

Imperfect Duties, Group Obligations, and Beneficence

S. Andrew Schroeder

assistant professor of philosophy, Claremont McKenna College
aschroeder@cmc.edu

Abstract

There is virtually no philosophical consensus on what, exactly, imperfect duties are. In this paper, I lay out three criteria which I argue any adequate account of imperfect duties should satisfy. Using beneficence as a leading example, I suggest that existing accounts of imperfect duties will have trouble meeting those criteria. I then propose a new approach: thinking of imperfect duties as duties held by groups, rather than individuals. I show, again using the example of beneficence, that this proposal can satisfy the criteria, explaining how something can both have the necessity characteristic of *duty*, while also allowing agents the latitude which seems to attach to *imperfect* duties.

Keywords

beneficence; collective responsibility; discursive dilemma; group agency; imperfect duties; Murphy, Liam; Singer, Peter

I. Introduction

Like many things philosophers write about, imperfect duties are easy to point to but hard to define. Imperfect duties are those – paradigmatically beneficence, charity, gratitude, mercy, and the like – that seem to carry with them a kind of latitude. Most people think that while the duty of charity requires us to give something to the less well-off, it doesn't require that we give as much as we can, and it doesn't require that our donations go to any particular person or organization. Similarly, we may be obliged to be merciful, but the duty doesn't specify the precise form our mercy must take, nor does it say exactly when or to whom we must exhibit it. On the other hand, perfect duties, like the duties to keep promises and refrain from violence, don't seem to allow the same discretion. The duty to repay a loan obliges the borrower to repay the precise amount agreed on, to a specific person,

often at a specific time. The obligation to refrain from violence forbids doing specific things to specific people.

That, though, is just to gesture at the distinction. Unfortunately, there's very little philosophical agreement about exactly what imperfect duties are. Rainbolt, for example, identifies eight distinct accounts in the philosophical literature and nevertheless leaves out several important proposals.¹ In fact, about the only thing agreed upon by (nearly) all parties to the discussion is that *if* there are any imperfect duties, then beneficence is one of them. In this paper, I propose to look at the duty of beneficence, to see if there is any way we can make sense of the idea that it is imperfect. My hope is that, since beneficence is typically taken to be the imperfect duty *par excellence*, this will shed light on the character of imperfect duties more generally – though that is a task I'll only be able to address briefly in this essay.

I'll begin my argument by noting three criteria that any plausible account of beneficence as an imperfect duty should meet (§II), then I'll quickly survey the traditional accounts of imperfect duties, suggesting that each will have trouble satisfying those criteria (§III). Next, I'll present a proposal that in its broadest outlines has been put forward by several philosophers, most notably Jonathan Cohen and Liam Murphy: that beneficence can be conceived of as a duty held primarily by groups, rather than individuals (§IV). I'll argue, though, that neither Cohen's nor Murphy's proposal works, because neither takes seriously enough the idea of a group obligation (§§V-VII). I'll then suggest one way that might be done, which promises to satisfy the three criteria (§§VIII-IX). Finally, I'll conclude by very briefly returning to the subject of imperfect duties more generally. I'll suggest that there is reason to think that at least some other classically imperfect duties are also cases of group obligation (§X).

II. Three Criteria

The reason that there have been so many different accounts of imperfect duties is that they all face serious – in my opinion, fatal – problems. The basic idea motivating a theory of imperfect duties is the intuition that there are certain types of action which we ought to do, don't need to do as much as possible, and have some discretion as to the particular occasions and

¹ George Rainbolt, 'Perfect and Imperfect Obligations', *Philosophical Studies* 98 (2000), pp. 233-56.

manner in which we do them. In the case of beneficence, many people accept that we ought to give to the poor, that we don't need to give as much as we can, and that we have some discretion about whether to give to Oxfam or to the Red Cross and about precisely how much to give. This discretion isn't total – giving one percent of your income is not enough, for example – but some leeway is allowed.

The most obvious problem with that initial characterization is that it leaves unclear precisely what the duty is. To say that something is a duty is to say that it is required – that we *must* do it, or that it is *necessary* that we do it. This modality allows for neither discretion nor vagueness. Just as a proposition can't be a "little bit" necessary or "kind of" contingent, an action must either be required or not. This means that (unlike other forms of advice) requirements must have precise boundaries, or must specify precise limits on behavior. Requirements mark the very bottom of the range of acceptable action.² If, therefore, imperfect duties are truly to be *duties*, they must satisfy this condition:

SPECIFICATION. *An account of imperfect duties should precisely identify what is required, what an agent must do.*

(Notice that this is not obviously true of other moral concepts. Nothing in the nature of goodness or virtue, for example, seems to require such precision.)

This criterion is clearly in tension with the original intuition behind imperfect duties, that they're ones that allow the agent some latitude. If we precisely specify what is required, then it doesn't seem like agents will have any latitude. We therefore next need to formulate a criterion that captures the type of latitude which is supposed to be characteristic of imperfect duties. Since in this paper I'm focusing on the duty of beneficence, I suggest that we do so by looking at a pair of classic cases, *Rescue* and *Aid*, which are usually taken to illustrate the distinction.³

First, imagine that you're walking past a shallow pond and see a child drowning. You could easily rescue the child at no danger to yourself, though it would mean ruining your new \$250 suit. Contrast this with a case of aid: many people are starving halfway around the globe. You could provide meaningful, long-term relief to any one of them by writing a \$250 check to

² I discuss this analysis of duty further in my *Divorcing the Good and the Right* (Ph.D. dissertation, Harvard University, 2009), ch. 1.

³ The cases were of course made prominent in Peter Singer, 'Famine, Affluence, and Morality', *Philosophy and Public Affairs* 1 (1972), pp. 229-43.

Oxfam. (\$500 would help two, and so forth.) Singer famously argues that there is no moral distinction between the cases, and that our duty in Aid is just as strong as our duty in Rescue. Proponents of imperfect duties, however, are generally motivated by the thought that cases like Rescue are more demanding than cases of aid. We're morally required to rescue the drowning child, but we're not obliged to send every extra \$250 we earn to Oxfam. And, further, an agent has a certain amount of freedom in responding to aid cases (for example, in choosing on what occasions to respond). Now, I don't think it is obvious that these intuitions are correct, but since they are basically fixed points among people who accept imperfect duties, let us say:

LATITUDE. An account of imperfect duties should specify the duty in Aid (and relevantly similar cases) so as to make it in some respect less onerous than the duty in Rescue (and relevantly similar cases), allowing agents some kind of discretion.

This condition asks for a formal description of the duty in Aid that shows how, consistently with SPECIFICATION, it could allow the agent latitude. Notice that it leaves open whether the duty in Rescue is categorized as perfect or imperfect. The standard analysis, I think, says that although both Aid and Rescue call for beneficence, the duty in Rescue is perfect, since it allows no latitude, while the duty in Aid is imperfect. We'd then say that the duty of beneficence, usually imperfect, becomes perfect in rescue cases.⁴ However, as Singer and others have pointed out, Rescue and Aid seem similar in many morally relevant respects. We might instead, therefore, say that the duty in Rescue is imperfect, despite the fact that it doesn't allow the agent latitude. On this view it would follow that, although imperfect duties *tend* to allow an agent latitude, latitude is not *constitutive* of imperfect duties. Both of these views are reasonable, I think, and so LATITUDE is neutral between them.

