 4)

6. *Therefore*, you may make it near certain that you cure the headaches of 999,999,999 people rather than save the life of someone else. (From 5)

7. If it is wrong to bring about a particular distribution, then it is wrong to make it near certain that this distribution obtains.

8. *Therefore*, you may cure the headaches of 999,999,999 people rather than save the life of someone else. (From 6 and 7)

Since line 1 is entailed by RISK TOLERANCE, and MANY-FEW contradicts line 8, the argument shows we cannot hold both principles.

5. Defending the iterative premise

To resist the argument, one would have to drive a wedge between the permissibility of taking a risk, and the permissibility of bringing about a distribution that is near certain to obtain as a result of repeating the risk. A natural way of doing so would be to deny the iterative premise:

2. If you may bring the chancy antidote in Poison, then you may bring the chancy antidote to any sequence of people in Poison.

One might object that you may take a risk occasionally, but that you may not repeat the risk so many times that there is a high chance that someone denies. For example, Derek Parfit claims

> We can usually ignore a very small chance. But we should not do so when we may affect a very large number of people, or when the chance will be taken a very large number of times. These large numbers roughly cancel out the smallness of the chance.[67]

Suppose we applied Parfit's reasoning to Poison. If you faced the case only once, then you may bring the chancy antidote since you can "ignore a very small chance." But you may not bring the chancy antidote to many people since then "the chance will be taken a very large number of times." As a result, I would expect that Parfit may be inclined to deny the iterative premise.

This line of thought may be tempting, but it springs from a mistake that we often make when thinking about risks. In evaluating a set of risky actions, we sometimes make a mistake of not appreciating the importance of the *independence* of the risks in question. Here is an example of the same mistake.[68] Suppose you leave your house, and worry that you have left the gas on. You judge there is a very small risk that you have. It is tempting to hold two thoughts. First, it is fine for you to take this risk *just this once*. The convenience of not going back to check is worth the risk. Second, it would not be fine for you to take this risk very many times. Were you to do so, it would become likely that one time you do leave the gas on, and this would not be worth many small conveniences. Conjoining these two thoughts, we might say that it is fine for you to tempt fate now and again, but that you may not tempt fate too much.

However, this is a mistaken way to think about risks. Suppose it is Tuesday and you are considering whether or not to go back to check the gas. Whether you may continue on your way without checking the gas depends only on the effects associated with this particular action—do the risks of harms and benefits associated with continuing without checking justify your doing so? Whether you may continue does not depend on whether you continued last week, or the week before. Nor does it depend on whether you go on to do so in the future. All that matters when deciding to go back to check on the gas on any particular occasion is the risk associated with this particular occasion. So if the sum of the conveniences is not worth the sum of the risks, then each convenience is not worth each risk.[69]

This point applies with equal force to Poison. The correct way to think about the risks is to look at the possible or actual effects of each action independently, and assess the permissibility of each action on these grounds. The permissibility of exposing someone to a risk does not depend on whether you have previously exposed others to this risk. The following principle is a plausible way of capturing this point:

---

[67] Derek Parfit, *Reasons and Persons*, (reprinted with further corrections), (New York: Oxford University Press Inc.), p. 75.

[68] Thanks to Richard Holton for suggesting this type of example.

[69] I assume here that the conveniences have constant marginal utility for you.

43

ONLY EFFECTS MATTER   When only one person is affected by a risky action, the permissibility of the action depends on the possible or actual effects of *that action* for *that person*.

Now, this principle entails the iterative premise. Consider a sequence of people, Marks, Roberts and Smith. May you bring the chancy antidote to each in sequence? This question would depend only on the risks involved for each person, individually. If bringing the chancy antidote does not create too large a risk of death for Marks, Roberts, and Smith individually, then it must be permissible for you to bring it to all of them in sequence. Conversely, if you are forbidden from bringing the chancy antidote to all in sequence, then this could only be because the risk for some individual was too high.

Note that it is crucial that we are dealing only with cases where the magnitudes of the risks, benefits and harms are independent of each other.

