

TWO DOGMAS OF DEONTOLOGY: AGGREGATION, RIGHTS, AND THE SEPARATENESS OF PERSONS

BY ALASTAIR NORCROSS

I. INTRODUCTION: THE SEPARATENESS DOGMA DESCRIBED

One of the currently popular dogmata of anticonsequentialism is that consequentialism does not respect, recognize, or in some important way account for what is referred to as the “separateness of persons.” William Shaw, a widely read philosopher at San José State University, refers to this charge as a “virtual mantra.”¹ The charge is often made, but rarely explained in any detail, much less argued for. In this essay, I will explain what I take to be the most plausible interpretation of the separateness of persons charge. I will argue that the charge itself can be deconstructed into at least two further objections to consequentialist theories. Of these two objections, I will argue that the first one, though often made, is untenable. I will also argue that the second objection, in its various forms, relies on distinctions whose moral significance is vigorously denied by almost all consequentialist theorists. I will thus argue that the separateness of persons objection poses no special threat to consequentialism.

I begin my examination of the separateness of persons dogma with its classical expression in the 1960s and 1970s. Most philosophers trace its origin to John Rawls. In *A Theory of Justice* (1971), he had this to say:

The most natural way, then, of arriving at utilitarianism . . . is to adopt for society as a whole the principle of rational choice for one man. . . . On this conception of society separate individuals are thought of as so many different lines along which rights and duties are to be assigned and scarce means of satisfaction allocated . . . so as to give the greatest fulfillment of wants. The nature of the decision . . . is not, therefore, materially different from that of an entrepreneur deciding how to maximize his profit . . . or that of a consumer deciding how to maximize his satisfaction by the purchase of this or that collection of goods. . . . This view of social co-operation is the consequence of extending to society the principle of choice for one man, and then, to make this extension work, conflating all persons into one through the

¹ William Shaw, *Contemporary Ethics: Taking Account of Utilitarianism* (Oxford: Blackwell, 1999), 128.

imaginative acts of the impartial sympathetic spectator. Utilitarianism does not take seriously the distinction between persons.²

Although Rawls popularized this charge, whatever it turns out to be, other philosophers were saying similar things around the same time. Indeed, nine years before the publication of *A Theory of Justice*, David Gauthier wrote that theories like utilitarianism

suppose that mankind is a super-person, whose greatest satisfaction is the objective of moral action. . . . But this is absurd. Individuals have wants, not mankind; individuals seek satisfaction, not mankind. A person's satisfaction is not part of any greater satisfaction.³

Likewise, Thomas Nagel claimed that consequentialism

treats the desires, needs, satisfactions, and dissatisfactions of distinct persons as if they were the desires, etc., of a mass person.⁴

We even find Rawls's erstwhile colleague and philosophical antagonist Robert Nozick jumping on the separation anxiety bandwagon:

There is no *social entity* with a good that undergoes some sacrifice for its own good. There are only individual people, different individual people, with their own individual lives. Using one of these people for the benefit of others, uses him and benefits the others. Nothing more. What happens is that something is done to him for the sake of others. Talk of an overall social good covers this up. (Intentionally?) To use a person in this way does not sufficiently respect and take account of the fact that he is a separate person, that his is the only life he has.⁵

II. A METAPHYSICAL MISTAKE?

So what is going on in all these complaints? There seem to be two possible interpretations of the charge that utilitarianism does not respect the separateness of persons. Utilitarianism is making either a metaphysical or a moral mistake. I will briefly consider the former possibility, but will concentrate on the latter, since it is clearly the only one that makes (even a little) sense.

² John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), 26–27.

³ David Gauthier, *Practical Reasoning: The Structure and Foundations of Prudential and Moral Arguments and Their Exemplification in Discourse* (Oxford: Oxford University Press, 1962), 126.

⁴ Thomas Nagel, *The Possibility of Altruism* (Oxford: Oxford University Press, 1970), 134.

⁵ Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), 32–33.

What would the metaphysical version of the charge be? Are anti-consequentialists really saying that consequentialists believe that there is a “super-person,” composed of everyone, with the same ontological status as a normal person? Though I am not exactly sure what this ontological commitment would be, I am pretty sure that most consequentialists do not have it. There is, perhaps, one notable contemporary exception. It is possible to get close to something that might be in the ballpark of this ontological commitment by adopting Derek Parfit’s approach to personal identity. Since the relations of psychological connectedness between our mental states that are important to us may vary and weaken over time, the connections between widely separated temporal stages of a person may be no stronger than the connections between contemporary stages of different persons. Parfit writes:

If we cease to believe that persons are separately existing entities, and come to believe that the unity of a life involves no more than the various relations between the experiences in this life, it becomes more plausible to be more concerned about the quality of experiences, and less concerned about whose experiences they are. This gives support to the Utilitarian view. . . . Since persons are not separately existing entities, the impersonality of Utilitarianism is less implausible.⁶

Parfit’s view on the nature of persons is interesting and challenging, and may well be correct. Parfit may also be correct in his claim that his view gives support to utilitarianism. For the sake of argument, let’s go even further, and suppose that Parfit’s view *entails* utilitarianism. Nonetheless, it is clear that it is not entailed *by* utilitarianism or other consequentialist theories. Let me explain. It is obviously possible to claim that the well-being of every sentient being counts equally, and that the consequences of actions for the aggregate well-being of all are the sole determinants of the moral status of actions, without being committed to the view that there is a super-being composed of all sentient beings. If Parfit’s view is false, any theory that entails the denial of the metaphysical separateness of persons is in trouble. However, as I demonstrated above, neither utilitarianism specifically, nor consequentialism more generally, does entail the denial of the metaphysical separateness of persons. Furthermore, Parfit’s view of the nature of persons is very much a minority view among consequentialists. Since Rawls, Gauthier, Nagel, and Nozick are (or were) smart and well-read philosophers, it is highly unlikely that they were unaware that consequentialism as such is not committed to the denial of the metaphysical separateness of persons. Even though the wording of their criticisms, especially those of Gauthier and Nozick quoted above, suggests the meta-

⁶ Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), 346.

physical version of the separateness of persons charge against consequentialism, a more charitable interpretation of their complaints suggests otherwise. I turn, then, in the next section to the more plausible version of the charge that consequentialist theories such as utilitarianism ignore the separateness of persons.

