

VIII*—THE DISTRIBUTION OF NUMBERS AND THE COMPREHENSIVENESS OF REASONS¹

by Véronique Munoz-Dardé

ABSTRACT In this paper, I concentrate on two themes: to what extent numbers bear on an agent's duties, and how numbers should relate to social policy. In the first half of the paper I consider the abstract case of a choice between saving two people and saving one, and my focus is on the contrast between a duty to act and a reason which merely makes an action intelligible. In the second half, I turn to the issue of social policy and investigate how if at all numbers can have a bearing there, if there is no obvious duty on individuals to save the greater number. My proposal is that it is not the bare numbers themselves (or even the ratio of claimants on either side of the dilemma) which explain our intuitions in such cases, but rather considerations of the extent to which each of us can make a reasonable claim on others. In short, I argue: numbers don't count, people do.

Many people hold that, in matters of distributive justice, the number of people affected through a policy is one of the main considerations for or against it. Likewise, now with an eye to an agent's duties to others, many think that individuals should be guided by the number of people affected by their actions. Those who hold that the number of people involved generates reasons or duties for us also recognize that other concerns bear on us too: the urgency or stringency of claims, questions of responsibility, fairness, autonomy, *et cetera*... However, when other moral considerations are equal, they affirm the following: the number of people affected by each of the courses of actions available to us is a fundamental, and indeed in the circumstances conclusive, consideration in deciding what to do. My aim, in this paper, is to question this proposition.

*Meeting of the Aristotelian Society, held in Senate House, University of London, on Monday, 7 February, 2005 at 4.15 p.m.

1. For extensive discussion of arguments in this paper I am grateful to Niko Kolodny, Mike Martin, Mike Otsuka, Joseph Raz, Tim Scanlon, and David Wiggins. I have also benefited from conversations with Ian Carter, Ronald Dworkin, Brian Feltham, Alexander Friedman, Sebastian Gardner, Brad Hooker, Mark Kalderon, Frances Kamm, Hallvard Lillehammer, Barry Loewer, David Owens, Paul Snowden, Zoltán Gendler Szabó, Georgia Testa, Andrew Williams, as well as comments from audiences at the UCL Colloquium in Legal and Social Philosophy, the *Institut d'Etudes Politiques* (Paris), Harvard, Sheffield, and the Collegium Budapest.

My main focus will be the role that numbers play in our thinking about duties and wrongness, and the type of reasons they provide for individual agents, on the one hand, and the ways they bear on social policy, on the other. The case for the weight of numbers in the realm of social policy may, perhaps, seem beyond doubt: no one could easily deny that we need some formula to provide guidance in difficult decisions involving conflicting claims of need and scarce resources. In such circumstances, it seems inevitable that considerations of numbers involved become salient. Many are then led to assume that numbers are one of the decisive considerations in laying out principles of distributive justice. As widespread as this assumption is, and as plausible as the intuitive grounds often offered for it seem to be, still there is a reason to think that our intuitions about various cases are not stable, but can be induced now to favour one kind of policy or action, now to favour another. For example, consider a case described by Thomas Scanlon:

Suppose that Jones has suffered an accident in the transmitter room of a television station. Electrical equipment has fallen on his arm, and we cannot rescue him without turning off the transmitter for fifteen minutes. A World Cup match is in progress, watched by many people, and it will not be over for an hour. Jones's injury will not get any worse if we wait, but his hand has been mashed and he is receiving extremely painful electrical shocks. (Scanlon, 1998: 235)

Most of us feel that even if a hundred million people were watching this game, we ought to interrupt the transmission to save Jones from his hour of severe pain. It is tempting to suppose that our intuitions in this case offer support to a general principle, namely that the number of people involved should not lead us to aggregate small benefits and let them simply outweigh severe burdens. But now consider a second example.

If no effective cure is found, one in a hundred million people will die of a particular disease when they reach their mid-teens. Were we to redirect all the resources presently devoted to education, arts and humanities, into research for a cure, a cure would be found, and the one in a hundred million teenager affected by it in each community would survive.

If any individual has a claim of need, surely such a teenager does. Yet many of us lack the intuition that we should so redirect resources to saving the person affected by the rare disease. At the very least, we hesitate before closing all the universities and all the schools in order to devote the money to research for saving this life.

There are important differences between the two cases. To be sure, depriving the young of school education is not on a par with missing out on a football game, even a thrilling one. However, given the similarities between these two cases in terms of distribution of burdens and benefits, and of the ratio of numbers on each side, the question remains: why do our intuitions go so firmly in opposite directions in the two examples? Few would be inclined to consider that the need of the teenager is *outweighed*. Yet we doubt that all these funds must be redirected for research to cure his illness. The question is how to explain this attitude.

Reflecting on cases of this sort, we might conjecture that a picture of our duties to others that starts from the assumption that numbers matter, and then operates with comparisons of aggregate costs and benefits, will not be able to comport with all of our intuitions about what is just. But how else are we to address dilemmas related to meeting urgent needs with limited resources? The purpose of this paper is to make some progress in the understanding of such dilemmas, and to sketch what might be a more intuitive theoretical framework for approaching them.

My discussion will concentrate on two aspects: how best to explain our intuitions that numbers matter, and the grounds of duties and principles. In the first half of the paper I consider the abstract case of a choice between saving two people and saving one, and my focus will be on the contrast between a duty to act and some reason which makes an action intelligible. Having come to question the existence of any duty related directly to numbers in the individual agent's case, we turn to the social case and investigate how if at all numbers can have a bearing there, if there is no obvious duty on individuals to save the greater number. In particular I will explore the threshold to reasonable demands we can make on others, and whether we should suppose that, in the social sphere, principles provide comprehensive reasons which determine the right way of acting. My proposal will be that it is not the bare numbers themselves (or

even the ratio of claimants on either side of the dilemma) which explain our intuitions, but rather considerations of the extent to which each of us can make a reasonable claim on others.