First, I assume that the size of the benefits and harms are the same, and that these are the only morally relevant features associated with taking these risks. For example, I assume, plausibly, that whether Marks dies does not affect how bad it is for Roberts to die. Moreover, I assume that the only morally relevant feature of their deaths is their well-being. That is, I assume that we are not dealing with a case where there are further undesirable effects that arise from them both dying. For example, I assume that we are not dealing with a case in which they are the last two remaining males in the species. In this case, there would be a further morally relevant consequence of both dying—no one would be able to procreate to continue the human race.

Second, I assume that the risks are statistically independent. For example, whether you give Marks the reliable or chancy antidote makes no difference to the size of the risk that you would expose Roberts to by bringing Roberts either antidote. Things would be different if the risks were cumulative:

> Ship. A ship is sinking with many children on board. There are two lifeboats that you can direct them to.
>
>> The reliable lifeboat is a large and sturdy boat. It is very safe and certain to reach the shore. It will take a day to reach land, by which time any children on board will be tired and cranky.
>>
>> The chancy lifeboat is a small and light boat. It will reach the shore in an hour. However, there is a one in one hundred million chance that it sinks if it has one child passenger on board. Furthermore, for each additional child that is on board, the probability of it sinking doubles.

Let us suppose, for the sake of argument, that it is permissible to take a risk of a child's death in order to spare this child several hours' crankiness if and only if the risk of death is below one in 30,000,000. (Again, my choice of these numbers is arbitrary, so you should feel free to change them. For our purposes, all that matters is that the risk is cumulative, and there is a threshold of risk that you may take in pursuit of this benefit.) Now, whether you have directed toddlers Timmy and Sammy to the reliable lifeboat

matters a great deal to whether you may direct toddler Vicky to this boat. For if both Timmy and Sammy are the only children on board the reliable lifeboat, then the probability of the boat sinking is one in 50,000,000. Were Vicky to board, this probability would rise to one in 25,000,000—an unacceptably high risk. But if neither Timmy and Sammy are on board, nor any other child, then the probability of the boat sinking with Vicky on board is merely one in one hundred million—a risk that we have stipulated to be acceptably low in light of the benefit to her.

In the Ship case, you may not repeat the risky action of directing a child to the chancy boat. This is because the risks are cumulative: exposing one child to a risk increases the risk that you could expose the next child to. Still, it is the risk of an action to each child that affects the permissibility of performing this action. In the Poison case, you may repeat the risky action, (if you may perform it at all), because it does not affect the risks that you expose other individuals to. Again, all that matters is the risk that each person faces.

6. A possible solution: do ex ante interests matter?

Assuming the Repetition Argument is sound, it shows that we cannot hold both of these claims:

> RISK TOLERANCE   You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else equal.)

> MANY-FEW   You may never provide small benefits to many people rather than save the life of someone else. (All else equal).

These claims seemed attractive because they could explain intuitions of ours. Let us focus on MANY-FEW. The reason why we were drawn to MANY-FEW is that it is a way of capturing an intuition about the case we encountered earlier:

> Islands. If you go to the North island, then you can save Jones's life. If you go to the South island, then you can cure the headaches of a billion people. You only have enough fuel to go to one island.

Can we explain the intuition that you ought to save Jones in Islands by appealing to a different principle that also tolerates bringing the chancy antidote to many people in Poison?

To do so, we should look for a morally important difference between the two cases. There is one difference that looks as if it may be important. In Islands, there is a single person, Jones, whom you would fail to save, were you to cure others' headaches. On the other hand, if you keep bringing the chancy antidote to sufficiently many people in Poison, then you would be near certain to fail to save someone, were you to cure others' headaches. But there is no single person whom you would be near certain to fail to save.

45

In light of this observation, let us coin some terminology. Let us say that a risky action is in someone's "ex ante interest" if this action produces a positive amount of expected utility for them.[70] In these terms, we can formulate this principle:

> EX ANTE VIEW   You may let someone die and provide small benefits
> to many when and only when your doing so is in
> everyone's ex ante interest. (All else equal).[71]

According to this principle, you may not fail to save Jones in Islands—doing so is not in Jones's ex ante interest. Moreover, this principle would also explain why you may repetitively bring the chancy antidote to people in Poison—doing so is in each person's ex ante interest, including the ex ante interest of anyone who turns out to die.