III. A MORAL MISTAKE: THE SEPARATENESS DOGMA DECONSTRUCTED

The most promising interpretation of the separateness complaints is that utilitarianism is making a moral mistake by treating the relationship between different people as similar to the relationship between different temporal stages (or even different cotemporal aspects) of one person, at least for the purposes of moral decision-making. Peter Vallentyne thus emphasizes what he calls “the *normative* separateness of persons,” and explains the trouble this makes for utilitarianism as follows:

[I]ndividuals have certain rights that may not be infringed simply because the consequences are better. Unlike prudential rationality, morality involves many distinct centers of will (choice) or interests, and these cannot simply be lumped together and traded off against each other.

The basic problem with standard versions of core consequentialism is that they fail to recognize adequately the *normative separateness of persons*. Psychological autonomous beings (as well, perhaps, as other beings with moral standing) are not merely means for the promotion of value. They must be *respected* and *honored*, and this means that at least sometimes certain things may not be done to them, even though this promotes value overall. An innocent person may not be killed against her will, for example, in order to make a million happy people significantly happier. This would be sacrificing her for the benefit of others.⁷

There seem to be several distinct ideas here: (i) Individuals have rights, which at least sometimes trump utility calculations. (ii) Individuals’ interests cannot simply be traded off against each other. (iii) Individuals must be *respected* or *honored*. Consider these claims in reverse order. A utilitarian may claim, with some justification, that the demand for equal consideration of interests embodied in the utilitarian theory (and other consequentialist theories) is precisely what it means to respect or honor individuals. It is only when I weigh your interests equally with the interests of all others whom I can affect that I adequately respect or honor you.

⁷ Peter Vallentyne, “Against Maximizing Act Consequentialism,” in James Dreier, ed., *Contemporary Debates in Moral Theory* (Oxford: Blackwell, 2006), 29; emphasis in the original.

Deontological constraints function to disallow the consideration of certain interests in certain circumstances. Thus, at least sometimes, they prevent us from respecting or honoring certain individuals.

At this point, the critic of utilitarianism will no doubt claim that I have (perhaps willfully) misunderstood (iii). In fact, he might claim that (i) and (ii) explain *what it means* to honor or respect individuals. Let us move, then, to claim (ii).

IV. THE REJECTION OF THOROUGHGOING AGGREGATION

Claim (ii) denies the aggregative feature of utilitarianism. For a utilitarian, the misfortunes of some can be outweighed by the fortunes of others. What gets a nonconsequentialist more agitated than a Promise Keeper in a Women's Studies class⁸ is the claim that the great misfortunes of a few could be outweighed by the fairly trivial fortunes of many. To use the standard example, aggregation seems to commit the utilitarian to the claim that the death of one innocent person could be outweighed by the relief of a sufficiently large number of minor headaches, and thus also to the claim that it could be permissible to kill an innocent person in order to relieve that number of headaches. Thus, we must deny aggregation.

The problem with the denial of trade-offs or aggregation is that even committed anticonsequentialists accept them in many circumstances. For example, suppose that Homer is faced with the painful choice between saving Barney from a burning building or saving both Moe and Apu from the building. Unless we want to follow John Taurek⁹ (and probably Philippa Foot,¹⁰ and possibly Judith Thomson)¹¹ to the funny farm,¹² we must

⁸ For those unfamiliar with the Promise Keeper phenomenon in the United States, a little explanation is in order. The Promise Keepers is an evangelical Christian organization for men. It teaches, among other things, that within marriage the man should be the head of the household and the woman should willingly submit to his leadership. The organization has been criticized by feminist groups in the U.S., such as the National Organization for Women, for (allegedly) encouraging inequality within marriages and teaching male superiority. Such doctrines rarely receive a sympathetic hearing in Women's Studies classes.

⁹ See John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6 (1977): 293–316.

¹⁰ See Philippa Foot, "Utilitarianism and the Virtues," *Mind* 94 (1985): 196–209.

¹¹ See Judith Jarvis Thomson, "The Right and the Good," *The Journal of Philosophy* 94, no. 6 (1997): 273–98.

¹² Taurek argues for the radical thesis that there is no moral reason to prefer the death of one person over the deaths of five persons (or even of five million persons). Foot argues that no sense can be made of one state of affairs being overall better than another from the perspective of morality. Foot's position seems to be in agreement with Taurek's in the following sense: it rejects the claim that I have a moral reason to prefer the death of one to the deaths of five, if that reason is supposed to be grounded in the claim that it is overall better that only one person dies than that five persons die. Thomson's position is less clear. She criticizes utilitarianism for its reliance on comparative judgments of the goodness of states of affairs, and in this respect seems to be sympathetic to Foot's position. However, a charitable reading of her essay (which she would no doubt reject) renders it as a defense of rule utilitarianism.

admit that it is clearly better for Homer to save the larger number, *precisely because it is a larger number*. The proponent of claim (ii) might try to accommodate this intuition by limiting the scope of trade-offs. For example, perhaps we are allowed to trade lives for lives (or similarly serious harms), but we are not allowed to trade lives for convenience. Homer can save the lives of Moe and Apu rather than Barney, but he cannot leave Barney to die in order to provide all the inhabitants of Springfield with a few minutes of extra free time every day.¹³ Thomas Scanlon tries such a move in his attempt to accommodate limited aggregation in his contractualist theory:¹⁴