Finally, imperfect duties seem to comprise something of a natural kind, somehow different from other, perfect duties. This is spelled out in a number of ways, but one common observation is that with imperfect duties, it is possible to do more than what's required, and doing so ordinarily has moral value. You can be more gracious, merciful, charitable, or beneficent than morality demands. You can't, on the other hand, keep more of your promise than morality requires, nor can you repay more than

⁴ See Violetta Igneski, 'Distance, Determinacy and the Duty to Aid: A Reply to Kamm', *Law and Philosophy* 20 (2001), p. 612; and John Witherspoon, *Lectures on Moral Philosophy*, ed. Collins (Princeton: Princeton University Press, 1912), p. 56.

what you owe. You can, of course, do more for a friend than you've promised and you can give more money to a creditor than you agreed to, but the surplus in each case is no longer regarded as an instance of promise-keeping or debt-repayment. Rather, it is typically described as beneficence or gratitude. Let us therefore say:

IMPERFECTION. An account of imperfect duties should say what distinguishes imperfect from perfect duties, in a way that explains why it is frequently possible to do more than what an imperfect duty requires and why such excess frequently has moral value.

On many theories, the solution to IMPERFECTION is closely related to the solution to LATITUDE: imperfect duties are those that allow latitude, and that "wobble room" is what allows an agent to do more than she's required to. But, as we'll see, not every theory works that way.

We have, then, three criteria that I think any adequate account of imperfect duties should meet. The account will need to *specify* the duty in a way that preserves *latitude* for the agent, and will need to explain what makes the duty *imperfect*. Now, all three of these criteria – especially LATITUDE and IMPERFECTION – are in need of interpretation. (What kind of latitude are we looking for? What counts as "less onerous"? How sharp should the boundary between perfect and imperfect duties be?) I'll set those problems aside for now, though. In the next section, I'll very briefly sketch what I take to be the six main philosophical approaches to imperfect duties, which I'll group into three families. I'll argue that members of the first family can't explain LATITUDE and that members of the second run into trouble with IMPERFECTION. Members of the third have a number of attractive features, but threaten to merely move the bump under the carpet. My discussion certainly won't constitute a refutation of these views. But it will, I hope, suggest that existing theories have enough *prima facie* problems to make it worthwhile to seek an entirely different option.

III. Six Accounts of Imperfect Duties

In the early modern period, imperfect duties began their life as imperfect rights, in the work of Hugo Grotius. This terminology makes much more sense of the label: a perfect right includes the authorization to use force and coercion to ensure it is respected. To have an imperfect right, on the other hand, is to be genuinely owed something, but not to have

license to use force to get it.⁵ So, an imperfect right is in a sense incomplete – it doesn't come with all the "stuff" that we might want in a right. Pufendorf modified Grotius's view in several respects, but kept basically the same framework, switching primarily to the language of perfect and imperfect *duties*, rather than rights.⁶ The Grotian/Pufendorffian view remained popular throughout the early modern period – see, for example, Home and in a sense even Kant.⁷ I'll group these theories together as *no-compulsion views*.

Related to no-compulsion views are what I'll call *no-rights views*, which say that imperfect duties are those that lack correlative rights. The basic idea is that with a duty like beneficence, you are required to give to the poor, but no one can claim a right to your charity. Thus, your duty has no correlative right. Pufendorf offers a version of this view, as do philosophers otherwise as different as Mill and O'Neill.⁸

What no-compulsion and no-rights views have in common is that they don't offer a promising way of handling LATITUDE. Insofar as they satisfy SPECIFICATION and say exactly what an agent must do – give \$1000 to Oxfam, say – there doesn't seem to be the discretion that LATITUDE calls for. We can say that the duty to give \$1000 is non-compellable, or that there is no specific person with a right to that money, but neither of those facts lessens or mitigates the obligation on the agent. She still must give \$1000 to Oxfam. Therefore, these views don't generate the kind of latitude which is taken to be characteristic of imperfect duties. The duty of aid is just as binding on an agent as the duty of rescue.

⁵ Hugo Grotius, *On the Law of War and Peace* (1625). See e.g. §§II.I.i-vii, II.XI.iii-iv, II.XXII.xvi.

⁶ Samuel Pufendorf, *On the Law of Nature and of Nations* (1672). See e.g. §§I.i.19-20, I.vii.x-11, III.iv.5-7. Although Pufendorf spoke primarily in terms of duty, other philosophers continued to speak of perfect and imperfect rights. See e.g. Francis Hutcheson, *An Inquiry into the Original of Our Ideas of Beauty and Virtue* (1726), §II.vii; and Adam Smith, *Lectures on Jurisprudence* (1762), lecture of 24 Dec. 1762. Imperfect rights still frequently appear in contemporary legal scholarship, though the mention is usually brief, because imperfect rights are frequently defined as those not legally compellable. (See Jonathan Law and Elizabeth A. Martin, 'Perfect and imperfect rights', *A Dictionary of Law*, 7th ed. (Oxford: Oxford University Press, 2009). But contrast 'Right', *Black's Law Dictionary*, 7th ed., ed. Bryan Garner (St. Paul, MN: Thompson/West, 1999).)

⁷ Henry Home, *Essay on the Principles of Morality and Natural Religion* (1751), §II.iii. Immanuel Kant, *The Metaphysics of Morals*, tr. Mary Gregor (Cambridge: Cambridge University Press, 1996), *Ak.* 6:386.

⁸ John Stuart Mill, *Utilitarianism* (Indianapolis: Hackett, 2001), pp. 49-50; Onora O'Neill, 'Children's Rights and Children's Lives', in *Constructions of Reason* (Cambridge: Cambridge University Press, 1990), p. 189; Pufendorf, *On the Law*, § III.iv.6.

The second family of accounts of imperfect duties tries to directly tackle LATITUDE by saying that imperfect duties are those that leave unspecified the occasions or conditions of their fulfillment. For example, you might have a duty to give to charity, without that duty specifying how much to give, to whom to give it, and so on.⁹ The problem here is obvious: if what the duty requires is unspecified, we haven't met SPECIFICATION. If there's nothing you *must* do, there's no actual duty. Imagine asking a series of questions like this: is it true that the duty requires me to give \$X, or is it permissible to give less? If the answer is always that giving \$X is not required, then there's no duty; there's no amount you *must* give. If, on the other hand, there is some X for which it is not permissible to give less, take the highest such X. The duty is then just one to give \$X – there's no sense in which you *must* give more, however virtuous such an action might be. We've lost the latitude with which this proposal began.

This line of thinking leads to what I take to be the most popular accounts of imperfect duties: *disjunctive* and *quota views*. These views try to satisfy SPECIFICATION and LATITUDE at the same time, by saying that imperfect duties are duties to do act < A or B or C or D >, or to do some specified number of acts from set {A, B, C, D}. Whereas a perfect duty to repay a loan requires giving \$500 to Mary on Monday, an imperfect duty of aid might require giving \$1000 to Oxfam *or* giving \$2000 to the Red Cross *or* spending 50 hours volunteering at a food bank.