II

Let us start with a simple case, even if described in the abstract. We have a choice between saving two people and saving one, and we cannot save all three. If everything else is equal, and given that numbers are the only salient feature, most people feel compelled to say that we must save the two. (In the rest of this discussion I will refer to this idea as the principle of saving the greater number, SGN for short.)

In her discussion of the doctrine of double effect, Philippa Foot states this conviction in the form of a manifest duty: 'We feel bound to let one man die rather than many if that is our only choice' (Foot, 1967: 9). (Foot is discussing a choice between giving a drug to a patient who needs it to save his life, or giving the same drug to five patients each of whom could be saved by one fifth of it.) In a brief and illuminating discussion of this point, Elizabeth Anscombe observes: 'Mrs Foot regards it as obvious that one must save the greater number'; nonetheless Anscombe claims to feel 'a curious disagreement about this' (Anscombe, 1967: 16).

My contention, in this paper, is that there are reasons to share Anscombe's 'curious disagreement'. That is: we may ask why there is such an obvious duty to save the greater number. To voice such a doubt is not to say that numbers should not count; for they may be a perfectly intelligible reason to act one way rather than another. Rather our initial question is whether numbers in the case under discussion do provide a conclusive reason from which we can derive a duty.²

Throughout this paper I will rely on this distinction between conclusive and merely intelligible reasons, so let me say one more word to clarify it. The idea of a conclusive reason, as I use it, is

2. In an article published a decade after Anscombe's discussion note, John Taurek too contests that numbers should count in situations such as the example given by Foot. I share some of Taurek's views, but for reasons which will become apparent, I disagree with one of his conclusions. I disagree, that is, that the right thing to say about this case is that numbers 'should not count'.

this. A *conclusive* reason in a given situation, for Φ -ing, is one which, when introduced into our deliberation, leads directly to the conclusion to Φ rather than not.³ Our question is whether in a choice of action the fact that there is a greater number on one side over another provides a *conclusive* reason for that course when all else is equal.

Before we proceed, let us state the problem under consideration a little more precisely. I assume that we have no prior commitment to consequentialism—to the view that the rightness or wrongness of actions is determined by properties of states of affairs brought about by acting. More exactly, I shall not start with the assumption that there is any *impersonal* value, realized in either option related to the numbers of people involved.⁴ Rather, the route I shall pursue is to ask whether we can find reasons for the agent to act in one way rather than the other which can be offered as a justification to all individuals involved.

The advantage of Anscombe's perspective for that purpose is precisely that on her reasoning the justifiability of SGN depends on *individuals'* reasons to accept it or to object to it. Someone may object that it does not make sense to impose such a constraint here, since SGN cannot be a duty to the individuals affected, and in particular to the individual(s) who will *not* be saved. Someone who made this objection would seem to be assuming that we have duties towards people only as the objects of our actions. To respond, a distinction is needed. I may have a duty to save, but not a duty to save *you*. A duty to save can be justified to each of the individuals involved even if not every single one of these individuals is going to be saved. The fact that I justify a principle *to* you does not by itself commit me to doing something *for* you. So our precise question is whether such a justification of a *duty* to SGN can be offered to all individuals involved. That is, an explanation of why the greater number *must* be saved which does not presuppose the principle of SGN, but rather seeks to establish it.

3. Note that a reason's being conclusive in one situation does not entail its being conclusive in significantly different cases. For more on this, see Section IV.

4. By impersonal value I mean one the worth of which obtains independently of its being a good for an agent or agents. Such a value would make one situation better than another without thereby having to be good for an agent or agents related to that situation.

III

Let us now turn to Anscombe's line of reasoning. Anscombe has two aims. The first (and probably most important) is to illustrate a scepticism towards a law-like conception of the moral 'ought'—previously elaborated in her *Modern Moral Philosophy* (1958). I shall return very briefly to this theoretical question in Section VI. For now I intend to concentrate on her second, more circumscribed, discussion, namely: who is wronged when we choose to save one, in a situation in which we can either save one person or several, but not all?

Her reasoning about the Foot case of the doctor faced with a choice between giving a life-saving drug to one or to five is this:

Suppose I am the doctor, and I don't use the drug at all. Whom do I wrong? None of them can say: 'You owed it to *me*.' For there might be nine, and if *one* can say that, all can; but if I used it, I let one at least go without and he can't say I owed it to *him*. Yet all can reproach me if I gave it to none. It was there, ready to supply human need, and human need was not supplied. So any one of them can say: you ought to have used it to help us who needed it; and so all are wronged. But if it was used for someone, as much as he needed it to keep him alive, no one has any ground for accusing me of having wronged *himself*.—Why, just because he was one of five who could have been saved, is he wronged in not being saved, if someone is supplied with it who needed it? What is *his* claim, except the claim that what was needed go to him rather than be wasted? But it was not wasted. So he was not wronged. So who was wronged? And if no one was wronged, what injury did I do? (Anscombe, 1967: 16–17)

I shall come back to this last question (Is there a wrong committed if no one is wronged?) in Section VI. But first let us focus on the question of which person is wronged, if we fail to save the greater number. Note that the question introduces a perspective which helps clarify what I said in passing in the previous section, namely that justifying *to* you a duty to save does not mean that there is a duty of saving *you*. For there is at least one person who is not going to be saved (the one who is on her own), and no one seems to want to argue that the right thing to do is to save no one and to let all die. So on every account of what is the justifiable thing to do there will still be at least one

person for whom nothing is done, and that person can still not oppose the principle by saying that she has been wronged. This may need spelling out a little further.

Thomas Scanlon, in particular, has set out to meet the challenge of justifying *to* the person isolated (the person who is not going to be saved if SGN is observed) the fact that the drug should go to the others. Suppose that we know independently that SGN is the right principle to adopt. Then in one sense, the sense in which all are wronged if the right principle is not observed, this person is wronged if the drug goes to her. And indeed I may feel aggrieved if benefits come to me in the wrong way, depriving others of what they are owed. But saying that the person who is isolated should feel 'wronged' if her life is saved (in violation of SGN) looks strange; and indeed Scanlon does not phrase it this way at all. What he says is that it would be unreasonable for the isolated person to reject SGN. But now another doubt arises. If Anscombe is right, if strictly speaking no one is wronged whether I choose to save the one or the many, why would it be unreasonable to reject SGN?