[I]t seems that our intuitive moral thinking is best understood in terms of a relation of “relevance” between harms. If one harm, though not as serious as another, is nonetheless serious enough to be morally “relevant” to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even “relevant to,” some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm.¹⁵

Scanlon rightly sees that it would be highly implausible to limit trade-offs to harms of exactly equal seriousness. It is clearly better that one person suffer some particular harm than that ten people suffer a harm that is only slightly less serious. However, Scanlon’s attempt, or any similar attempt, to limit the scope for trade-offs faces some serious problems. First, it is fairly clear that the relation of moral relevance does not obey the principle of transitivity. To see why, suppose, first, that it does. Consider now a descending scale of finitely many different harms, from the most serious, such as death, all the way down to the most trivial, such as a minor temporary headache. The difference in seriousness between any two adjacent harms is no larger than is necessary for the lesser harm to be clearly less serious than the greater harm. Suppose, also, that for every harm on the scale above the most trivial, there is some lesser harm that is relevant to it. Call this second assumption the “continuity assumption.” Transitivity and continuity together entail that the most trivial

¹³ I leave the reader to fill in the details of this and other examples involving the endlessly fascinating inhabitants of Springfield.

¹⁴ See Thomas Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998). For a more comprehensive critique of Scanlon’s attempts to accommodate limited aggregation, see Alastair Norcross, “Contractualism and Aggregation,” *Social Theory and Practice* 28, no. 2 (2002): 303–14.

¹⁵ Scanlon, *What We Owe to Each Other*, 240.

harm is relevant to the most serious harm, precisely the result that the notion of moral relevance is intended to avoid. Can we preserve transitivity by rejecting continuity? This would involve finding a break (or breaks) in the scale between two harms, such that the harm directly below the break is not morally relevant to the harm directly above the break. Given that the difference between any two adjacent harms is as small as is compatible with the harms being morally distinct, the postulation of a break in the scale would run directly counter to the intuition that suggested the notion of moral relevance in the first place. Where could such a break plausibly occur? The most likely candidate would be just below death. There is, we might think, something special about death. As Clint Eastwood says in the film *Unforgiven*, "It takes away all a man has, and all he's gonna have." Unpleasant as even severe mutilation is, perhaps it is still worse that one person die than that any number are mutilated. This might be the view of death espoused by those students in introductory classes who claim that life is "invaluable" or "infinitely valuable," but is it really plausible? Can anyone who really considers the matter seriously honestly claim to believe that it is worse that one person die than that the entire sentient population of the universe be severely mutilated? Clearly not. Perhaps the break in the sequence of harms could occur at some later point. Perhaps there is some harm short of death that is worse than any number of any lesser harms. This seems even more implausible, though, than the claim that death is worse than any number of any lesser harms.

We must, therefore, conclude that the relation of moral relevance, if it is to do the work intended for it by Scanlon, does not obey strict transitivity. So what? If the notion of moral relevance were supposed to constrain our judgments of all-things-considered betterness, this would be a serious problem. Although some brave souls have seriously entertained the possibility that "all-things-considered better than" is not a transitive relation,¹⁶ the sheer implausibility of the suggestion makes the standard objections to utilitarianism, Kantianism, or contractualism appear trivial by contrast. However, Scanlon suggests the notion of moral relevance as part of an account of what principles are reasonably rejectable, and thus

¹⁶ For examples of the attempt to deny transitivity for "all-things-considered better than," see Larry Temkin, "Intransitivity and the Mere Addition Paradox," *Philosophy and Public Affairs* 16 (1987): 138-87; Larry Temkin, "A Continuum Argument for Intransitivity," *Philosophy and Public Affairs* 25 (1996): 175-210; Warren Quinn, "The Puzzle of the Self-Torturer," *Philosophical Studies* 59 (1990): 79-90; and Stuart Rachels, "Counterexamples to the Transitivity of 'Better Than,'" *Australasian Journal of Philosophy* 76, no. 1 (1998): 71-83. Temkin's 1996 article uses the same central example as Rachels's article, but Temkin's explanation for the supposed intransitivity is the same as the one he provides in his 1987 article. Quinn does not explicitly claim that "better than" is intransitive, but his arguments, if successful, entail that a utilitarian should deny the transitivity of "better than." I discuss Temkin's 1987 article in Alastair Norcross, "Intransitivity and the Person-Affecting Principle," *Philosophy and Phenomenological Research* 59, no. 3 (1999): 769-76. I discuss Temkin's 1996 article and Quinn's 1990 article in Alastair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135-67.

of which options are permissible, obligatory, or forbidden. To demonstrate that, even in this context, the failure of transitivity leads to highly implausible results, I need to consider an example.

Suppose, for the sake of argument, (a) that the loss of both arms is less serious than but morally relevant to death; (b) that a broken leg is less serious than but morally relevant to the loss of both arms, but not morally relevant to death; (c) that in a choice between saving one life and preventing one thousand people from losing both arms, it is obligatory to aid the larger group; and (d) that in a choice between preventing one thousand people from losing both arms and preventing one million people from breaking a leg, it is obligatory to aid the larger group. (The choice of examples is unimportant.) Consider now three different choices: (1) Save one life or prevent one thousand people from losing both arms. (2) Prevent one thousand people from losing both arms or prevent one million people from breaking a leg. (3) Save one life or prevent one million people from breaking a leg. From (b), (c), and (d), it follows that it is obligatory to aid the larger group in (1) and (2), and the smaller group in (3). So far, so good. But what happens when we are faced with all three options in one choice? Should we choose to save one life, or to prevent one thousand people from losing both arms, or to prevent one million people from breaking a leg? No answer here seems satisfactory.