Many philosophers have held views like this, with a notable early account being Richard Price's.¹⁰ Michael Stocker points out the problem with any such proposal, though: *all* duties turn out to be imperfect.¹¹ Although your duty to give \$500 to Mary seems to be non-disjunctive, in fact you have many options: you could give her one hundred \$5 bills, five \$100 bills, or you could write her a check. You could give her the money with your left hand, with your right hand, or deposit it directly into her bank account. You could deliver it with a smile on your face, or with a scowl. And so forth. Stocker argues that virtually all duties will be infinitely disjunctive. We therefore haven't yet satisfied IMPERFECTION: there is no significant difference between duties like beneficence and loan-repayment. Now, Stocker

⁹ This may be what many advocates of no-rights views really have in mind: since there is no specific person with a right to your charity, your duty of beneficence leaves it unspecified what, exactly, you should do with your donation.

¹⁰ Richard Price, *A Review of the Principal Questions in Morals*, reprinted in *British Moralists*, vol. II, ed. D.D. Raphael (Indianapolis: Hackett, 1991), §§720-723.

¹¹ Michael Stocker, 'Acts, Perfect Duties, and Imperfect Duties', *Review of Metaphysics* 20 (1967), pp. 507-17.

and Rainbolt offer revisions to the disjunctive view, arguing that it can be rehabilitated by saying that only certain types of disjunctiveness make a duty imperfect: giving with the right vs. the left hand doesn't intuitively seem on a par with giving to Oxfam vs. giving to the Red Cross. But spelling out exactly what the difference is is tricky and leads to a number of problems, as Stocker and Rainbolt admit.¹²

The final two schools of thought on imperfect duties view them as duties to adopt general maxims,¹³ or to cultivate dispositions.¹⁴ While you have a perfect duty to repay a loan, you have an imperfect duty to adopt a general maxim of beneficence or to cultivate a charitable disposition. (These views are of course distinct, but I group them together for reasons which should become clear.) I think there are a number of virtues to this kind of approach. Foremost among them is that we can meaningfully talk about whether even the very poor have satisfied their duty of beneficence. On the previous views, where imperfect duties are duties to perform actions, someone who is very poor will neither have satisfied nor violated the duty of beneficence: not having any resources to give away, the duty won't apply to her at all. On a maxim or dispositional view, however, we can ask whether or not she has the appropriate attitude. Has she made it her end to help those in need, when she's able? Has she cultivated a charitable disposition? Is she prepared or willing to help, should the opportunity arise? We can ask those questions of her, even though we know that she's never been in position to employ those maxims or dispositions. This strikes me as a very attractive result.

Nevertheless, there is an important respect in which maxim and dispositional views merely relocate the problem of imperfect duties, without solving it. Both maxims and dispositions can be differentiated by their content. A disposition to give \$1000 to charity is not the same disposition as one to give \$2000 to charity. Similarly, a maxim of spending every Sunday helping the poor is not the same as a maxim of spending one day each week helping the poor. In order to satisfy SPECIFICATION, a maxim or dispositional view will need to say exactly which maxim(s) or disposition(s) an agent is

¹² See Stocker ('Acts'), Rainbolt ('Perfect and Imperfect'), and also Daniel Statman, 'Who Needs Imperfect Duties?' *American Philosophical Quarterly* 33 (1996), pp. 211-14.

¹³ Kant is obviously the source of this view. See e.g. *Metaphysics of Morals*, Ak. 6:388-9.

¹⁴ This view is not often explicitly offered as a solution to the problem of imperfect duties, but it frequently plays that role. See e.g. Richard Miller ('Beneficence, Duty and Distance', *Philosophy and Public Affairs* 32 (2004), pp.357-83) on beneficence, and Roslyn Weiss ('The Moral and Social Dimensions of Gratitude', *Southern Journal of Philosophy* 23 (1985), pp. 491-501) on gratitude.

required to have. That will in turn involve specifying the content of the maxim or disposition. The problem is that doing this requires spelling out the conditions in which a morally-motivated agent will act.¹⁵ In particular, the maxim or disposition will need to invariably call for action in rescue cases, but call for action only sometimes in aid cases. But this is just the challenge posed by the LATITUDE criterion. Maxim and dispositional views therefore can't, by themselves, solve the problem of imperfect duties.

This can be seen more clearly, perhaps, from the perspective of an agent engaged in deliberation. Imagine that I, a morally-motivated Kantian, am faced both with a drowning child and with the chance to donate to Oxfam. I know that I'm required to have a maxim of beneficence, but faced with the drowning child and the chance to donate to Oxfam, I need to know something specific about the maxim: does it require me to save the child, and does it require me to donate to Oxfam? In order to satisfy the criteria we've established, the maxim will need to be *specific* enough to rule out omitting the rescue (i.e. it can't merely say "help some people sometimes"), but it will also need to allow a certain kind of *latitude* in aid cases (i.e. it can't say "help people whenever they're in need"). In other words, we get something that looks an awful lot like the original problem of imperfect duties, when we try to determine the content of the maxim in question.¹⁶ (A similar point could be made concerning dispositional views.)

There are, of course, further ways that the maxim or dispositional theorist could respond to these concerns, but I think that what I've said suffices to show that such views, like the other accounts we looked at, face non-trivial problems. I therefore think it is worth pursuing another, far less familiar way of understanding imperfect duties: as duties held by groups.¹⁷

¹⁵ Robert Noggle rejects this, claiming that a Kantian view need not specify the precise end to be adopted – and says that this is a *virtue* of the approach. But as far as I can tell, he only says this because he assumes that the alternative is something like a disjunctive or quota view. ('Give Till It Hurts? Beneficence, Imperfect Duties, and a Moderate Response to the Aid Question', *Journal of Social Philosophy* 40 (2009), pp. 1-16.)

¹⁶ See Barbara Herman, *The Practice of Moral Judgment* (Cambridge, MA: Harvard University Press, 1993), pp. 65-66; Thomas Hill, *Human Welfare and Moral Worth* (Oxford: Oxford University Press, 2002), p. 208; and especially Karen Stohr, 'Kantian Beneficence and the Problem of Obligatory Aid', *Journal of Moral Philosophy* 8 (2011), pp. 45-67 for discussions of the Kantian view that try to address this worry. And see Statman, 'Who Needs Imperfect Duties?', pp. 216-17, for an objection along the lines I've outlined here. The next move for a Kantian, I think, would be to try to combine, say, a maxim view with a quota view: agents have a duty to adopt a general maxim that has the form of a quota. This seems to me to be an implausible strategy for a number of reasons, which I don't have the space to go into here.

¹⁷ Although I've presented these three families as distinct, that's a gross oversimplification. Kant's view, for example, arguably fits in all three categories. Nevertheless, I think

IV. Beneficence as a Group Obligation: Murphy and Cohen

The idea that beneficence is best thought of as a duty belonging to groups is, in one form, quite commonplace. It is often said that it's the government's or the church's job to care for the poor. The more robust view, though, that beneficence is a duty that *we* have as a group is mentioned in several places,¹⁸ but it gets spelled out most clearly in Jonathan Cohen's article "Who is starving whom?" and Liam Murphy's important book *Moral Demands in Nonideal Theory*, which give broadly similar accounts.¹⁹

Cohen presents a number of cases like this: twenty people live on a desert island. Ten have plenty of food; ten have very little. Eight of the rich residents each give away some of their food to a poor resident. As a result, eight of the poor are fed and two are still near starvation. Surely, Cohen says, if the two are to blame anyone, they ought to blame the two rich ones who gave nothing. The responsibility for their hunger lies with those two, not with the eight who gave to the poor – even if those eight still have a surplus of food. Cohen concludes that your duty to the poor is a duty to do your share, where your share is determined by thinking about what you would be required to do, were everyone to do her part.