Scanlon's answer is that it is each 'additional' person who is wronged. In an argument which draws on Kamm's views,⁵ Scanlon asks us to consider a series of cases: zero to one, one to one, one to two, zero to two. The reasoning against the permissibility of saving either a larger or a smaller group of people is this:

In such a case, either member of the larger group might complain that this principle did not take account of the value of saving his life, since it permits the agent to decide what to do in the very same way that it would have permitted had he not been present at all, and there was only one person in each group. The fate of the single person is obviously being given positive weight, he might argue, since if that person were not threatened then the agent would have been required to save the two. And the fact that there is one other person who can be saved if and only if the first person is not saved is being given positive weight to balance the value of saving the one. The presence of the additional person,

5. See (Kamm, 1993): Chapters 5 and 6; Scanlon refers in particular to pages 116–117. Kamm briefly considers Anscombe's challenge (Kamm, 1993: 119), but argues that a process of balancing of personal interests on each side would decide in favour of SGN. (She thinks that if the presence of an additional person makes no difference, that person is not given due recognition of her equal significance.)

however, makes no difference to what the person is required to do or how she is required to go about deciding what to do. This is unacceptable, the person might argue, since his life should be given the same moral significance as anyone else's in this situation (which is, by stipulation, a situation in which no one has a special moral claim). (Scanlon, 1998: 232)

If we don't get distracted by the language of 'weights', Scanlon's point is simple and appealing, namely that we need a different procedure for one-one cases and one-two cases. (Just as we have a different procedure for zero-one and one-one cases.)⁶ This is the 'additional' person's case, and Scanlon's case for SGN: the additional person would be wronged if we adopt the same procedure for a one-one case and a one-two case. But now we have just pushed the question one step back, without getting rid of it. For in order to agree that we need to treat one-one and one-two cases differently, we need already to believe that numbers matter.

Scanlon's thought here is attractive in that we seem to have a sequence of different cases in which we need at each stage a different procedure. We can represent this in the form of a table:

	How Many to be Saved?		
i	0	0	Save no one
ii	1	0	Save the one
iii	1	1	Save either
iv	2	1	Save the two

However, the appeal of this way of presenting things is misleading. For what is the salient difference at each stage? Faced with this table, Anscombe would answer that if we attend to the duty to save, the need for modifying the procedure between i) and ii) and ii) and iii) is apparent, but not so between iii) and

6. This insistence on an adequate procedure is closely linked to Scanlon's reliance on reasonableness: the test of a principle is not whether it makes some people better off than they would be in its absence, but whether it is the best procedure for deciding conflicting claims, a procedure that not even these who stand to lose by it can reasonably reject.

iv). She would contend, therefore, that the table should only have three rows thus:

	Is there human need?		
i	No	No	Save no one
ii	Yes	No	Save the one
iii	Yes	Yes	Save either

So Anscombe can agree with Scanlon that different cases need different procedures. But they disagree on what constitutes the difference between cases; they disagree, that is, whether the salient feature of each case is the *number* of people in need, rather than the mere human need to be helped.⁷

Is this too quick? Someone attracted by Scanlon's justification for SGN might object that, contrary to what I just claimed, numbers are not presupposed in his reasoning. The justification he offers for SGN is addressed to individuals. Surely we need to justify our action to the additional person; her life must be given *some* consideration, and this supports saving the two without presupposing any aggregative judgment. At no point, in other words, does the justification rely on saying to the person isolated, who is not going to be saved: 'We must save the two because they are *two* and you are just one.' A first worry is that the tie-breaking argument might then not generalize beyond the case of just one versus two.⁸ Be this as it may, Anscombe could agree that the additional person's life must be taken into account, and add that, in her reasoning, the person's life *does* count. The additional person's life is a claim that the rescuer

7. Maybe some will disagree. They will argue that the fact that Anscombe has one less row shows precisely that she is abstracting from an important feature of the case (the fact that there are different number of people on each side). But Scanlon is also abstracting from certain features of the cases under consideration. For example, he is not changing the procedure whether the person to be saved is on the right or the left hand side. (He treats in the same way 1-0 cases and 0-1 cases, 2-1 cases and 1-2 cases and so forth.) Why? Of course, because the spatial distribution of need does not matter. But just in the same way, unless you have *already* established the fact that the number of people in need matters, you cannot assume a relevant difference between rows iii and iv.

8. Scanlon raised just this concern in personal correspondence. The thought is that in cases of n versus $n + 1$, where n is greater than 1, we cannot think of a *tie* that the additional person is breaking without making a genuine aggregative judgment.

must take into consideration; it is human need which works in him as a reason. Its *failing* to count can only be illustrated by an aggregative thought: namely that there is one on one side and two on the other, and that this should generate a different policy.

Scanlon seems to be presupposing that SGN is a correct principle, since the distinctions he draws implicitly assume that numbers are significant to the choice made, rather than establishing that numbers must count for us even when we start out just from the perspective of all parties involved. Hence, it is not clear how we could establish SGN once we press Anscombe's question.

IV

What are the consequences of adopting Anscombe's perspective?

To begin with, her way of envisaging this problem avoids the distraction caused by the language of weights in the assessment of reasons (especially when we are addressing the relevance of numbers as a reason to adopt a course of action over other possible courses of action). Otherwise one might give in to the temptation to think that we have as many reasons to save as there are people on each side and that the reasons can be 'added'. So in the case of two people against one, other things being equal, the reasons to save the two people would seem to 'outweigh' the reason to save the one. This way of thinking about reasons should be resisted, even if it seems tempting to many people in the particular case under discussion.⁹ Two reasons may sometimes combine, but reasons do not become conclusive by juxtaposition; and one conclusive reason in favour of a course of action is sufficient.¹⁰ In this case Anscombe suggests that we have one

9. For a recent example see (Parfit, 2003: 378). Parfit asks us to consider a case in which we, equipped with a lifeboat, could save either White, stranded on one rock, or five other people, stranded on another. He then writes: 'The five must be allowed to argue that...their five reasons together *outweigh* White's.' (Emphasis added. Parfit uses this argument precisely against what he calls 'Scanlon's Individualist Restriction'.)