Consider the possibility that one of the options (say, saving the life) is obligatory. But now suppose that, just as you are about to save the life, it becomes impossible for you to prevent the million people from breaking a leg. Perhaps the largest group is farther away than the other two, and your fuel tank is punctured by a jagged rock on the road to the one person. You are still able to save either the one or the thousand, but you cannot reach the million in time. Now you find yourself faced with choice (1), in which it is obligatory to save the thousand and forbidden to save the one. But this is very strange. You were about to do your duty, virtuously eschewing both forbidden alternatives, when one of the forbidden alternatives by chance became unavailable, as a result of which the other forbidden alternative became obligatory, and the previously obligatory alternative became forbidden. We should, if at all possible, avoid having to swallow such an unpalatable consequence. The same reasoning applies, *mutatis mutandis*, to the hypothesis that either of the other alternatives is obligatory in the three-option choice. Perhaps, then, each option is permissible in the three-option choice. But the implausibility of this can be demonstrated by the very same thought experiment. You are about to perform the perfectly permissible act of saving a life, when one of your other permissible alternatives becomes unavailable by chance. Now it is no longer permissible to save the life. A further possibility is that each option is forbidden in the three-option case. But this is even more unpalatable than the previous suggestions. Not only would we have to accept that a previously forbidden alternative can become obligatory by the

chance deletion of another forbidden alternative, but we would also have to accept the existence of situations in which an agent, through no fault of her own, cannot help but do wrong. What is more, such situations may be very common. Both through the agency of charities and through our own efforts, many of us are able to bring many different types and levels of aid to others.

V. AXIOLOGICAL AGGREGATION

Up until this point, I have been ignoring a potentially important ambiguity in describing the anticonsequentialist rejection of aggregation. The consequentialist seems to be committed to both axiological and deontic aggregation. For those unfamiliar with this terminology, a little rough explanation is in order (the rest can skip this bit). Axiology is the study of value. It is concerned with theories of the good, and what makes for a good state of affairs. Deontology is the study of duty¹⁷ and is concerned with questions about what choices are required, forbidden, or permitted. Axiological aggregation involves the claim that harms and benefits can be traded off against each other in determining the overall goodness (or badness) of a state of affairs. Deontic aggregation involves the claim that harms and benefits can be traded off against each other in determining which choices are required, permissible, or forbidden. Given the structure of consequentialist theories, a commitment to axiological aggregation entails a commitment to deontic aggregation. If a state of affairs with a large number of small benefits and a small number of large harms is better, *ceteris paribus*, than one with a small number of large benefits and a large number of small harms, then it will always be at least permissible to choose the former over the latter.

However, other ethical approaches may, at least in theory, separate the question of axiological aggregation from that of deontic aggregation. If we accept at least the limited axiological aggregation that even nonconsequentialists acknowledge (for example, preferring fewer deaths to more deaths, or preferring small amounts of some harm to much greater amounts of a slightly less serious harm), is there any other plausible way to block the unrestricted axiological aggregation that the separateness of persons objection finds so disagreeable? At first glance, it appears the answer is no. Recall the continuity assumption. It seems highly plausible that there are misfortunes that are worse than mild headaches, that nonetheless can be individually outweighed by a sufficient number of mild headaches. This is relatively uncontroversial. A mild ankle sprain is a good candidate for such a misfortune. Likewise, it is pretty clear that there are misfor-

¹⁷ Confusingly, “deontology” is also the catch-all name for the most common family of anticonsequentialist moral theories. This is confusing because, understood as the study of duty, deontology also encompasses most versions of consequentialism, which do, after all, provide an account of moral duty.

tunes that are worse than mild ankle sprains, that nonetheless can be individually outweighed by a sufficient number of mild ankle sprains. Perhaps a broken ankle is such a misfortune. Even though it is worse that one person break her ankle than that she mildly sprain it, it is worse that many people have mild ankle sprains than that one has a broken ankle. But this process of escalation can be continued. For each misfortune short of the worst possible one, there is a worse misfortune that can be individually outweighed by a sufficient number of the lesser one. In particular, it seems plausible that there is some misfortune short of death, perhaps some kind of mutilation, that can, if suffered by enough people, outweigh one death. Consider now a sequence of judgments, S , that begins as follows: one death is better than n^1 mutilations; n^1 mutilations are better than n^2 X s (where X is some misfortune less bad than mutilation). S continues with the first term of each comparison being identical to the second term of the previous comparison, until we reach the last two comparisons: n^{m-2} broken ankles are better than n^{m-1} mild ankle sprains; n^{m-1} mild ankle sprains are better than n^m mild headaches. If we have S , we can conclude, by the transitivity of "better than," that one death is better than n^m mild headaches. Thus, unrestricted aggregation seems to be the only alternative to denying the transitivity of "all-things-considered better than."

VI. THE ASYMPTOTIC GAMBIT TO BLOCK AXIOLOGICAL AGGREGATION

But perhaps I have been too hasty. In reply to an essay in which I appealed to the reasoning of the previous section to argue that consequentialists are indeed committed to the claim that some number of mild headaches is worse than one death, Erik Carlson produced an ingenious argument to the contrary.¹⁸ He suggests that something like the principle of diminishing marginal utilities might apply to harms themselves, and, furthermore, that there might be an upper bound to the cumulative disvalue produced by aggregating any particular kind of harm. Perhaps each type of harm, when aggregated, would move asymptotically toward its upper bound. Thus, for example, ten mild headaches might not be ten times as bad as one mild headache, and there may be no number of mild headaches whose aggregate disvalue is one hundred times worse than one. The exact details of the suggestion are not important, but so long as the upper bound on the disvalue of headaches falls short of the disvalue of one death, the continuity assumption will not license the postulation of a true S , the sequence that takes us all the way from a death down to headaches. To see why, consider a simplified sequence that satisfies the

¹⁸ See Erik Carlson, "Aggregating Harms—Should We Kill to Avoid Headaches?" *Theoria* 66, no. 3 (2000): 246–55.

continuity assumption and the diminishing utility suggestion. Suppose that there are just three types of harm—headache, mutilation, and death—having disvalues of 1, 10, and 100, respectively. Suppose further that each type of harm has an upper bound on its aggregate disvalue of 15 times its individual disvalue. Thus, no number of headaches will have an aggregate disvalue of more than 15; no number of mutilations will have an aggregate disvalue of more than 150; and no number of deaths will have an aggregate disvalue of more than 1,500. It is true that some number of mutilations is worse than one death, and that some number of headaches is worse than one mutilation; but no number of headaches is worse than one death.