Murphy arrives at basically the same conclusion, though his argument runs somewhat differently. For reasons that it would take us too far afield to consider, he claims that because beneficence is an agent-neutral duty, it follows that "[w]e could say, somewhat tendentiously, that a principle of beneficence is directed to agents as a group, whereas other [agent-relative] moral principles are directed to agents individually." However,

[T]here are two questions that naturally arise for a principle directed to people as a group that do not arise for principles directed to each person individually. It can be asked, first, how responsibility for promoting the common aim should be distributed among the members of the group and, second, how an agent's responsibility is affected when another member of the group shirks her responsibility. In other words, when a group pursues a common aim, there are questions about the distribution of responsibility and about the responsibility of complying agents in situations of partial compliance.²⁰

grouping theories in this way is useful, even if some accounts end up being placed in multiple groups.

¹⁸ See Jan Narveson, *Morality and Utility* (Baltimore: Johns Hopkins University Press, 1967), pp. 231-8; and Bertrand Bandman, 'Do Future Generations Have the Right to Breathe Clean Air?' *Political Theory* 10 (1982), pp. 95-102.

¹⁹ Jonathan Cohen, 'Who is Starving Whom?' *Theoria* 47 (1981), pp. 65-81; Liam Murphy, *Moral Demands in Nonideal Theory* (Oxford: Oxford University Press, 2003).

²⁰ Murphy, *Moral Demands*, p. 75.

Murphy argues that the answer to the first question will be given by what most promotes the good. In other words, we should divide group responsibility in whatever way would be best. The second question is answered by what Murphy calls the Compliance Condition. The Compliance Condition ends up being very long and detailed, mainly due to problems in finding a good metric for demandingness, but Murphy's loose formulation will suffice for our purposes:

An agent-neutral moral principle should not increase its demands on agents as expected compliance with the principle by other agents decreases. Demands on an agent under partial compliance should not exceed what they would be (all other aspects of her situation remaining the same) under full compliance from now on.²¹

The basic idea is that morality should demand no more of me than it would demand of me if everyone were fulfilling their duties.

Thus, we have a spelled-out version of the proposal that beneficence is a duty attaching to groups: a group (say, society) has a duty to feed the poor. As a member of that group, I'm assigned a share based on what would be best (Murphy) or on considerations of fairness and equity (Cohen). That, and no more, is my duty.

This proposal begins by recognizing a duty of beneficence that has a precise and non-arbitrary extent: the duty is to feed all the poor (as opposed to, for example, feeding some of the poor). Therefore, this view will satisfy SPECIFICATION. The duty, however, is held by the group, which gives our answer to IMPERFECTION: imperfect duties are those an agent holds in virtue of its membership in a group, as a means of discharging an obligation of the group's.²² An individual will be able to do more than what is required of her whenever others don't do their part, and doing so will typically be of moral value, since it will be an instance of duty-fulfillment (though not the fulfillment of her own duty). In other words, being more beneficent than one is required to will typically amount to doing someone else's duty for her and will therefore often have moral value. For these reasons, we can expect that any view of imperfect duties as group obligations will have no trouble meeting SPECIFICATION and IMPERFECTION.

The real action will be around LATITUDE. According to Cohen and Murphy, by the time the group's duty is filtered down to the level of individual obligation, each of us has a duty only to do our share. In Rescue, if

²¹ Ibid., p. 77.

²² Notice that this means that group agents can also have imperfect duties, if they are members of still larger groups (super-groups) with obligations.

you walk by the child and you're the only one around, then "your share" is to save the child, and so you're required to save the child. Your duty with regards to famine relief, though, might only be to donate a thousand dollars a year, not to save every life you could. Thus, we seem to have satisfied LATITUDE.

The success is short-lived, however. Although they may solve the simple Rescue case, neither Murphy's nor Cohen's proposal can handle a number of other, equally compelling cases. For example: Megan and Nick are walking by a shallow pond. Two children are drowning in the pond. Megan runs in and saves one, but Nick ignores the situation and keeps walking. Is Megan required to run back in to save the other child?²³ According to Cohen and Murphy, the duty here is divided between the two: Megan and Nick each have a duty to save one child. Megan isn't required to do more than what she would be required to do, were Nick to do his duty. So, it is permissible for Megan to continue on after saving one child, leaving the other to drown. Both Murphy and Cohen acknowledge that their views have this consequence, and both seem to wish that they could avoid it. However, each concludes that there is no way to modify his theory to handle the case without returning to the implausibly over-demanding view of beneficence, and so each bites the bullet.²⁴

If thinking of beneficence as a group obligation necessarily had this consequence, it wouldn't be an especially promising proposal. In an effort to conform to our intuitions on certain problem cases, it diverges wildly from them in others. Now, Murphy in particular has a very long and careful argument for his Compliance Condition, so we shouldn't reject his theory out of hand, but nevertheless it no longer seems *prima facie* more attractive than the other views of imperfect duties we discussed in §III. Fortunately, there's reason to think that Murphy and Cohen haven't explored all the territory.

²³ This case is from James Rachels, 'Killing and Starving to Death', *Philosophy* 54 (1979), pp. 162-3.

²⁴ Cohen, 'Starving', pp. 76-8; Murphy, *Moral Demands*, pp. 127-33. Actually, Murphy's analysis of this scenario is a bit more complicated, in ways that sometimes yield more plausible but sometimes more implausible results. For example, Murphy acknowledges that on his view it follows that anyone who would on balance be better off under full compliance with morality isn't required to help *at all* in the real world. So, if you would have a higher level of well-being under full compliance, you have *no* obligation to do anything for anyone else. Since it is not entirely implausible to suppose that many people would be better off in a world of full compliance (as compared to the actual world), Murphy's principle of beneficence would never require such people to save a drowning child, even if doing so meant no more than tossing a life preserver. (Murphy tries to soften the blow of this conclusion a bit, but his efforts don't strike me as very successful.)

V. Group Agency

In order to see why, it is necessary to take a brief detour into the philosophy of (group) agency. I take it to be relatively uncontroversial that groups can act, intend, desire, believe, and the like. We talk that way all the time, and there are many philosophical theories of group action, group intention, and so forth. What is controversial, however, is whether these group states must be reducible to similar states in the group's members. Quinton expresses the positive view:

We do, of course, speak freely of the mental properties and acts of a group in the way that we do of individual people. Groups are said to have beliefs, emotions, and attitudes and to take decisions and make promises. But these ways of speaking are plainly metaphorical. To ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members... To say that the industrial working class is determined to resist anti-trade union laws is to say that all or most industrial workers are so minded.²⁵

Many philosophers seem to accept something like Quinton's position, though not always as simplistically as he puts it.

Other philosophers, however, take a non-reductive approach, according to which certain group states are not reducible to similar states in the group's members. For example, Searle says,

[C]ollective intentional behavior is a primitive phenomenon which cannot be analyzed as just the summation of individual intentional behavior; and collective intentions expressed in the form "we intend to do such-and-such", and, "we are doing such-and-such" are also primitive phenomena and cannot be analyzed in terms of individual intentions expressed in the form "I intend to do such-and-such" or "I am doing such-and-such".²⁶

These non-reductive views seem to me to be right, though I won't argue for that claim here. For the remainder of this paper I'll simply assume that groups can instantiate at least some properties of agents in a non-reductive way. Most importantly, I'll take it for granted that groups can properly be held morally responsible, or can have moral obligations, non-reductively.