Why would it be morally significant that an action is in everyone's ex ante interest? A possible answer is the contractualist claim that our actions must be capable of being "justified to *each* affected person," as James Lenman puts it.[72] Some people argue that you could justify your action to someone if the action was in this person's ex ante interest.[73] In a similar vein, one could motivate the EX ANTE VIEW by stressing the importance of *hypothetical consent*. Someone would be able to hypothetically consent to an action if this action was in their ex ante interest.

7. Problems with the EX ANTE VIEW

The EX ANTE VIEW might seem like the way to reconcile our intuitions about Poison and Islands. But there are two significant problems with the view.

The first problem is that it is not clear that the view is well motivated. There are well known controversies about contractualism and hypothetical consent that I will not enter

---

[70] We should calculate expected utility in the traditional method—we multiply the utility of each harm or benefit by the probability that it results, and then sum these products.

[71] Michael Otsuka suggests a view similar to this one. Within the context of his view that the justification of actions depends on the complaints that people might have against them, he proposes that

> for any individual who ends up suffering harm as the result of risky activity, we might
> allow for the discounting of his complaint against suffering harm by the improbability
> that *he in particular* would suffer harm in the following restricted set of circumstances: those
> in which it is in his ex ante self-interest to be exposed to such risks when one takes into
> account the expected benefit to him of such risky activity.

Since Otsuka appears to aim to justify everyday actions that expose people to risk, I take him to hold that we should "discount" this person's "complaint" so much that the actions in question are permissible.
Otsuka, "Life and Limb," manuscript, p. 15.

[72] Lenman, "Contractualism and Risk Imposition." Lenman denies that an action's being in someone's ex ante interest is sufficient for being justifiable to that person. He thinks that in addition the action must not preclude taking reasonable precautions, and must be capable of being guided by others' safety. Thus, he suggests that a contractualist adopt the "Aim Consistency Principle," which

> would then say something like this: my act of imposing a risk on you may be rendered
> permissible by furthering some purpose that does not in principle preclude my taking
> reasonable precautions against your coming to harm; other acts I might perform are
> rendered impermissible when the purpose from which they might seek a warrant is
> inconsistent with your safety's being one of my guiding aims. (p. 111)

[73] See Johann Frick, "Health Resource Allocation Behind a Natural Veil of Ignorance," manuscript.

into.[74] In addition, independently of these worries about contractualism and hypothetical consent, there is a further worry as to why *ex ante* interests should matter. Suppose a risky action ends up harming somebody. Why would it make a moral difference that at *an earlier time* this action was likely to be in their interest? Why should we be concerned with the likely effects of an action rather than its actual effects? These are questions for the contractualist as much as the rest of us: why would actions be justified to someone on the basis of the actions' likely effects rather than their actual affects?

The second problem is that the EX ANTE VIEW is committed to counterintuitive claims. Consider:

> Wheel. There is a giant roulette wheel in the sky. It has a billion slots that are connected to chutes. The chutes are directed at individuals who have swallowed poison. The poison will make them have a headache and die, if untreated. To make the wheel work, you have to press a button. This will give the wheel an indeterministic spin, and then release the contents of each slot into a chute. Before you do so, you have to make an antidote to put in the slots. You only have enough resources to choose between these options:
>
>> Vials For Everyone. You make a billion vials of an antidote that will certainly save each recipient's life, but do nothing for each recipient's headache.
>>
>> Vials That Cure Headaches. You make a billion minus one vials of an antidote that will certainly save each recipient's life and cure this person's headache.[75]

The net effect of choosing Vials That Cure Headaches rather than Vials For Everyone is that 999,999,999 people have their headaches cured and one person dies. Intuitively, you may not choose Vials That Cure Headaches. However, the EX ANTE VIEW predicts that

---

[74] One objection is that an action is justifiable to someone because it is permissible, and so contractualism gets the order of explanation wrong.