10. Scanlon would seem to concur, for he follows Kagan in warning against the 'additive fallacy' and the 'slippery metaphor of "weight"' in the domain of reasons' (Scanlon, 1998: 397, n. 35). Unfortunately, Scanlon takes it back immediately, and succumbs precisely to this metaphor. For he claims that, in the case under consideration, 'the significance of the additional life does take this form' (of an additional 'weight' which decides the case by tipping the balance of reasons).

and the same conclusive reason on each side, namely the human need to be saved. (The reason is 'conclusive' because it establishes that it would be wrong to save no one, that *each* of the people in need would be wronged if help were entirely withheld.)

The second, and related, aspect brought out by Anscombe's discussion is that *even in the pure abstract case* in which there is no relevant factor other than each person's need to be saved, we do not have a duty to save the greater number. Even, that is, if everything else *is* equal, we have a choice between saving two people and saving one, and that is our only choice, there is nothing wrong about taking *either* course of action, at least from the perspective of what we owe to each of the people involved.

This contradicts an apparently appealing way of thinking of the case at hand: Concern for particular individuals generates constraints on the application of aggregative standards (and on the related impulse to minimize harm). These side constraints should be firmly imposed when, and in as far as, aggregative reasoning yields unjust or unfair results; e.g., ignoring the separateness of individual losses in a given distribution, because of the focus on combined benefits. But when such threats to the concerns of individuals are not in play, for example as in cases in which losses or injuries are evenly balanced and the *only* difference between taking one course of action and another is the number of people saved, then aggregative standards should be applied without restriction. On this line of thought, the *only* principle that can govern your action in our example is SGN. If everything else is equal, that is, you have a duty to save the greater number.

There is something right in the intuition here, as we shall return to below: in cases of social policy in which numbers do differ markedly there are often grounds which compel us to act to save what is, in fact, the greater number (if not *because* they are so much more numerous). For all its popularity, however, I don't think that we should accept the reasoning here, or the claim that, in the pure abstract case, numbers 'tip the balance'. First, this kind of reasoning seems to assume that there is some form of impersonal value in play here which recommends saving the greater number. And one might question whether there is any such impersonal value as related to quantity. One might criticize consequentialist reasoning, that is, at the point at which

it introduces values to guide our actions, rather than waiting to rein in what would otherwise lead to unintuitive consequences if left unconstrained.

Second, and here more pressingly, one may wonder whether just any value will ground a duty of action. *Must* the balance be tipped? That is, even if we grant that there is some impersonal value which recommends the saving of the greater number over the smaller, does this generate a duty?¹¹ Surely all that would have been established is that it is *permissible* to save the greater number, not that it is *imperative* to do so. For, unless the mere fact that there is a reason to act on one side which is not echoed on the other requires us on pain of wrongdoing to choose that course of action, no duty will have been established.

Let me take these two thoughts together a little more slowly. There are delicate issues here to do with the relation between reasons, wrongness and duty. Suppose some impersonal value is present which favours saving the greater number over the smaller, does this give us a conclusive reason? We might just say that such a value gives us what Anscombe calls ‘a perfectly intelligible reason’ (Anscombe, 1967: 17) to act on one side which is not echoed on the other. The reason is ‘intelligible’ in that it is not a ‘bad’ reason. If your reason for saving the two rather than the one is that they are rich, the one on her own poor, and that the need of poor people doesn’t work in you as a consideration, this would constitute a ‘bad reason’ (that is: no reason at all).¹² But if you say that you saved the two because they were two rather than one, because you want to maximize the value of saving, then people unmoved by aggregative considerations could respond: ‘In the same situation I would not have taken this course of action *for that reason*, but I understand why you did it.’ And this is also all that

11. Note that I am merely granting the possibility that there is an impersonal value so related to quantity for the sake of argument.

12. A bad reason is not merely irrelevant. If one saved one party rather than the other for bad reasons, what one would do would be wrong. Contrast what Kamm would call ‘irrelevant utilities’ (Kamm, 1993: 101). Unlike Kamm’s irrelevant utilities, which may be a good reason to act in other circumstances, a bad reason is an illegitimate consideration. (The idea of irrelevant utilities is illustrated by a case in which you have a choice between saving A’s life and saving B’s. Alongside B is C who has a sore throat. The drug that can save B’s life can also cure C’s sore throat. Kamm thinks that the sore throat of the third person should not tip the balance.)

the person affected by the consideration that there are more people on one side can say if you are moved by another, equally intelligible, perfectly good, reason to take the opposite course of action. She cannot, that is, blame you for wrong-doing. So we can now return to the possibility left open by Anscombe of a wrong committed without anyone being wronged. Even given the existence of an impersonal value related to quantity, a wrong could only be committed if this value grounds a duty. And we have seen no reason to suppose any such impersonal value grounds such a duty.

But does it really make sense to suppose that someone moved by the numbers involved is acting merely on a reason which makes intelligible their course of action, rather than acting on something which requires them to move in some way? The questions of there being a conclusive reason and of there being a moral reason, or a duty, seem to be closely aligned. First, one might be tempted to suppose that, given that we have highlighted a reason to favour the one course of action, and that everything else is equal, then the agent will anyway be required to act in that way. That is, if one thinks that an agent is always required to do whatever he or she has most reason to do, then the existence of a reason on the one side with no compensating reason on the other would always require saving the greater number. To endorse this temptation is to accept a 'maximizing' conception of reason: that we are rationally required to do what we have most reason to do. Such a conception has no room for a distinction between a merely intelligible reason for an action, one on which one may act but need not, and a conclusive reason. Such a stance is surely controversial. After all, it is not immediately obvious that someone who feels the attraction of eating a peach, but in the end fails to do so, though no other consideration intervenes in their thoughts, thereby exhibits practical irrationality.