²⁵ Anthony Quinton, 'Social Objects', *Proceedings of the Aristotelian Society* 75 (1975), p. 17.

²⁶ John Searle, 'Collective Intentions and Actions', in *Intentions in Communication* (Cambridge, MA: MIT Press, 1990), p. 1. Other non-reductive views can be found in Margaret Gilbert, *On Social Facts* (Princeton: Princeton University Press, 1992); Russell Hardin, *Morality Within the Limits of Reason* (Chicago: Chicago University Press, 1988); and Austen Clark, 'Beliefs and Desires Incorporated', *The Journal of Philosophy* 91 (1994), pp. 404-25; among many others.

That is, I'll assume that we can sensibly say that a group is morally responsible for something, where that isn't shorthand for saying that all or most members of the group are individually morally responsible for that thing. This assumption shouldn't be unfamiliar – we often, for example, assign moral responsibility to nations in a way that is at least arguably non-reductive.²⁷

VI. The Discursive Dilemma

Before returning to the main argument, one more detour is necessary, to discuss a phenomenon called the doctrinal paradox or discursive dilemma.²⁸ Consider a group of three people charged with reaching judgments on three propositions, P, Q, and R, and on their disjunction, P or Q or R. Individually, they judge as follows:

	P?	Q?	R?	P or Q or R?
A	no	no	yes	yes
B	no	yes	no	yes
C	yes	no	no	yes

Notice that each member of the group has a consistent set of judgments: each endorses the disjunction if and only if she endorses at least one disjunct. If we follow Quinton's proposal, though, and equate the judgment of the group with the majority judgment on each proposition, the group is in the irrational position of rejecting P, rejecting Q, and rejecting R, yet *accepting* their disjunction. The group is therefore faced with the prospect of: (1) retaining the irrational set of judgments, (2) suspending judgment on one of the propositions, or (3) endorsing some claim rejected by the majority of its members – i.e. following some non-majoritarian procedure for determining the group's judgments.

²⁷ See e.g. Onora O'Neill, 'Global justice: whose obligations?' in *The Ethics of Assistance*, ed. Chatterjee (Cambridge: Cambridge Univ. Press, 2004), pp. 242-259, and David Miller, 'Holding Nations Responsible', *Ethics* 114 (2004), pp. 240-68.

²⁸ This phenomenon was first discussed by Lewis Kornhauser and Lawrence Sager, looking at problems facing multi-judge courts ('Unpacking the Court', *Yale Law Journal* 96 (1986), pp. 82-117, and subsequent works). Philip Pettit's recent work has brought it to the attention of philosophers. See 'Groups with Minds of Their Own', in *Socializing Metaphysics*, ed. Schmitt (Oxford: Rowman & Littlefield Publishers, 2003); 'Responsibility Incorporated,' *Ethics* 117 (2007), pp. 171-201; and 'Rationality, Reasoning and Group Agency', *Dialectica* 61 (2007), pp. 495-519, which overlap to a significant degree.

Pettit argues that real groups, like corporations or clubs or political parties, will frequently face problems like this. In such cases, the consequences of (1) will sometimes be dire, and (2) will often not be an option. For example, suppose that A, B, and C are the governing board of a corporation, and that P, Q, and R are separate cost-cutting measures – perhaps laying off workers, reducing benefits, and deferring compensation. All members of the board agree that some cost-cutting measure must be enacted, or the corporation will fold. In other words, they endorse the disjunction. The group can't merely suspend judgment on one of the propositions, nor can it retain the (irrational) judgments reached by a majority vote on each issue. To do either of those would itself doom the company to bankruptcy. Therefore, the board will be forced to adopt some version of (3), above.

So, groups will sometimes be forced to find another, non-majoritarian way to arrive at group judgments. How should they do that? Various impossibility theorems prove that no procedure – majoritarian or otherwise – that respects certain plausible constraints can guarantee a rational verdict.²⁹ If a group is to have a rational set of judgments in cases like this, those judgments can't be formed through any bottom-up aggregation of the judgments of the members. Instead, resolving this sort of inconsistency requires some kind of non-mechanical, group-level process – plausibly, one that looks analogous to the processes involved in individual reasoning.³⁰

Pettit concludes from this that group rationality is best thought of in a non-reductive way. I've already said that I'm *assuming* that group rationality can be thought of non-reductively, so the conclusion I'd like to draw is more narrow. The impossibility results proven by Pettit, List, and others show that if a group wants to avoid irrationality at the group level, it will need to do so by way of a group-level process. We might have thought that if we found group-level irrationality, we could, at least in principle, trace it to irrationality at the individual level – or, equivalently, that a rational group judgment could always be “read off” of a collection of rational individual judgments. But the impossibility results show that's not right.

²⁹ See Christian List and Philip Pettit, 'Aggregating Sets of Judgments: An Impossibility Result', *Economics and Philosophy* 18 (2002), pp. 89-110, and the works cited therein.

³⁰ See Pettit ('Personal and Sub-personal Reason') for a survey of several options. He endorses what he calls a “straw vote procedure.” Basically, the inconsistent propositions are identified, discussed, and put up for another vote, in the hope that some individuals will revise their judgments, perhaps in the spirit of “doing what's best for the group”. If that doesn't work, Pettit suggests the inconsistency be “quarantined”, though it's not entirely clear how that is to be done.

Resolving group irrationality requires a process – it requires that the group *do* something.

VII. Non-Reductive Group Obligations

Now we can return to Cohen and Murphy. The problem with their views is that each construes the group reductively, at least as far as obligation goes. To say that a group is obliged to do X, for both of them, is essentially to say that the obligations of the members add up to X. That's not to say that groups play no role in Cohen's and Murphy's accounts – in both cases, groups play an essential role in the argument. But at the level of obligation, both Cohen and Murphy offer reductive views: the group obligation gets neatly decomposed into individual obligations and doesn't have any further role. Once we recognize the possibility of a non-reductive group obligation, however, we can see that there is a wider range of possibilities. A group may have an obligation that is best thought of "top-down", rather than "bottom-up". To say that a group has a duty to X may be to say more than that all the members of the group individually have obligations that add up to X. Individual obligations may not exhaust the deontic story.

On reflection, this should have been obvious. Cohen and Murphy's neat way of dividing up obligations is artificial in at least two respects. Consider an everyday case of group responsibility. Alice, Bob, and Carol have played together for years as a piano trio. They frequently rehearse in David's house, because of his drawing room's exquisite Steinway and wonderful acoustics. In return, the trio agrees to paint David's house. Just as work is to begin, Alice falls ill and is unable to paint. David returns home to a two-thirds painted house and asks Bob and Carol what happened. It would be rather unsatisfying for him to hear, "Well, don't look at us – we each did our third! Alice is the one you have a complaint against. Talk to her."

First, the minor point: it's not obvious that having each paint a third is the way the obligation should be divided. Perhaps when the trio formed cellists were in great demand, and so in order to coax Alice to join their group, Bob and Carol told her that she'd have to shoulder only a quarter of the group's responsibilities. Or perhaps the group, being of a rather whimsical spirit, ordinarily divides all burdens according to a dice roll. The point is that though one-third each may seem like the natural or salient division of labor, there's no reason that the group has to operate in that way. It may be, therefore, that Alice was never on the hook for one third.