A second objection is that contractualism can only manage to be a rival to consequentialism at the expense of giving up core intuitions about cases. For example, David Wasserman and Alan Strudler argue that Scanlon has failed to show how his version of contractualism explains why the numbers should count. Wasserman and Strudler, "Can a Nonconsequentialist Count Lives?" *Philosophy & Public Affairs* 31 (2003): 71-94.

[75] James Lenman considers a similar government policy consisting of performing two actions. First a lottery is held. Then the winners are given a benefit and the losers are killed by the police. Lenman, "Contractualism and Risk Imposition," p. 102. He proposes this as a counterexample to the claim that *policies* are justified when they are in the ex ante interest of all. Later, he considers "Russian Roulette Cases," which involve a single *action*. Some of these seem to be a variant of the objectionable cases where:

> the push of a button sets the lottery in train and some complex, entirely automatic mechanism then brings it about that the losers are killed without the involvement of the police or any other human agency after the button is pushed. (p. 112).

His remarks imply that he would consider it impermissible to push the button in such a case, and that this would refute the claim that an action is permissible when it is in the ex ante interest of all. Lenman claims that his "Aim-Consistency Principle" (see previous footnote) can explain why this is so, on the grounds that

> it is hard to see how [the] redeeming feature [of not precluding the taking of precautions] could be present at least in any remotely realistic example involving Russian roulette or elaborate, automated, survival-lottery machines. (p. 112).

you may. This is because this option would be in the ex ante interest of all. (I assume that a one in a billion chance of death is small enough for headache relief to compensate for it. Feel free to make the probability smaller until you agree.)

8. Conclusion

These problems with the EX ANTE VIEW fall short of a refutation. But they do raise trouble for it. If we do not find the EX ANTE VIEW satisfactory, then we are left in a puzzling situation. Is there an alternative way of reconciling our intuitions about Islands and about everyday actions that expose potential beneficiaries to minuscule risks of death? If so, what is it?

If no satisfactory way of reconciling these intuitions is forthcoming then we must revise some of our intuitions. I will briefly explore the reasons for revising the intuition about Islands, and the significance of doing so. Then I will do the same for our intuitions about everyday risky actions.

The first option would be to revise our intuition about Islands. We might motivate doing so by noting that it is an intuition about an incredibly large number of people. Our ability to compute such large numbers may be limited. For example, I know that my attempts to imagine a million people are suspiciously like my attempts to imagine a billion. Maybe we are unable to appreciate the moral significance of helping so many people.

Revising the intuition about Islands would be important for our moral theory. It would affect what view we should take of a utilitarian approach to beneficence. The utilitarian approach endorses "aggregation"—it aggregates benefits to individuals into total amounts of welfare, and recommends that you bring about the greatest total amount of welfare possible. Some of us thought that cases like Islands were the death knell for the utilitarian approach to beneficence. Moreover, we thought that the fall of utilitarianism was progress in moral philosophy. If our intuition about Islands is misguided, then our rejection of the utilitarian approach would be called into question.

The second option would be to revise our intuitions about everyday risky actions. We could motivate this option by noting that these intuitions concern incredibly small probabilities. Our ability to compute such small numbers may be limited. For example, I know that my deliberations about a one in a billion risk are suspiciously like my deliberations about no risk at all. Moreover, some people's intuitions about risks depend on whether these risks are described quantitatively or qualitatively. For example, they may fear a risk of "one in a million," but be comfortable with a risk "equivalent to that of being hit by lightning." Indeed, empirical work suggests that in general we reason poorly about probabilities: we adopt heuristics, and commit probabilistic fallacies.[76] Lastly, our intuitions about the permissibility of taking certain risks may be influenced psychologically by the fact that we *do* take these risks. The psychological literature on "cognitive dissonance" finds that we are under psychological pressure to avoid the unpleasant feeling of believing both that a behavior is wrong and that we engage in this

---

[76] See, for example, Amos Tversky, and Daniel Kahneman, "Judgment under uncertainty: Heuristics and biases," *Science* 185 (1974): 1124–1131.

behavior.[77] This puts pressure on us to rationalize our actual behavior: we are under pressure to judge that our behavior is permissible. So it may be that these intuitions about everyday risks are less trustworthy than we might hope.