Suppose, then, that we allow that there can be a reason which could move an agent to act in a certain direction, but does not require that they so act. Making salient that the numbers involved is indeed a value for us (if it is) would establish at least this. Of course the defenders of SGN require more: that in the circumstances there is a duty on the agent, that he or she would otherwise have acted wrongly; and moreover that from this fact we can identify a principle that we can also see to be operative

in other circumstances. How should we bridge the gap between these? Well, yet another temptation is simply to say that surely if the numbers involved give us a reason, then they give us a *moral* reason to act. And how could a moral reason be other than conclusive?

Given our current stage of enquiry, this temptation must be resisted. What could it be to say that the reason in question is moral, other than to say that in the circumstances it gives us a duty? And of course, that there would be a duty to save the greater number, and that there are no countervailing considerations, would seem to offer conclusive grounds for action. But to reason in this way is simply to assume that there is a duty, rather than to argue for one, or explain why one should arise.

Once we press Anscombe's question, we can see that, in the purely abstract case where all that differs is the numbers of agents involved, however slightly, it is just not obvious that one must save the greater number. There seems to be no reason to offer independent of the difference in number. Even if we can see the difference as indicating an intelligible reason for acting one way rather than the other, nothing shows that it must be conclusive in action.

I have been writing as if the existence of a duty requires just that it would be wrong for an agent to have acted otherwise *in those particular circumstances*. But typically talk of a principle of action is taken to imply consequences for actions in rather different circumstances. And that seems to be part of the rationale for defending SGN (a principle) by appeal to an abstract situation which could never actually be realized. Yet without a demonstration of there being a duty in the narrow sense that I have used above, in these peculiar circumstances, as yet we have no reason to suppose that there is any general moral principle connected to numbers.¹³ (I'll return to the question of duties and principles in Section VI.)

13. Parfit argues that Taurek's argument against SGN ignores the existence of agent-centred restrictions on the obligations on us (Parfit, 1978). While that may be a fair observation to make against Taurek's strategy of testing the existence of an obligation against one's claims of interest in a friend, it does not touch the worries expressed here, which focus rather on how the existence of either duty or principle could be demonstrated in the first place. (Parfit himself might well be unmoved by such a complaint since he finds aggregative considerations so compelling in the first place.)

V

I have argued so far that all that we can say, even in the most abstract case in which all that distinguishes the options for action is the numbers to be saved, is that it is permissible and not required that one save the greater number. This result might look disappointing and lacking in direction. In particular, some people may feel dissatisfied because they want to know what we *ought* to do in that case. They are unhappy with the claim that it is permissible to save either side, that the only duty we have is not to withhold help. So they would urge us to choose between two alternative ways of thinking about this problem. I just mentioned (and rejected) one of them, namely that we affirm SGN, but only if everything else *is* equal. (Then numbers ‘tip the balance’.) The other alternative is to say that it is a *bad* reason to save the two because you want to save more rather than less.¹⁴ Let me add a very brief comment on how the reasons for rejecting the first of these alternatives also lead us to reject the idea that it would be *wrong* to let numbers count.

As we have seen, what is problematic is why there should be a duty for an agent to save the greater number. That leaves entirely open whether numbers generate a reason on which an agent may act. That is to say: we don’t want to narrow the discussion down to the stark alternatives of duties or bad reasons, without further argument.

There is no reason to say that someone who is moved to save the two rather than the one because he wanted to save more rather than less was thereby moved by a *bad* reason. Likewise we should be suspicious of the demand that you should devise a procedure, perhaps flipping a coin, in order to decide whom to save.¹⁵ This would be to endorse a principle according to which you have a duty to save whoever wins in the coin-flipping procedure. In the pure case, the case in which the *only* salient feature which differentiates two incompatible courses of action is the number of people on each side, we have indeed run out

14. Taurek seems committed to this position (at least where the choice is between saving five people and saving a friend): ‘The numbers, in themselves, simply do not count for me. I think that they *should not count* for any of us’ (Taurek, 1977: 310, emphasis added).

15. This is the procedure suggested by Taurek (Taurek, 1977: 303ff).

of reasons to affirm that it would be wrong to save the greater number.

That is: in the pure abstract case, our conclusions are merely that there is a duty to save one party or the other, that it is permissible not to save one party in saving the other, and that there may be further reasons an agent has that make intelligible his or her choice in saving one of the parties.¹⁶

VI

Anscombe concludes her discussion of these issues with an expression of scepticism about systematic theories of reasons. She writes, ‘When I do action A for reasons R, it is not necessary or even usual for me to have any special reason for doing-action-A-rather-than-action-B, which may also be possible’ (1967: 17). So her whole discussion may appear as an instructive illustration of a more general distrust of the ambition to come up with a systematic set of rules and principles to settle all moral dilemmas.

My aim, in this paper, is far more modest, but let me say two words on this aspect. Anscombe’s scepticism regarding moral obligation, her claim that we ought to jettison the moral sense of ‘ought’ may be overstated. But there is something sound in her suspicion of rigoristic convictions stated in the form of rules on which every judgement of right and wrong is then supposed to rely. This is indeed a very strange way of thinking of moral principles. A much more plausible way of understanding principles, as Thomas Scanlon urges, is as the conclusion about the status of various reasons, a form of shorthand for a complex structure of reasons.¹⁷ What all principles will share, then, is

16. Note that my emphasis has been throughout on wrongness and duties. In strict ignorance of whether one will be in the smaller or the bigger group, so that one is only moved to maximize one’s chances of survival, one’s probability of ever being saved is maximized by a principle of always saving the greater number. But there is no obvious way in which this observation connects with Anscombe’s question.

17. ‘In making judgements of right and wrong we are drawing on this complex understanding, rather than applying a storable rule, and this understanding enables us to arrive at conclusions about hard and difficult cases, which no rule would cover’ (Scanlon, 1998: 201). For a similar perspective on principles see (Kolakowski, 1999: 31): ‘But if an absolute injunction against lying is both ineffective and potentially in conflict with other, more important moral imperatives, how are we to go about finding a general principles which will take account of those times when lying is permissible? The answer . . . is that there is no such principle: no general rule can take account of every conceivable moral circumstance and provide an infallible solution.’

that they contain in crisp form what is effectively a more complex rationale for imposing constraints on certain forms of action.