Much more importantly, however, even if the trio had a policy of dividing burdens evenly, divisions of responsibility within a group aren't so neat, viewed from outside the group. Even if Alice owes it to the group to paint one third of the house, if she doesn't follow through it seems perfectly appropriate for David to hold *the group* responsible. He'd likely reply to Bob and Carol, "Look, you and Alice can sort this out however you want – that's your business. As far as I'm concerned, though, you *all* owe me a painted house and I haven't gotten it yet. None of you are welcome to rehearse here until the painting is complete." Though it may be the case that Alice is the one who let the group down, the whole group has let David down, and the responsibility relation here isn't transitive.

This case is not at all unusual. When a member of a choir, basketball team, or consulting firm is having a bad day, the other members of the group are usually expected to pick up the slack. However we ultimately analyze these situations, I think it is clear that Cohen and Murphy's suggestion of a complete division of obligation doesn't square with the way we ordinarily think about group responsibility. We should therefore look for a better, "top down" or holistic way to model the function from group responsibility to member responsibility.

VIII. How to Think about Beneficence as a Group Obligation

Suppose some group has a duty of beneficence. How should we understand that obligation? Initially, we might try to resist any division of the obligation, saying that a group obligation is held by the group alone – it doesn't filter down to the individual level at all. This, though, would be implausible, for a number of reasons. Most obviously, it wouldn't solve the problem that motivated our rejection of the Cohen/Murphy proposal: if Nick doesn't rescue the second drowning child, Megan wouldn't be obliged to step in. In fact, on this view neither Nick nor Megan was obliged to help in the first place, since the obligation to save the children was held only by the group. This seems wrong.

So, our account of group obligation must allow that such obligations filter down to the individual level. But we can't have them filter down in the way Cohen and Murphy do. The worry, suggested by cases like that of Alice, Bob, and Carol, is that once we leave the Murphy/Cohen position, we'll be forced to say that group responsibilities filter down *in total* to the individual level. In other words, we'll be forced to say that if a group of which I'm a member has a duty to feed the poor, then I therefore have a duty to feed the

poor. Accordingly, if you all don't do your part, I'm obliged to do as much as I can, and there won't be any distinction between rescue and aid cases. That is, we won't have a solution to LATITUDE. Our challenge, then, is to give a philosophically compelling account of the relationship between group and member obligation that satisfies LATITUDE. I'll now sketch such an account.

An individual faced with an obligation – say, to repay a loan – first needs a plan to fulfill that obligation. But that's not enough: if you plan to drive to the bank via Main Street but find it closed to traffic, you oughtn't just throw up your hands and go home. You have an obligation to find an alternate route. In addition to having a *primary plan* for fulfilling an obligation, then, an individual should also have or be prepared to formulate a *back-up plan*, in case things go wrong. A group agent faced with an obligation will, by the same reasoning, also need both a primary plan and a back-up plan. Among the more interesting things that can go wrong for a group agent (that can't, in quite the same way, go wrong for an individual agent) is that some members of the group may not do their part, as spelled out in the primary plan. So, one crucial element of a group's back-up plan will be a specification of what to do in case of less than full compliance with the primary plan.

So far, this basically follows Murphy's analysis, which we saw earlier, in saying that a group obligation calls for an initial division of responsibility, as well as a division of responsibility in cases of partial compliance. Murphy then looks to a *theory* of group obligation to supply the answers. He thinks that general moral principles dictate the proper plan and back-up plans. I instead think we should look to the *group*. Just as there might be a number of reasonable ways for an individual to fulfill an obligation, there may be a number of reasonable ways for a group to fulfill its obligations. Faced with a collection of reasonable plans, a group (just like an individual) will need to decide which one to adopt.

What, exactly, will that entail? What is it for a group to *decide* on some course of action? The specifics will vary depending on the theory of group agency one adopts, and so I won't attempt to spell out a full theory of group decision here. (That is the subject for another essay.) But our discussion of the discursive dilemma tells us something about what that theory must look like. We might have thought that a group's decision would in some way be dependent on the decisions or preferences of its members. That is, we might have thought that we could, at least in principle, "read off" a group's decision from the mental states of its members. The impossibility results we saw earlier, however, show that this isn't right. Even given a group composed entirely of members with individually rational judgments, there is no way to "read off" a group judgment that is guaranteed to be rational.

Put another way, there is no function which takes as inputs individually rational judgments and outputs a rational (group) judgment. Instead, a group level process is, at least in many cases, required. Therefore, any adequate theory of group decision will require a group process, at least in cases with substantial disagreement.

Applying this to the case of group obligation, suppose that every member of some group agrees that the group has an obligation, and accordingly endorses collectively acting to discharge the obligation. (In the earlier presentation of the discursive dilemma, this is like endorsing the disjunction.) And further, every member of the group is individually rational, in that each endorses at least one specific, reasonable plan for discharging the obligation. (This is like endorsing one of the disjuncts.) Even in such a situation, there is no way, even in principle, to sum up the individuals' attitudes into a rational collection of group judgments. The impossibility results show that if we are to reach a rational group verdict – that is, if the group is to decide to discharge its obligation and to settle on a particular plan for doing so – it will need to do so by way of a group process. Pettit argues that the reasonable candidates for this process will involve doing something at a group level which looks analogous to reasoning at the individual level: for example, the individuals might come together, propose options, offer reasons, debate, and so forth, culminating in a series of votes.³¹ Whatever the precise form of the process, though, the important point is that until the group actually engages in it, there will be no fact of the matter as to what the group's plan is.

This, then, is the general explanation for why imperfect duties like beneficence, if they are group obligations, can seem so undefined or “fuzzy”. If a group with a duty of beneficence hasn't actually engaged in a group-level process in order to determine its plan or back-up plan, then there is no fact of the matter about what the plan is. A member of that group therefore isn't assigned any particular part of the obligation to discharge, and so enjoys latitude as to what, precisely, to do.

In order to see exactly how this works, and to make this discussion less abstract, let me return to the two scenarios we discussed earlier, which Murphy's and Cohen's proposals were unable to distinguish. Suppose that the duty of beneficence is held by society in general and that it demands both that the world's hungry be fed, and that drowning children be rescued. What is Megan, a member of society, to do about the unfed, and what is she

³¹ Pettit, 'Personal and Sub-Personal Reason'.

to do about the second child that Nick failed to rescue? Three possibilities are relevant.

First, suppose Megan is part of a society which has already settled on reasonable plans and back-up plans (or at least has procedures for doing so, when necessary). Call such a group *well-organized*. In a well-organized society, I take it that Megan ought to do whatever the plans ask of her. If some members of society fail to comply with the primary plan, then Megan ought to do whatever the relevant back-up plan requires of her. If some members of society fail to put the back-up plan into action, then she ought to do her part in the back-up back-up plan. In the case of the second drowning child, the back-up plan will presumably require Megan to save the second child. In the case of famine relief, Megan may be required to make a significant contribution, but in a group as well-organized as this one, presumably the machinery will be in place to minimize the burden and distribute it widely. For example, perhaps money would be collected through mandatory, progressive taxation and redistributed as necessary. This leaves open the possibility that Megan will have to make extreme sacrifices, but I take it that morality can sometimes demand extreme sacrifices.