However, if we revised our intuitions about everyday actions, then we would have to change the way we live parts of our everyday lives, on pain of acting impermissibly. We would have to be much more cautious than we are now when helping people. We would have to avoid all beneficial actions that expose the potential beneficiaries to any risk of death at all. This is not an exciting prospect.

So revising neither the intuition about Islands nor our intuitions about everyday risks is attractive. But without a satisfactory principle that lets us reconcile these intuitions, we face the threat that we may have to embrace one of these unattractive options. We would face the threat that our lives should include giving no lifts to the airport, serving no sushi, and planting no kisses. Or worse, we face the threat that the utilitarian approach to beneficence was correct all along, and that the progress some of us thought we were making in moral philosophy was illusory.

---

[77] See, for example, Leon Festinger, *A theory of cognitive dissonance*, (Stanford, CA: Stanford University Press, 1957).

Appendix to Chapter Three
Subjective and objective risks

The main argument of Chapter Three is the Repetition Argument. It aims to show that we cannot hold both of these claims:

RISK TOLERANCE   You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else equal.)

MANY-FEW   You may never provide small benefits to many people rather than save the life of someone else. (All else equal).

In this appendix, I will discuss the issue of whether the risks in question are metaphysical or doxastic in nature. I will argue that it makes little difference, if any, to my previous discussion.

Before I can make this point, I need to set up some machinery. Let me begin by invoking the orthodox distinction between objective chance and credence.[78] Objective chance is metaphysical in nature. Objective chance is surprisingly difficult to define,[79] but we do have an intuitive grip on the notion: we have an intuitive grip on the notion that a coin toss is *itself* chancy, regardless of our attitudes towards the toss. Metaphysical chance is different from doxastic credence. Your credence in a proposition is your degree of belief that this proposition is true.

Let us illustrate the difference between objective chance and credence with an example. Suppose you do not know whether there is an even or odd number of words in yesterday's Boston Globe. You are equally inclined to believe that the number is even or odd. In fact, there was an odd number. The objective chance of there being an even number is zero. But your credence in the proposition that there is an even number is half.

In light of the distinction between objective chance and credence, let us distinguish objective risk from subjective risk. An action has an "objective risk" of a harm just in case there is an objective chance that the action leads to this harm. By contrast, an action carries a "subjective risk" of a harm from your point of view just in case on the basis of your evidence you ought to assign some credence to the proposition that the action leads to this harm.

We can use the previous example to illustrate the difference between objective risk and subjective risk. Suppose that you are offered a bet: you win $10 if there was an odd number of words in yesterday's Boston Globe, and lose $10 if there was an even number. Again suppose there was an odd number. Taking this bet has no objective risk, as there is no objective chance that there was an even number—there is no objective chance you can lose the bet. But in light of your evidence, you should have a credence of 0.5 in the

---

[78] The locus classicus is David Lewis, "A Subjectivist's Guide to Objective Chance," in Richard C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, 1980, (Berkeley CA: University of California Press).

[79] Take the definition that an event is objectively chancy when previous events do not determine whether it occurs. Even this definition is controversial, since some people hold that statistical mechanics involves objective chances, even on the assumption that determinism is true. See Barry Loewer, "Determinism and Chance", *Studies in History and Philosophy of Modern Physics* 32B(4) (2001): 609-620.

proposition that there was an even number of words. Therefore, from your point of view, taking the bet exposes you to a subjective risk of losing $10.

Armed with this machinery, we can ask the question of whether everyday actions are objectively risky or merely subjectively risky. I will argue that (i) our intuitions that support RISK TOLERANCE are intuitions about objective risks; (ii) in light of this the Repetition Argument uncovers a conflict between this thesis understood as a thesis about objective risks and MANY-FEW; and (ii) whether or not everyday risky actions are objectively or merely subjective chancy does not affect how problematic it is to reject RISK TOLERANCE when it is understood as a thesis about objective risks.