If we understand principles in this way, there may be an infinite number of them (Scanlon, 1998: 201). But they can all be traced back to the wrong that would be committed in the absence of the principle (to the reasons for imposing certain constraints). Principles thus understood are not, then, a way of organising all the reasons we have.

In earlier sections I have claimed that if we have a choice between saving two people and saving one (and we cannot save all three), then, other things being equal, it is permissible to save either side. We have seen reasons to express doubts about the presence of a principle here. Were there such a principle in the most abstract case, it should be clear to us that there was a duty to save the greater number. But short of simply affirming the existence of such a duty, the defender of SGN seems at the very best only to highlight for us a consideration which makes intelligible why one chooses the larger party over the smaller; and we have seen no ground for why this should be required of us. This strongly suggests that either course of action is permissible. Let me now make more precise the idea of permissibility in the light of the conception of principles I just invoked.

In saying that it is permissible to take either course of action I am saying that, in this case, neither action is ruled out by any conclusive reason which dictates what we have a duty to do or refrain from doing. This is not to say that there is a single principle which determines for either course of action whether it is permissible or not to act in that way. An agent may find a reason for preferring one course of action over the other, a reason which makes intelligible her choice without thereby giving us a conclusive reason for so acting. Someone who feels moved to save the two in the most abstract case is moved by just such a non-conclusive reason. We should not think that the realm of duties and obligations occupies the entire domain of reasons for action. Perhaps most would readily grant this. What I have argued here is that *once* one grants this, one should also recognize there is a significant gap between identifying something which can act as a reason for an action, and identifying any

corresponding duty or obligation. And this is the terrain in common between Anscombe's scepticism and my more cautious response to the case in hand.¹⁸

VII

A doubt now arises. How does what I have said about the dilemma facing an individual relate to dilemmas facing us in social policy, where the numbers involved are much greater? How if at all can numbers have a bearing in the social case, if there is no duty on individuals to save the greater number? Many will hold that when we shift the perspective from the choice which faces a single agent to (arguably) more pressing and constant choices in terms of public policy the case for SGN is obvious. For these choices involve limited resources, for the application of which numbers seem an obvious relevant factor.

Let us try to make this thought more precise. I have been concerned all along with an agent faced with the choice between saving two people and saving one. The central claim that it is permissible to save either side meets with resistance from people who understand the choice of the agent as a choice between doing some good (saving one) and doing more good (saving two). In previous sections I voiced doubts regarding this broadly consequentialist way of understanding the case at hand. But many will think that this is too quick.

It may for instance be argued that the pure case we have discussed obscures quite how strong the intuition that numbers count is. For consider. If numbers do not count, then the example of an agent faced with a choice between saving two people and saving one can be modified; the decision can now be whether to save one or to save a billion people. The claim that it is permissible for such an agent to save either side in the modified example does look seriously fanatical. Surely, this line of argument continues, this shows that numbers *do* count, and the burden of proof is on whoever denies this.

18. I shall return to the question of the comprehensiveness of reasons in Section VII. To anticipate, I suggest that in the social sphere too, principles provide a framework within which there can be distinct intelligible reasons, rather than comprehensive reasons which determine the rightness or wrongness of all ways of acting.

That we can seriously doubt the permissibility of saving either side in the modified example sounds right; however, it is not clear what this line of argument succeeds in establishing. The person proposing this argument is inflating the number of people on one side, in order to insist that numbers do matter. This, however, just raises the puzzle for us of why you have so to inflate numbers to make the case obvious. For remember: the initial case was so framed that the only salient difference between the options was the number of people involved on each side. We might then wonder whether an alternative explanation is available of why the case introducing such a large differential is so compelling.

Imagine a modification of the Foot case of a choice between giving a drug to a patient who needs it to save his life, or giving the same drug to five patients each of whom could be saved by one fifth of it. Now the decision is between, say, devoting resources to prevent the spreading a fatal disease, which would otherwise affect and kill a billion people, and devoting the very same amount of resources to prolonging the life of one child. This puzzle seems to have exactly the same structure as the Foot case: giving resources to a billion people who could be saved by one billionth part of it, and giving all these resources to one person who needs it all to save her life. However, the superficial structural similarity is illusory. I want to suggest that the source of our reaction to this second case does not lie in the weight of numbers; it is, rather, a response to a fact that all theories of distributive justice must deal with, namely that principles of justice operate in any human societies within the constraints of limited resources available even to meet significant needs. This, in and by itself, means that there is a limit on the claims that a person can have on others. There are restrictions, that is, on the reasonable share each person can claim of scarce social resources available, even to meet her vital needs.

Let us take stock. The defender of SGN highlights the strength of consequentialist reasoning thus. We might accept that there are counterintuitive consequences to pure aggregative considerations when these lead to the lesser benefits of many people outweighing more urgent interests of a few. This imposes some side constraints. But when there are similar harms facing two groups of people, numbers 'tip the balance'. If you are

not moved when numbers are relatively small, just imagine the possibility of saving a great number of people, and balance it against saving just one person. I have resisted the analysis of the first case; and I deny that any intuitions in the second case support SGN as applied to the first case because the two cases have a different structure. Where our intuitions are moved when numbers are inflated this is for a reason other than the mere difference in numbers between the options.

To see this, return to one of our initial examples: the case of the one in a hundred million teenager with a fatal disease. It is a terrible thing for the teenager not to be cured and so to die so young. Nevertheless, faced with such a dilemma, we do not think that we are required to abandon resourcing education, the arts and universities in order that he should survive. Why don't we find this vital need here a compelling reason to act?