The second case to consider is one in which Megan's society has decided on the relevant plans and back-up plans, but the plans are unreasonable. Call such a group *disorganized*. Now, a plan could be unreasonable in one of two ways: first, it might be *instrumentally* unreasonable, proposing a means insufficient to discharge the obligation. (Instructing Megan to save the second child by praying to the Water Spirits would be instrumentally unreasonable.) The more interesting case is of a *morally* unreasonable plan. The idea of a morally unreasonable plan will prove important later, so I'll take a few moments to explore it here, before returning to Megan's predicament.

On most theories of group decision, certain decisions or plans will count as substantively unfair or illegitimate, even if they are arrived at in the procedurally proper manner. For example, if a large majority were to vote to assign exclusive responsibility for funding the state's police force to a small, middle-class minority, that would (all else equal) be an unfair distribution of the burden. Even if majority vote is ordinarily an appropriate way to reach a group verdict, this decision is unfair in virtue of its content. It is not morally acceptable (all else equal) for a relatively wealthy group to assign some joint financial obligation to a relatively poorer one. Such a plan is morally unreasonable.

Notice that the problem of morally unreasonable plans arises only for groups. Since (I assume) any theory of group decision-making won't require

unanimity, that means that a group's plans will sometimes impose burdens on agents who never individually consented to them. (This, obviously, can't happen in the case of individual plans.) Most of us think that sometimes such impositions can be warranted. (Just because you disagree with the speed limit doesn't mean that you may drive as fast as you like.) But other times an imposition is not permissible, even if it is arrived at in the same manner – by a majority vote, say. The key question, then, is: what distinguishes the substantively permissible impositions from the impermissible ones? What makes one group plan morally reasonable and another morally unreasonable, if both are settled upon in the same manner?

We can say a few things here about the answer. As we saw above, plans that distribute a burden in a grossly unequal way will be substantively unfair, and therefore morally unreasonable. Plans that violate certain rights of a group's members will presumably count as morally unreasonable, absent the individual consent of those members. (It would be morally unreasonable to vote to pay the policing budget by selling some fraction of the population into slavery.) Beyond that, there will be many difficult cases. (Is a flat tax a morally reasonable way to distribute a financial burden? If not, how progressive must a tax be, to be morally reasonable?)

Rather than attempting to provide a full account here of what would make for a morally reasonable plan, I will pass the buck by noting that this problem is not specific to the case of beneficence and group obligations. It is a problem that needs to be addressed by moral and political philosophers, regardless of its role in a theory of group obligation. We need to know when it is okay for a group to impose burdens on non-consenting members, since that is what states do all the time. Fortunately, much work (for example, in the Rawlsian liberal tradition) has already been done on the problem of how to deal with reasonable disagreement. Accordingly, for the remainder of this paper I will assume that the group obligation account can take on board whatever theory of reasonableness proves to be the correct one.³² (The specific claims I make about morally reasonable and unreasonable plans are ones that I believe could be accepted from a wide range of different positions.)

³² In fact, the group obligation account needs to appeal only to a small portion of a theory of reasonableness. Whereas the Rawlsian needs to determine whether some particular end (e.g. universal health care or unemployment benefits) is an acceptable one to impose in the face of disagreement, when we deal with group obligations the only relevant questions concern the means. If a group has an obligation to save drowning children or avert a famine, that end is not up for debate. The only relevant question for our purposes is: *given that the group must avert the famine*, in what way should the accompanying burden be allocated?

Let's return, then, to Megan. Suppose she is a member of a group that has unreasonable plans for discharging its obligations. Faced with such plans, it seems to me that Megan ought to simply ignore them, and act as if the group didn't have a plan at all. Following through with an instrumentally unreasonable plan won't fulfill the obligation, and following through with a morally unreasonable plan will involve treating some group members in morally unacceptable ways. While these seem to me like good reasons to ignore a plan, that conclusion isn't obviously true. In at least some cases, it seems plausible to think that citizens have an obligation to abide by unjust or pointless laws. Nevertheless, I will set this worry aside and move on to the third possibility mentioned above – that of a society without plans to fulfill its obligations – since that case is the most interesting and (I'll argue) the one that gives group obligations the kind of latitude taken to be characteristic of imperfect duties.

The third case, then, is of a society with no plan or back-up plan. I'll call it *unorganized*. I take it that this typically describes our society, at least when it comes to back-up plans. Whether or not we have a well-specified primary plan for dealing with our collective obligations, we certainly don't have well-specified plans for dealing with non-compliance. In the drowning child scenario, the obvious primary plan is for Megan and Nick to each save one child.³³ Megan, though, now needs a back-up plan, since Nick has continued down the road. Even though there is no well-specified back-up plan, Megan knows the group is obliged to have one. In the present situation there is only one reasonable back-up plan: *Megan saves the second child*. Given the constraints of time (there isn't much until the child drowns) and resources (Megan is the only one in a position to help, since Nick has continued down the road), there aren't other options. So, whatever (instrumentally reasonable) back-up plan the group were to adopt, it would say that Megan is obliged to save the second child. Therefore, she ought to save the child.

Now, consider Megan's duty with respect to famine relief. It is unclear to me whether our society has any primary plan for feeding the hungry. Generously, we might suppose that we plan to feed the hungry through voluntary donations to charitable organizations. If that is our primary plan,

³³ There may be some sense in which this is the only reasonable primary plan, absent some other explicit arrangement. In that respect, we might call it a *default plan*. I suspect that in many situations there will be a salient default plan, which will relieve groups from the burden of having to explicitly reach judgments concerning those issues.

though, it has clearly failed, and so we are in need of a back-up plan.³⁴ Unlike in the drowning child case, however, here there are many reasonable back-up plans. One appropriate response to noncompliance might be to try to induce compliance through incentives. Another might be to try to coerce compliance through government taxation or seizure of resources. Another might be for the compliers to pitch in to pick up the slack. These plans (along with many others) could be the subject of reasonable disagreement. Some of the plans will involve placing further significant financial obligations on Megan, but some of them won't – some will place the additional burdens elsewhere. The group, therefore, has multiple reasonable back-up plans, among which it hasn't chosen. As we saw earlier, group decision requires some kind of process, and so until that process has been undertaken there is no fact of the matter as to the group's back-up plan. Thus, there is no fact of the matter as to what Megan is obligated to do. She is not obligated to give away every extra dollar she earns, since some reasonable back-up plans don't require it of her.³⁵

This argument shows, then, that we can use the notion of group obligation to ground a duty of beneficence that explains why Megan is required to save the second drowning child but not required to give most of her income to famine relief – that is, we can satisfy *LATITUDE*. And we've done this without sacrificing either of the other two criteria. Imperfect duties differ from perfect duties in that they are held primarily by the group. It is frequently possible for an individual to do more than an imperfect duty requires of her, since it is frequently possible to take on more than one's assigned share of the load. Thus, this proposal meets *IMPERFECTION*.

The proposal also meets *SPECIFICATION*, though it is a bit trickier to see why. Recall that *SPECIFICATION* says that the content of the duty should be specified precisely. Here, Megan's duty to the second drowning child is straightforwardly precise, but her share of the group's obligation to feed the hungry is indeterminate – which seems not to sit nicely with a requirement of precision. Indeed, I concede that, at the individual level, Megan's duty with regard to famine relief is not well-specified. However, it is important to remember that (in our example) the duty to feed the hungry does not lie in

³⁴ If this generosity is misplaced and we don't even have a primary plan, then the analysis that follows still holds, replacing 'primary' for 'back-up'.