(i) Our intuitions that support RISK TOLERANCE are intuitions about objective risks

Consider RISK TOLERANCE. When we distinguish objective and subjective risk, we can ask which type of risk we tolerate, and make the thesis more specific, accordingly. This principle was motivated by our intuitions about everyday actions. We intuitively judge that we may give people lifts in a car, that we may serve them raw fish, and that we may kiss them. (Although hopefully we do not perform these actions at the same time). Our intuitions about these everyday actions are, I think, based on the assumption that the actions in question are objectively risky. So I take it that these intuitive judgments support this thesis:

> OBJECTIVE RISK TOLERANCE You may expose someone to a negligible objective risk of death in order to otherwise provide this person with a small benefit. (All else equal.)

(ii) The significance of the Repetition Argument is unaffected by the fact that RISK TOLERANCE is a thesis about objective risk

So our intuitions support OBJECTIVE RISK TOLERANCE. The Repetition Argument shows that this thesis is inconsistent with MANY-FEW. To satisfy yourself of this, suppose that the risks of bringing the chancy antidote in Poison are objective risks. If you repeat these objective risks enough times, you will near certainly end up with a distribution of many small benefits and one death. Appealing to this consideration, the Repetition Argument finds that OBJECTIVE RISK TOLERANCE is inconsistent with MANY-FEW.

(ii) Rejecting OBJECTIVE RISK TOLERANCE is problematic

So we cannot hold both OBJECTIVE RISK TOLERANCE and MANY-FEW. Suppose we give up OBJECTIVE RISK TOLERANCE. Our naive view of everyday risky actions is that they are objectively risky. If the naive view is correct, and we reject OBJECTIVE RISK TOLERANCE, then we lack a justification for these everyday actions.

But whether the naive view is correct will turn on some tricky issues in physics and metaphysics. Consider the following two claims. First, the effects of these everyday actions are determined. Second, if an event is determined, then it is not objectively chancy. Presently, there seems to be little consensus among the experts concerning whether each

of these claims is true.[80] For the sake of argument, let us suppose that they both are true. In this case, everyday actions are merely subjectively risky.

This might seem to present an escape for us. The Repetition Argument shows that we cannot hold OBJECTIVE RISK TOLERANCE and MANY-FEW. The reason why we want to hold onto OBJECTIVE RISK TOLERANCE is to justify everyday risky actions like giving people lifts in a car. We could reject OBJECTIVE RISK TOLERANCE and still hold that these everyday actions are permissible because the actions are merely subjectively risky.[81] Call this response, "The Escape."

Alas, The Escape does not work. In addition to the permissibility of an action, we should also consider whether we would be blameworthy for performing the action. The Escape is no defense against the charge that we are blameworthy for performing these everyday risky actions. Let me say why.

When people correctly apportion their credences on the basis of their available evidence, the blameworthiness of their actions depends on their evidence.[82] To see this, suppose Alf buys a packet of Domino's sugar, which he adds to someone's tea. The packet he buys contains an untraceable poison. Still, he acts blamelessly. This is because his evidence was that the sugar was harmless, and he based his beliefs on the evidence.

Because the blameworthiness of an action turns on an agent's *evidence*, an agent can be blameworthy for performing a subjectively risky action, even if the action is not objectively risky. Now, consider the following principle:

BLAME

You are blameworthy for taking a subjective risk of a certain magnitude just in case you are not permitted to take an objective risk of the same magnitude.[83]

This principle strikes me as very attractive. If this principle is correct, then OBJECTIVE RISK TOLERANCE is true just in case the following thesis is true:

---

[80] As things stand, it seems that the experts disagree about (i) whether the correct interpretation of quantum mechanics is indeterministic; (ii) whether indeterminism could affect macroscopic events like human actions; (iii) whether actions could be objectively chancy even on the assumption that determinism is true.

[81] This raises the question of whether we may perform merely subjectively risky actions. There is a controversy among moral philosophers about this issue. Some hold that the permissibility of these actions depends on your available evidence. See Thomas M. Scanlon, *Moral Dimensions*, (Cambridge MA: Belknap Press of Harvard University Press, 2009). Others deny that the permissibility of merely subjectively risky action depends on the agent's evidence. Instead, they say that the permissibility of such an action depends on the action's actual effects. See Judith Jarvis Thomson, *Normativity*, (Chicago and La Salle, IL: Open Court Press, 2008).