As I said, this argument is ingenious. It is also unsuccessful. The first thing to note is that the suggestion cannot simply be an application of the commonly accepted principle of diminishing marginal utility. It may be, as a causal matter, that further headaches, or the further duration of a headache, becomes less unpleasant for the sufferer. My argument, though, concerns many minor unpleasant experiences, which are all (at least roughly) equally unpleasant. Furthermore, the anticonsequentialist intuition is that one death is worse than any number of minor headaches, even (or especially?) when spread out among any number of different people. It is highly implausible to suggest that the headache of the trillionth person is somehow less *unpleasant* than the otherwise identical headache of the first person, just because a whole bunch of other people have already experienced one (or are currently experiencing one). So the suggestion must be that experiences that are equally unpleasant diminish in *badness* as they are aggregated. That is, the disutility of x headaches is less than x times the disutility of one headache.

Further, and this is a question raised by Carlson in discussing his suggestion, do we assume that the badness of harms diminishes cumulatively, starting with the first such harm ever experienced, or does the diminishing start fresh with each choice? Either option is unacceptable. If the former, then your current headache may be only infinitesimally bad (if there is an upper bound to the total possible headache badness, as there would have to be to counter my argument), because countless people before you have experienced headaches. What is even more absurd, with a fine-grained enough categorization of harms, one fairly trivial harm could be worse than an intuitively much more serious one. Let me explain. Suppose that a particular mildly unpleasant nasal itch has never before been experienced. Perhaps it is only caused by a rare combination of English and French cuisine that no chef has yet been brave (or foolish) enough to attempt. Suppose the first such itch has a disutility of 2, with an upper bound of 200. Now suppose that the only way to prevent someone, say Mary, from losing a leg involves producing the first instance of the nasal itch. Let's say that the loss of a leg has an initial disutility of 300, with an upper bound of 30,000. Intuitively, it is much worse to lose

a leg than to experience the mildly unpleasant nasal itch. But if we apply the cumulative version of the diminishing utility suggestion, we get a strange result. Suppose that, in the course of history, so many people have lost legs that the cumulative disutility has passed 29,999. Now, each additional loss has a tiny disutility, well below 1. It now appears that it is worse for Mary to suffer the nasal itch than to lose a leg. But this is clearly absurd. We must, then, consider the version of the diminishing utility suggestion in which the diminishing starts fresh with each choice.

But this version also leads to unacceptable results. Consider again the simplified spectrum of harms encompassing just headaches, mutilations, and deaths. Recall that the initial disutility of a mutilation is 10, with an upper bound of 150, and the initial disutility of a death is 100, with an upper bound of 1,500. Suppose that the aggregate disutility of two deaths is 199, and the aggregate disutility of two thousand mutilations is 149. According to the current interpretation of Carlson's suggestion, it is clearly worse that two people die than that two thousand people are mutilated, and thus one should choose the two thousand mutilations over the two deaths, if faced with the choice. But suppose also that the aggregate disutility of one thousand mutilations is 145. In this case, it is clearly worse that one thousand people are mutilated than that one person dies, and thus one should choose the one death over the one thousand mutilations, if faced with the choice. But now we are faced with the ridiculous possibility that we could reverse moral judgments by splitting one choice into two. Whether it is better to kill two than to mutilate two thousand could depend on whether one could first choose between killing one and mutilating one thousand, and then choose between killing the other and mutilating the remaining one thousand. Let us add a few details to the example. Suppose that two thousand people are in danger of suffering mutilation from a disease. However, if two other people, currently trapped in a mineshaft, die, a cure can be synthesized from their bodies that will prevent the two thousand mutilations. You can save the two in the mineshaft by pressing a button in front of you, or you can let them die. You cannot save only one. If you let them die, the two thousand mutilations will be prevented. What is the better course of action? On Carlson's suggestion, the answer would seem to depend on the details of how the cure will be synthesized from the two bodies. If each body can provide a cure for one thousand people, the better course is to choose the two deaths. In effect, you are twice choosing between one death and one thousand mutilations. However, if each body provides half the cure for all two thousand people (and half a cure without the other half does no good), the better course is to choose the two thousand mutilations. In this case, you are choosing one time between two deaths and two thousand mutilations. This result is, as I said, absurd. If you set out to bring about the best state of affairs in all your choices, your decisions could differ depending on whether you were able to split your choices up, or perhaps

simply to think of your choices as split up. It is interesting to note that a similar objection applies to Judith Thomson's attempt to limit trade-offs between rights and utility in *The Realm of Rights*.¹⁹ Carlson's suggestion, then, is technically ingenious, but morally indefensible.

Before we leave the topic of axiological aggregation, it is worth remembering that we commonly accept trade-offs between lives and much lesser values, such as convenience. For example, we allow public projects such as building a bridge in order to make travel between two places more convenient, even when we know that several people will die in the course of the construction. Likewise, even most anticonsequentialists do not demand that highway speed limits be lowered to the optimal point for saving lives, even though the advantage of higher speed limits is merely increased convenience for many motorists.²⁰