One possible explanation of our intuitions in this case is a consequentialist thought, now played out in the social domain. Namely, that the claims of need of an individual have to be balanced against other kind of needs in society, and hence can be outweighed by sufficient number of claims of other people.¹⁹ But that is not the only explanation available, or the most intuitive or plausible. An alternative way to think of these questions is to envisage reasonable demands made of individuals, and their reasonable objections to particular schemes of distribution. This is to ask what constitutes not only a fair, but also a *reasonable*, way of sharing resources. In asking this question we insist upon considering individuals both within the perspective of making claims, and within the perspective of having demands made upon one. From this we can derive the idea of what I would call a *reasonable threshold* beyond which I cannot make demands on others, and they cannot make demands on me. This is not necessarily something which we can determine by any algorithm, but it is something we can recognize through the reciprocal process of considering the two perspectives and the reasonable resistance to or insistence on a given claim.

Given this way of looking at things, the reasonableness of the demand depends both on how the various parties will be affected,

19. For a detailed discussion of dilemmas involving conflicting claims of needs, see Griffin, 1986.

and on their share of common resources. To understand how we make this kind of judgments, we need to draw a contrast between an almost absolute evaluation of something as a misfortune, and what reasonable claim one thereby has on others to prevent or alleviate that misfortune. It is a dreadful misfortune to die at fifteen, well before the natural span of human life. However the claim that the teenage victims of the extremely rare genetic disease have on resources available to save lives is not a simple function of the severity of their misfortune. The teenager has a claim on common resources, as does any other individual in need of help. At the same time, there is a threshold above which they cease to have a claim: a threshold above which the share they would claim of overall resources would become unreasonable, for it would make unreasonable demands on others.

Now the idea of a threshold seems to invite a debate concerning the exact point after which we are no longer required to help others, even if they are in desperate need. In particular, the idea of a reasonable threshold could be interpreted as a variant of consequentialism with side constraints. (So the idea would be to arrive at a general principle determining whether, and if so at which level, others can hold back from making demands of individuals. The problem thus understood would be to determine systematically how to combine maximising considerations and considerations of relative costs.) But so interpreting the idea of a reasonable threshold would constitute a misunderstanding. In the rest of this section, I want to sketch a little further the contour of this idea, and to return to the comprehensiveness of reasons.

Why would it be unreasonable for a teenager affected by the extremely rare illness I envisaged above to claim all the available resources in his country? How are we to understand the limits on the claims of needs that this person has on others? It seems that *in the circumstances* others are not unfair in not throwing all the resources available at him, that it would be *selfish* for this teenager to demand so much of others. But this thought has to be interpreted carefully.

My talk of thresholds here should not be interpreted in terms of a fixed quantity derived from costs and benefits across society; nor more generally should we suppose that there is some algorithm to which we can appeal which will determine

for us when a claim is above or below the threshold. It is true that some may prefer to read this point in a consequentialist manner; asking for a general principle which establishes a precise level of cost at which it is reasonable for individuals to reject demands of a certain kind made on them. My own suspicion is that this is a vain and wrongheaded hope. A judgement is required of what we can reasonably demand and withhold in particular circumstances, and it doesn't seem that 'any plausible theory could eliminate the need for judgments of this kind' (Scanlon, 1998: 225). The idea is rather this: we recognize that others have a claim on us, and we on others; but we also have an understanding that in particular circumstances, with scarce resources and a given pattern of claims, I do not occupy a position from which it would be appropriate to demand more from others. Exactly how these matters are to be determined will vary from circumstance to circumstance as in all practical matters.

The idea of a *reasonable threshold* thus defines an upper limit to the share of overall resources that we think can reasonably be demanded by any single individual. Because this threshold is typically reached in cases which involve a choice between saving either a very small or a very large number of people in need of help, it might be tempting to think of it as being decided by the weight of numbers; but this does not seem the best explanation. While a consequentialist reading of this intuition would make clear that numbers count, on our alternative reading, that is not what matters, but rather the constraints imposed by the scarcity of resources in the actual circumstances, and the reasonable demands you can make on others.

By acknowledging these two limits (what we cannot reasonably demand, what agents can reasonably withhold as things stand), we are appropriately respecting persons (and being respected by them). This is where the key contrast with a seemingly similar line of thought, which leads to consequentialism, is to be located. One of the key appeals of consequentialist reasoning in social policy is the attractiveness of not being selfish in one's demands on joint resources. But if we take this thought to justify a maximizing conception of distribution of welfare or resources, we have starkly counter-intuitive consequences. For then the scope for the partiality

of our concerns, which feeds into the justification for the very thought that it would be selfish to demand so much of others, is itself undermined.

In contrast, the idea of there being a reasonable threshold beyond which I cannot make demands on others, and they cannot make demands on me, leaves a space for such partial concerns untouched. In other words: not all of an individual's ends are up for consideration, not everything can be calculated in terms of utility or disutility, without losing sight of the very reason why we began distributing in the first place. This, I suggest, is the way to understand the limits on claims of needs in the example of the teenager and the extremely rare illness.

We now see in the social sphere an echo of the thought we already raised in relation to individual moral action, namely that we should not suppose that the principles of right action provide comprehensive reasons which determine the rightness or wrongness of all ways of acting. Rather, they provide a framework within which there can be distinct, and even conflicting, intelligible reasons which can move agents in one direction or another. And now at the social level, we see a need for principles of distribution which can negotiate the reasonable demands we can make on each other, given the scarcity of resources, while allowing for each to pursue their own particular projects.

VIII

The introduction of the idea of a reasonable threshold might appear unsatisfactory, for at least two reasons. The first of these is the difficulty of determining whether, and if so at which exact point, we reach such a threshold. We saw that this objection originates from a misunderstanding of the idea. But now I want to consider a second objection, more germane to the case under discussion.

The objection I have in mind concerns the difficulty of explaining what, for lack of a better word, I will call 'intermediate' cases. Let me explain. Thus far we considered the case of an agent faced with a dilemma between saving two people and saving one (when he cannot save all three). We saw that all three people have a *claim* to be saved, and that it is permissible for the agent to save either side. As for the decision to save

either one or a billion, I suggest that there is a threshold after which we cease to have a claim on others. But neither of these two diagnoses is available if we move to a third type of choice. Consider the following case.