[82] I pass over complications that arise concerning the blameworthiness of an agent's conduct on the basis of a belief that is not backed by the available evidence.

[83] More precisely:

> You are blameworthy for performing an action that on the basis of your evidence you ought to assign a credence of p that it leads to harm, h, and a credence of (1 - p) that it leads to benefit, b, just in case you may not perform an action with an objective chance p that it leads to harm, h, and an objective chance (1 - p) that it leads to benefit b.

SUBJECTIVE RISK TOLERANCE   You are not blameworthy for exposing someone to a negligible subjective risk of death in order to otherwise provide this person with a small benefit. (All else equal.)

So if we give up OBJECTIVE RISK TOLERANCE, then we must also give up SUBJECTIVE RISK TOLERANCE. Therefore, if we take The Escape, we still have to concede that our everyday risky actions are blameworthy. In light of this concession, it would be a pyrrhic victory to claim that these actions are permissible.

Bibliography

Adams, Robert, 1984, "Saints," *The Journal of Philosophy* 81: 392 – 401

Anscombe, G. E. M., 1967, "Who is Wronged?" *The Oxford Review* no. 5: 16-17

Aristotle, 1999, *Nichomachean Ethics*, translated by Terence Irwin, (Indianapolis: Hackett Publishing Company, Inc.)

Baron, Marcia, 1987, "Kantian Ethics and Supererogation," *The Journal of Philosophy* 84: 237-262

Batson, C. Daniel and Shaw, Laura L., 1991, "Evidence of Altruism: towards a pluralism of pro-social motives," *Psychological Inquiry* 2: 107-122

Broome, John, 1990 - 1991, "Fairness," *Proceedings of the Aristotelian Society* New Series 91: 87-101

- 2005, "Does rationality give us reasons?" *Philosophical Issues* 15: 321–37

- manuscript, *Rationality Through Reasoning*

Buss, Sarah, 2006, "Needs (Someone Else's), Projects (My Own) and Reasons," *The Journal of Philosophy* 103: 373 – 402

Cummiskey, David, 1996, *Kantian Consequentialism*, (New York: Oxford University Press)

Doggett, Tyler, forthcoming, "Saving the Few," *Nous*

Doris, John, 2002, *Lack of Character: Personality and Moral Behavior*, (New York: Cambridge University Press)

Festinger, Leon, 1957, *A theory of cognitive dissonance*, (Stanford, CA: Stanford University Press)

Foot, Philppa, 1985, "Utilitarianism and the Virtues," *Mind New Series* 94: 196-209

Frick, Johann, manuscript, "Health Resource Allocation Behind a Natural Veil of Ignorance."

Geach, Peter, 1956, "Good and Evil," *Analysis* 17: 32-42

Hare, Caspar, 2009, "Perfectly balanced interests," *Philosophical Perspectives* 23: 165 - 176

Hassin, Ran, Uleman, James, and Bargh, John, (eds.), 2005, *The New Unconscious*, (New York: Oxford University Press)

Herman, Barbara, 2007, *Moral Literacy*, (Cambridge MA: Harvard University Press)

Hill Jr, Thomas E, 2002, *Human Welfare and Moral Worth*, (New York: Oxford University Press)

Hirose, Iwao, 2001, "Saving the Greater Number without Combining Claims," *Analysis* 61: 341-2

Hosein, Adam, 2009, *The Significance of Fairness*, Doctoral Thesis, M.I.T.