VII. BLOCKING DEONTIC AGGREGATION: RIGHTS

It appears that the anticonsequentialist must read the writing on the wall and accept the permissibility of unrestricted axiological aggregation. Any harm, no matter how serious, can, in theory, be outweighed by a sufficiently large number of trivial benefits. If the separateness of persons objection is to pose a serious problem for utilitarianism, it must reject unrestricted deontic aggregation; that is, it must postulate deontic restrictions that deny the permissibility of always bringing about the greatest aggregate utility. Vallentyne's claim that individuals have rights that sometimes trump utilities (discussed in Section III above) is probably the most common attempt to do this. Utilitarianism is criticized for failing to distinguish between the following pair of cases (adapted from examples first concocted by Philippa Foot):²¹ In *Springfield Rescue I*, Homer must choose whether to save Barney, who is trapped on one side of Springfield, or both Moe and Apu, who are trapped on the other side. He cannot save all three, and no one else can save any of them. In *Springfield Rescue II*, Homer, and no one else, can save both Moe and Apu, who are trapped on the edge of Springfield Gorge. However, in order to reach them in time to save them, he must run over and kill Barney, who is trapped on a narrow

¹⁹ See Alastair Norcross, "Rights Violations and Distributive Constraints: Three Scenarios," *The Pacific Philosophical Quarterly* 76, no. 2 (1995): 159–67. Thomson's argument appears in Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, MA: Harvard University Press, 1990), 166–67. Thomson claims that we may be justified in violating certain rights, if the violation produces enough good, but only so long as the good produced is distributed appropriately. In particular, the good cannot be the sum of very tiny increments of good for a large number of people. Thomson describes this as the thesis that "where claims are concerned, the sum of goods across people does not count. . . . In still shorter form, where claims are concerned, the numbers do not count."

²⁰ For detailed discussion of both these points, see Norcross, "Comparing Harms."

²¹ Philippa Foot, "Killing and Letting Die," reprinted in Bonnie Steinbock and Alastair Norcross, eds., *Killing and Letting Die*, 2d ed. (New York: Fordham University Press, 1994), 280–89.

segment of the only road leading to the gorge. We are supposed to agree that Homer may choose to save Moe and Apu in *Springfield Rescue I*, but not in *II*. If he saves Moe and Apu in *II*, he will violate Barney's right not to be killed. Respecting the separateness of persons requires that we regard Barney's right not to be killed as trumping the utility of saving Moe and Apu. But why should we think that Barney has the right not to be killed, but that Moe and Apu don't have the right to be saved? And if Moe and Apu do have the right to be saved, why isn't it worse to violate their rights—there are, after all, two of them—than to violate Barney's right? More generally, for any right the nonconsequentialist might propose as trumping utility, why isn't there a corresponding right that would tell in favor of promoting utility?

When Philippa Foot first presented her more boring versions of the previous examples,²² she did so as part of an attempt both to explain the distinction between killing and letting die, and to argue for the moral significance of the distinction. She presents a distinction that focuses on the question of whether someone is "the agent" of harm to someone else. When the harm in question is death, this distinction corresponds roughly to the killing/letting die distinction. The distinction, according to Foot, is between originating or sustaining a fatal sequence, on the one hand, and allowing such a sequence to run its course, on the other. It is often permissible, she claims, to bring about a harm by the latter method that could not permissibly be brought about by the former. What explains this moral difference? The different types of agency receive their moral significance via their connection with different types of rights:

For there are rights to noninterference, which form one class of rights; and there are also rights to goods or services, which are different. . . . Typically, it takes more to justify an interference than to justify the withholding of goods or services.²³

Originating or sustaining a harmful sequence will usually involve the violation of a right to noninterference, whereas allowing such a sequence to run its course will, at most, involve the violation of a right to goods or services. The former type of right is stronger than the latter, so the former type of agency is less likely to be permissible than the latter. According to Foot, then, if Moe and Apu do have the right to be saved, it is not as important (strict, stringent, etc.) as Barney's right not to be killed. In general, if the rights view is to present a genuine alternative to conse-

²² For insomniacs, Foot's versions are titled "Rescue I" and "Rescue II." In Rescue I, we can save either five people in danger of drowning in one place or one person drowning somewhere else. In Rescue II, we can save the five drowning people only by driving over and killing someone who is trapped on the road. They appear in Foot, "Killing and Letting Die."

²³ *Ibid.*, 284.

quentialism, negative rights not to be harmed in some way must be stronger than the corresponding positive rights, if any, to be aided in avoiding such harm. More specifically, the duty not to harm in a certain way must be stricter than the corresponding duty to prevent such harm.

Foot's claim about the relative strictness of positive and negative rights and duties has a good deal of intuitive support. My right not to be poisoned, for example, does *seem* stronger than my right, if any, to be given the food I need to survive. Recall, though, that Foot is trying to explain the moral significance of the killing/letting die distinction. Clearly, though, as an explanation of a morally significant difference between killing and letting die, this appeal to different types of rights simply diverts the question. If the moral difference between positive and negative rights is to provide a satisfactory explanation of the moral difference between killing and letting die, we also need an explanation of the former difference. *Why* is my right not to be poisoned stronger than my right to be given the food I need to survive? Why, in general, does it "take more to justify an interference than to justify the withholding of goods or services"? The answer that springs most readily to mind is that it is worse, in general, to harm me than to fail to benefit me—worse, in particular, to kill me than to let me die. This answer, of course, can be of no help to Foot's approach, since it merely takes us in a circle. She appeals to the relative stringency of different types of rights to explain the moral significance of the distinction between killing and letting die. She obviously cannot then appeal to the moral significance of the distinction between killing and letting die as part of an account of the relative stringency of different types of rights. Of course, she does not commit this error. She simply *asserts* that negative duties are more stringent than positive duties, and thus that it takes more to justify an interference than a withholding. This assertion, completely lacking in argumentative support though it is, does help us to locate what is left of the separateness of persons dogma after the denial of axiological aggregation has been rightfully jettisoned, and to reveal what is really at issue between the dogma's followers and consequentialists.