A landslide occurs, trapping people in a mine. They are trapped in two separate caves. In one of them, there is one person. In the second, there are a hundred. Experts reckon that there will shortly be a second landslide, which will make it quite impossible to rescue any of the people involved. They advise that there is just about enough time to rescue either the one or the hundred, but definitely not to rescue people in both caves.

Here too, it seems irresponsible for rescuers to decide to save the person isolated in her cave, rather than the hundred. Is this conclusion reached because of the weight of numbers? That may be one reading of it. But there are reasons to think that it is not as simple as that. Say that the landslide had occurred slightly differently, and there were forty one people in the first cave and sixty in the second, rescuers might concentrate hard on other salient aspects, such as the respective accessibility of both caves. Indeed, in justifying their decision to the victims, it is much more natural for the rescuers to use this reason (accessibility), rather than saying: 'Unfortunately, you are only forty one, and they are sixty.'

It would be hasty therefore in this example as in the others to assume that it is best explained by SGN. But how else might we explain it? Given the urgency, the difficulty and the uncertainties of the rescue, it is difficult for us not to be moved by the thought that the rescue may fail. The caves may collapse before there is a chance to save anyone. Or there may be enough time to save some, but not all of the hundred people trapped in the second cave. So *in the circumstances*, we think of the case as involving a choice between possibly meeting *no* claim, and meeting at least *some* claim. As a result, what are maximized are our chances to meet some claim rather than none. So we think that we should direct efforts towards the cave with a hundred people, in the hope to save at least some.

Indeed, just such an intelligible reason is what may move us *not* to save the greater number in cases considered earlier, cases

of one versus five. Suppose you are in a boat facing a choice between rescuing five people stranded on a rock, and one person stranded on another, and that the water is raising quickly, so you do not have enough time to save them all. You are not certain that your boat will be able to rescue the five without eventually sinking under the excess of weight in the middle of the ocean. Your worry is that you may end up saving no one, therefore wasting your capacity to help. So you go and rescue the single person.

What this shows is that it is very difficult to maintain pure intuitions about cases in which everything is symmetrical apart from the numbers involved. What the intermediate cases have in common with the examples from social policy is that it is doubtful that we can screen off our reactions to the descriptions from our background assumptions about the actual resources likely to be available in the circumstances, and other features of the situation which are liable to introduce considerations which can tip the balance in one direction rather than the other, without numbers themselves directly contributing to the decision one makes. The only real test for the significance of numbers *per se* is precisely to employ the pure (and artificial) examples we considered in the first half of the paper. But, as we saw, in those cases it is just not obvious that there is a *principle* of SGN. And as we have seen, nothing about these further cases shows a distinctive role for such a principle either.

IX

To conclude: in this paper, I have set out to throw doubt on the idea that our principles of social policy are in fact directly sensitive to numbers of people involved. There are cases in which we have a duty to help what is, in fact, the greater number, but this is not so because they are more numerous. Numbers seem to be salient in such cases only because we have a way of describing them in short hand, where the reasons which really move us are concerns related, but not identical, to the ratios of the people involved.

I have offered evidence for this conclusion by moving back and forth between the individual perspective and that of social

policy. We might hope that a pure case in which nothing else but numbers is salient would establish the basis for a social principle. However, although the pure case of an agent facing a choice between saving one and saving two people allows us to isolate numbers, when we do that it doesn't seem that a demonstration of a duty to save the greater number is possible. No one is wronged by saving either.

In the second half of this paper, I have considered cases in which numbers really do seem to have an impact on our decisions at the level of social policy. Many have been tempted to suppose that the duty which would be discovered in the more abstract cases explains our intuitions in the realm of distributive justice. That is not my diagnosis. I have argued that there is a significant difference between the two kinds of case—namely that cases in the social sphere impinge on the question of the scarcity of common resources, and consequently the reasonable threshold on one's claims of need.²⁰ Now, this thought too can be interpreted in a consequentialist way—one may think that the only way to address a competition of needs is to aggregate them. Over the last few pages, I began to sketch an alternative way of thinking about reconciling individual need and common scarcity. Particularly in the context of formulating social policy, we need to find a perspective which pays proper attention both to individuals' partial concerns and to their willingness to join together in non-selfish schemes of co-operation.

*Department of Philosophy,
University College London
Gower Street
London WC1E 6BT
v.munoz@ucl.ac.uk*

20. The urgency and severity of claims of need in the social case sometimes overshadow how we think of the individual agent having to decide whom to give vital aid to. Thus some find it tempting to think of the aid given by the individual agent, too, as a scarce resource. This is a mistake. There may be several other people equipped to give aid. Each of them has the same duty to help for as long as help is needed and/or someone else has discharged this duty. None of them is a scarce, common, resource.

REFERENCES

- Anscombe, G. E. M. (1967): 'Who is Wronged?' *The Oxford Review*, n. 5, pp. 16–17.
- Foot, Philippa (1967): 'The Problem of Abortion and the Doctrine of Double Effect.' *The Oxford Review*, n. 5, pp. 5–15.
- Griffin, James (1986): *Well-being: Its Meaning, Measurement and Moral Importance*. Oxford: Clarendon Press.
- Kamm, Frances (1993): *Morality, Mortality*. Oxford: Oxford University Press, Vol. 1.
- Kolakowski, Leszek (1999): *Freedom, Fame, Lying and Betrayal. Essays on Everyday Life*, London: Penguin.
- Parfit, Derek (1978): 'Innumerate Ethics', *Philosophy and Public Affairs* (4), vol. 7, pp. 285–301.
- Parfit, Derek (2003): 'Justifiability to Each Person,' *Ratio* (XVI) December 2003.
- Rawls, John (1971): *A Theory of Justice*. Cambridge Mass., Harvard University Press.
- Scanlon, T. M. (1998): *What We Owe to Each Other*. Cambridge Mass., Harvard University Press.
- Taurek, John M. (1977): 'Should the Numbers Count?' *Philosophy and Public Affairs*, Vol. 6, pp. 293–316.