Hsieh, Nien-Hê, Strudler, Alan, and Wasserman, David, 2006, "The Numbers Problem,"

Philosophy and Public Affairs 34 : 352-372

- 2007, "Pairwise Comparison and Numbers Skepticism," *Utilitas* 19: 487-504

Kamm, Frances, 1993, *Morality, Mortality Vol. I: Death and Whom to Save From It*, (New York: Oxford University Press)

Kamm, Frances, 2005, "Aggregation and Two Moral Methods," *Utilitas* 17: 1-23

Kant, Immanuel, 1991, *The Metaphysics of Morals*, translated by Mary Gregor, (New York: Cambridge University Press)

Kant, Immanuel, 1998, *Groundwork of the Metaphysics of Morals*, translated by Mary Gregor, (New York: Cambridge University Press)

Kavka, Gregory, 1979, "The Numbers Should Count," *Philosophical Studies* 36 (1979): 285 – 294

Kumar, Rahul, 2001, "Contractualism on Saving the Many," *Analysis* 61: 165 – 170

Lenman, James, 2008, "Contractualism and risk imposition," *Politics Philosophy Economics* 7: 99 - 122

Lewis, David, 1980, "A Subjectivist's Guide to Objective Chance," in Richard C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, 1980, (Berkeley CA: University of California Press).

Loewer, Barry, 2001, "Determinism and Chance", *Studies in History and Philosophy of Modern Physics* 32B(4) : 609-620

Lübbe Weyma, 2008, "Taurek's No Worse Claim," *Philosophy and Public Affairs* 36: 69-85

Meyer, Kirsten, 2006, "How to be Consistent without Saving the Greater Number," *Philosophy & Public Affairs* 34: 136-146

Muraven, Mark, and Baumeister, Roy F., 2000, "Self-regulation and depletion of limited resources: Does self-control resemble a muscle?" *Psychological Bulletin* 126: 247-259

Murphy, Liam, 2003, *Moral Demands in Nonideal Theory*, (Oxford: Oxford University Press)

Norcross, Alastair, 1998, "Great harms from small benefits grow: how death can b e outweighed by headaches," *Analysis* 58: 152 - 158

Otsuka, Michael, 2004, "Skepticism about saving the greater number," *Philosophy & Public Affairs*, 32: 413-26

   - 2006, "Saving Lives, Moral Theories and the Claims of Individuals," *Philosophy & Public Affairs* 34: 109-135

   - manuscript, "Life and Limb"

Parfit, Derek, 1987, *Reasons and Persons*, (reprinted with further corrections), (New York: Oxford University Press Inc.)

   - 2003, "Justifiability to Each Person", *Ratio (new series)* 16: 368 - 390

   - manuscript, *On What Matters*

Raz, Joseph, 2003, "Numbers, with and without Contractualism," *Ratio (new series)* 16: 346-367

Scanlon, Thomas M., 1998, *What We Owe to Each Other*, (Cambridge MA: Harvard University Press).

- 2009, *Moral Dimensions*, (Cambridge MA: Belknap Press of Harvard University Press)

Singer, Peter, 1972, "Famine, Affluence and Morality," *Philosophy and Public Affairs*, 1: 229 – 243

Stuart Mill, John Stuart, 1998, *Utilitarianism*, ed. Roger Crisp, (New York: Oxford University Press, originally published 1861)

Taurek, John, 1977, "Should the Numbers Count?" *Philosophy & Public Affairs*, 6: 293-216

Temkin, Larry, 2003, "Egalitarianism Defended," *Ethics* 113: 764-782

Thomson, Judith Jarvis, 1971, "A Defense of Abortion," *Philosophy and Public Affairs* 1: 47-66

- 1993, "Goodness and Utilitarianism," *Proceedings and Addresses of the American Philosophical Association* 67: 145-159

- 2008, *Normativity*, (Chicago and La Salle, IL: Open Court Press)

Thomson, Judith Jarvis, and Harman, Gilbert, 1996, *Moral Relativism and Moral Objectivity*, (Cambridge MA: Blackwell Publishers Inc.)

Tversky, Amos, and Kahneman, Daniel, 1974, "Judgment under uncertainty: Heuristics and biases," *Science* 185: 1124–1131

Urmson, J. O., 1958, "Saints and Heroes," p. 64, in *Essays in Moral Philosophy*, ed. A. I. Melden,, (University of Washington Press)

Wasserman, David, and Strudler, Alan, 2003, "Can a Nonconsequentialist Count Lives?" *Philosophy & Public Affairs* 31: 71-94

Wolf, Susan, 1982, "Moral Saints," *The Journal of Philosophy*, 79: 419–439