VIII. BLOCKING DEONTIC AGGREGATION: CONSTRAINTS AGAINST USING

I said above that the appeal to rights is the most common way to reject unrestricted deontic aggregation. It is also possible to postulate deontic constraints on aggregation without specifically appealing to rights. Consider a variation on a case popularized by Judith Thomson, that I call *Homer on the Bridge*: An out-of-control train is speeding toward Apu, Barney, and Moe, who are trapped on the track. On a bridge over the track, between the train and the trapped trio, stand Homer and Ned. Ned realizes that he can stop the train by pushing Homer off the bridge and

onto the track, but Homer will be killed by the collision (Ned is too small to stop the train by sacrificing himself). Many (most?) nonconsequentialists claim that it would be impermissible for Ned to push Homer off the bridge, even if that would result in a better overall state of affairs than the alternative in which Apu, Barney, and Moe all die. Some may claim that there is a (possibly defeasible) constraint against *using* some to benefit others. This seems to be at least part of Nozick's objection to utilitarianism in the passage I quoted above (at the end of Section I). Here is the relevant part again: "Using one of these people for the benefit of others, uses him and benefits the others. Nothing more. What happens is that something is done to him for the sake of others." Applying this to *Homer on the Bridge*, if Ned pushes Homer off the bridge, he will use Homer to save the lives of Apu, Barney, and Moe.

A crucial question at this point is whether Nozick's objection is to the existence of a certain kind of causal structure in the world—a structure in which harm to some produces benefit to others—or to an agent's actively bringing about this causal structure. Does the constraint against using incorporate the doing/allowing distinction? What would Nozick say about the following variant of *Homer on the Bridge*? Ned sees Homer unconscious on the track in the path of the runaway train. Further down the track, Apu, Barney, and Moe are also unconscious on the track. If the train hits Homer, it will kill him, but will stop before it reaches the other three. Ned can, if he hurries, drag Homer out of the path of the train, but he cannot reach any of the other three in time. Is Ned obliged to drag Homer out of the way? On the one hand, if the constraint against using incorporates the doing/allowing distinction, Nozick (and other nonconsequentialists) may say that this case differs in a morally significant way from *Homer on the Bridge*. If Ned does not drag Homer out of the way, he is not *using* Homer to benefit the others, even though Homer's death does, in some sense, produce the benefit to the other three. On the other hand, if the constraint against using is doing/allowing symmetrical, it would seem that Ned must drag Homer out of the way. But this would be a strange result indeed. Consider the situation in which the other three are closer to the train than is Homer. Since they are slighter of stature than Homer, the train will kill all three of them, before stopping short of Homer. Ned is in a position to drag all of them out of the path of the train, but cannot reach Homer in time. In this case, Ned must drag the three of them out of the way, with the result that Homer dies. In this case, that is to say, Ned must act so that whoever is furthest from the train ends up dead!

For an even more disturbing illustration of the results of accepting a doing/allowing symmetrical version of the constraint against using, consider a variant of Thomson's loop version of the trolley case. In *Homer in the Loop*, a runaway train is headed toward Homer, who is unconscious on the track. If the train hits Homer, it will stop. If Homer were not on the

track, the train would loop around to a side-track, on which Apu, Barney, and Moe are unconscious, and would kill all three of them. It is possible for Ned to switch the train to the side-track before it reaches Homer. If the train is switched to the side-track, it will not loop around to the portion of the main track occupied by Homer, whether or not the other three are on the track. It is not possible for Ned to remove any of the four unconscious people from the track. In this case, a doing/allowing symmetrical version of the constraint against using would seem to oblige Ned to switch the train to the side-track occupied by three people, rather than allow it to continue on the main track occupied by Homer. If the three are killed, their deaths do not produce the benefit to Homer, whereas if Homer is killed, his death does produce the benefit to the other three. So, in this case, Ned must divert the train from the few toward the many! The absurdity of this result can be seen if we imagine the initial positions to have been reversed. If Apu, Barney, and Moe had been on the main track, and Homer on the side-track, then Ned would still have been obliged to switch the train to the side-track, this time killing Homer. Thus, whether Ned can act so that few die or many die depends on the structure of the track. It is, of course, possible to bite the bullet and accept this consequence. That would require a kind of fetishizing of causal structures that I find inexplicable, but that I suspect is quite common in nonconsequentialist thinking. Nonetheless, I strongly suspect that the prohibition against using that at least some associate with the separateness of persons dogma is strongly doing/allowing asymmetrical.

Before I leave the topic of a prohibition against using, I should note that my description of the causal facts in *Homer in the Loop* (and other examples) could be challenged by an adherent of a counterfactual theory of causation. If the claim that Homer's death "produces the benefit" to the other three is understood as the claim that Homer's death causes the other three to live, then certain versions of the counterfactual theory of causation entail that the deaths of the other three would "produce" the benefit to Homer. After all, if they hadn't died, he would have. This criticism is well-taken. For all I know, a version of the counterfactual theory of causation that has this result is the correct one. However, if it is, the whole notion of a prohibition against using fails at the first hurdle. Consider the well-known transplant case, which illustrates what many would consider paradigm cases of using and not using. If the doctor takes the organs from the one healthy person to save the five sick people, she has used the one to benefit the others. If she instead lets the five die, she has not used the five to benefit the one. However, a theory of causation that gives the result in *Homer in the Loop* that the deaths of the three would produce the benefit to Homer of keeping him alive, would also give the result in the transplant case that the deaths of the five would produce the benefit to the one of keeping him alive.

If the appeal to rights is to present a genuine alternative to consequentialism, something like Foot's claim about the relative stringency of different types of rights and duties must be part of it. Claims that negative rights and duties are (at least usually) stronger than positive rights and duties (if any such exist) will have to be grounded in an account of the alleged moral significance of the general distinction between doing and allowing, of which the distinction between killing and letting die is a specific example. Without an appeal to this distinction, we will not be able to explain why Moe and Apu do not have a right to be saved that is of equal strength with Barney's right not to be killed in the *Springfield Rescue* cases. Likewise, if we fall back on a prohibition against using in an attempt to block deontic aggregation, we must appeal to the doing/allowing distinction to explain why Ned is not obliged to switch the train away from the solitary Homer *toward* the other three in *Homer in the Loop*. This topic is the subject of much debate, which I do not have the space here to recapitulate.²⁴ My own view is that no one has yet produced anything like a convincing argument that there really is a morally significant distinction between doing and allowing.

One more possibility for erecting deontic barriers to at least some aggregative reasoning is to appeal to the alleged deontic significance of the intending/foreseeing distinction, perhaps as an alternative attempt to explicate the prohibition on using persons as means. Perhaps it is permissible to run over and kill Barney en route to saving Moe and Apu, so long as Barney's death is foreseen but not intended. It would not, however, be permissible to kill Barney as a means to saving Moe and Apu, because in that case Barney's death would be intended as a means, and not "merely" foreseen. Although I do not have the space to explore this topic here, I deal with it extensively elsewhere.²⁵ My view is, first, that the distinction between what is foreseen and what is intended as a means is notoriously difficult to explicate, and second, that even if the distinction does hold up to metaphysical scrutiny it is neither deontically nor aretaically relevant.²⁶ This is not to say that intentions are never aretaically relevant. Final intentions are fairly central to judgments of character. But the kind of intentions that might be thought to underlie the separateness of persons' objections are not final intentions, but rather what are intended as means. The anticonsequentialist objects to the fact that consequentialism licenses harming some *in order to* benefit others. If this objection is to be understood in terms of intentions, it is the intention to harm some *in*

²⁴ See, for example, Jonathan Bennett, *The Act Itself* (New York: Clarendon Press, 1995); Steinbock and Norcross, eds., *Killing and Letting Die*; and Alastair Norcross, "Killing and Letting Die," in R. G. Frey and C. H. Wellman, eds., *The Blackwell Companion to Applied Ethics* (Malden, MA: Blackwell Publishing, 2003), 451–63.

²⁵ See my "The Road to Hell," unpublished manuscript, latest draft available on my Web site (<http://spot.colorado.edu/~norcross/>).

²⁶ To say that a distinction is aretaically relevant is to say that it is relevant to a moral evaluation of character. The Greek word "arete" is usually translated as "virtue."

order to benefit others that is thought to be problematic. This concerns what is intended as a means (the harm to some), rather than what is intended as an end (the benefit to others). The anticonsequentialist does not object to the intention to promote the good as such, but rather to the lack of constraints on that intention.

IX. CONCLUSION

I have argued that the separateness of persons dogma, at least as that notion is understood in the anticonsequentialist literature, can be deconstructed into (at least) two different dogmas. First, there is the untenable rejection of thoroughgoing axiological aggregation. Second, the most common attempts to block deontic aggregation require the acceptance of the moral significance of the doing/allowing distinction. If these two dogmas really are at the heart of the demand for respecting the separateness of persons, it should be no surprise that consequentialists are cheerfully and unrepentantly lacking in this respect. As I have demonstrated, even nonconsequentialists must accept unconstrained axiological aggregation, as the rationally preferable alternative to either rejecting the transitivity of “better than,” or rejecting the coherence of all-things-considered value judgments at all. Further, while not strictly entailed by the structure of consequentialist ethical theory, a rejection of the significance of the doing/allowing distinction is one of the central features of every version of consequentialism of which I am aware.²⁷ Once the appearance of intellectual depth and ethical insight has been stripped from the separateness of persons claim, it is revealed as the all-too-familiar and all-too-unconvincing criticism that it is.²⁸

Finally, I would like to note that it is somewhat Orwellian that the demand to “respect” persons entails accepting a distinction—between doing and allowing—that serves in practice to justify, or at least rationalize, *neglect* for persons. A powerful psychological reason for the widespread belief that killing is worse than letting die concerns the cost of rejecting that belief. Millions die every year in all parts of the world, many of them young children, as a direct or indirect result of extreme poverty. Modern relief agencies, such as CARE and UNICEF, have made it very

²⁷ It may be possible to construct a consequentialist theory that is sensitive to this distinction (see Alastair Norcross, “Should Utilitarianism Accommodate Moral Dilemmas?” *Philosophical Studies* 79, no. 1 [1995]: 59–83), but I know of no one who embraces such a theory.

²⁸ There are, of course, other suggested nonconsequentialist constraints on maximizing the good. For impressively intricate examples, see Frances Kamm, *Intricate Ethics* (Oxford: Oxford University Press, 2007), esp. chap. 5; and for criticisms of some of them, see Alastair Norcross, “Off Her Trolley? Frances Kamm and the Metaphysics of Morality,” *Utilitas* 20, no. 1 (2008): 65–80. I do not mean to suggest in the current essay that the two (or three, if we include the putative deontic significance of intentions) dogmas that I have focused on exhaust the disagreement between consequentialists and their opponents. My claim is that the dogmas I have discussed are the most central to the separateness of persons dogma.

easy for those of us who are even mildly affluent to save significant numbers of them. Most of us do very little to help. If we reject the belief that killing is worse than letting die, it is hard to see how we can judge our behavior as anything less than abominable. This clearly constitutes a powerful motivation, though not a respectable reason, to believe that the distinction between killing and letting die carries considerable moral weight. Though this does not itself constitute a decisive argument against the significance of the doing/allowing distinction, it should make us suspicious of our intuitive acceptance of it, especially in the absence of any good arguments in its favor.

If there is no morally significant difference between killing and letting die in particular, and between doing and allowing in general, it is that much harder to justify our neglect of the underprivileged, both in our own country and abroad. We might well be forced to conclude that most of us who possess even modest resources are seriously at fault for not doing more to help others. This conclusion would certainly be painful. But we should not try to mask a self-interested attempt to avoid it with high-sounding talk of respect. The unpleasantness of a moral conclusion is not evidence for its falsity, or even for its unacceptability. If we have to choose between a position that is rationally ungrounded and one with painful implications, we should grit our teeth and choose the latter.

Philosophy, University of Colorado at Boulder