³⁵ It is possible that any reasonable back-up plan will impose some additional burden on Megan, so it may be that she is obligated to increase her contribution. (I suspect, in fact, that any reasonable plan would require that most residents of affluent nations give far more than they actually do.) But it is surely not the case that *all* reasonable plans would require Megan to give away most of her income, so she is not obliged to do so.

the first instance with Megan. The group is the primary bearer of the duty, and the group's obligation is precise: feed all victims of the famine. As we noted earlier, moral obligations invoke a species of necessity, and that doesn't seem to permit them to be "fuzzy". The account I've outlined here meets this constraint, since it addresses the duty to the group and precisely specifies what must be done. Though this may leave other details undetermined (e.g. how the group should accomplish the task), that doesn't make the requirement itself imprecise.

In fact, the situation is no different than with a familiar, individual duty, such as to repay a loan. Both duties specify a precise end, and say that it must be brought about. Both duties also leave to the agent a choice in selecting the means and manner. (The individual can pay cash or write a check, and can drive to the bank via Main Street or via State Street. The group can send cash or food, and can divide the obligation among its members in any number of ways.) In this way, we can see that although from Megan's perspective the duty to feed the hungry looks indeterminate, the duty itself is really no different in kind from other, familiar duties – aside from the fact that it is incumbent on a plural agent, rather than an individual. Thus, the group obligation account is able to meet SPECIFICATION.

Before moving on to a pair of concluding sections, I want to highlight one additional virtue of the proposal I've offered here. On Murphy's and Cohen's proposals – and on many of the theories of imperfect duties we canvassed in §III – once you've done your individual share, you are morally in the clear. If others don't do their part, the moral black mark is entirely on them; you are justified in having a completely clean conscience. On the view I've proposed, however, that isn't the case. Complying agents aren't let off the hook completely, if others fail to comply. Even if I've done my part in the primary plan, or done the minimum required of me by the various reasonable back-up plans that the group could adopt, if others don't do their part *we* will still be in moral default. So, even if *I* have been as beneficent as I'm obliged to be, and accordingly can have a clean conscience in one respect, *we* will have failed to fulfill our obligation, and so I will be a party to a moral failure.³⁶ What an individual should do in such cases of collective default is the subject for another essay. But, briefly, it seems plausible to say

³⁶ Note that this result still holds even if some other agent is obviously responsible for the failure. If, for example, the government squanders tax revenue which it was supposed to use to feed the hungry, *we* are still in moral default. Just as Alice's failure to paint doesn't let the rest of the group off the hook, the government's failure to use its tax revenue properly doesn't leave us fully in the moral clear.

that if a group is failing to fulfill some duty because it lacks the appropriate plans or back-up plans, its members have an obligation to take steps (voting, lobbying, protesting, etc.) to see to it that the group becomes organized so that it can more effectively discharge its obligations. So long as a group of which you're a member is failing to fulfill its obligations, you have a duty to help to reconstitute that group, to ensure it does fulfill its obligations.

IX. Do Groups Have a Duty of Beneficence?

In this paper I've argued that thinking of beneficence as a duty held by groups has several attractive features, so long as we think of the duty in a non-reductive way. One thing I've conspicuously *not* argued for, however, is the claim that any groups actually have a duty of beneficence. Let me now make a few remarks about what form that argument would need to take.

There are two substantially different ways to argue for a duty to alleviate poverty. On the first, the mere fact that some are in need establishes an obligation on other agents to help. This kind of justification likely wouldn't be amenable to my proposal. If *all* agents have a duty to help the poor, then it turns out that individuals, as well as groups, have that duty. The motivation for pursuing an account of beneficence as an imperfect duty was to explain why individuals don't bear the heaviest or most demanding burdens. On this way of grounding an obligation to alleviate poverty, though, even if you don't have a duty to give all your income to charity in virtue of your *group* obligations, you do have a duty to give all your income to charity in virtue of your *individual* obligations. We'd be adding a less-demanding obligation on top of a very demanding one, not in place of it.

On the other hand, we might try to ground a duty of beneficence in the fact that certain agents bear responsibility for the hardships of others, either because they caused or failed to prevent those hardships or because they agreed to accept that responsibility.³⁷ This justification for a duty of beneficence is much more friendly to the group obligation account. Citizens

³⁷ See e.g. Thomas Pogge, *World Poverty and Human Rights* (Cambridge: Polity Press, 2002) for an argument of this type. If these sorts of considerations are the ones underlying our duties to the poor, then (as Pogge stresses) the duty is no longer a positive duty to aid, but rather a negative duty not to harm (or a duty to fulfill contractual agreements, etc.). 'Beneficence' therefore may no longer seem like the proper label. I agree that this is a somewhat revisionary suggestion. That said, this analysis would apply to most of the cases that are central to philosophical discussions of beneficence, and arguments like Pogge's are

of affluent countries contribute in all sorts of ways to poverty, hunger, and ill health throughout the world, through the purchase of sweatshop-produced goods, unfair or exploitative treaties, restrictive drug licensing agreements, and so forth. It's plausibly true, though, that my marginal contribution to these problems is insignificant, and therefore that I don't bear responsibility for rectifying them. But it's undoubtedly true that I'm a member of many groups that do make significant marginal contributions, and so those groups do have a duty to help. Similarly, while I never agreed to provide health care for Americans unable to afford it, I'm a member of a group that collectively has.³⁸ More needs to be said to flesh out these arguments, of course, but I think the general strategy has promise, and it clearly explains why it is that groups are usually – but not always! – the ones with these sorts of far-reaching obligations.

X. Imperfect Duties as Group Obligations?

At the beginning of this paper I said I was hoping to use a study of beneficence as a way of getting at the problem of imperfect duties. Since beneficence is usually taken to be the quintessential imperfect duty, my hope was that an analysis of beneficence might help us to formulate a general theory of imperfect duties. I've argued that, if we think of it as a duty held primarily by groups, we can make sense of a duty that has the basic structure usually taken to characterize an imperfect duty of beneficence.

frequently treated as players in the beneficence debate – rivals to traditional accounts grounded in positive duties. (See Garrett Cullity, *The Moral Demands of Affluence* (Oxford: Oxford University Press, 2004), pp. 8-9, 196-9.) For these reasons, although my proposal may be revisionary it doesn't strike me as objectionably so.

³⁸ As this last case suggests, I think the group obligation approach to beneficence is a natural ally of the social contract tradition. There will certainly be some occasions in which we can indict groups for, for example, failing to prevent children from falling into ponds. (We should have set up a fence around the pond, or hired a lifeguard.) But other times, no such negligence will be present. In such cases, we should appeal to the idea of an implicit agreement. We, as a society, have agreed to help children drowning in ponds – and miners stuck in caves, hikers lost on mountains, etc. The same basic thoughts apply to international natural disasters. Frequently, we are at least partially responsible for the consequences of an earthquake or hurricane, e.g. when the international economic order has left a country's infrastructure especially vulnerable to flooding. But other times, the best story might be a contractual one. We have explicit compacts concerning maritime rescue (1974 International Convention for the Safety of Life at Sea (5.33.1, 5.7); 1982 U.N. Convention on the Law of the Sea (98.1-2)), and I don't think it is implausible to think that we have implicit agreements to help one another during other crises.

This, of course, is not to say that the proposal I've offered is a perfect fit. (As I note in a footnote, above, the duty may look more like a negative duty than a positive one, for example.) But I think – since we have yet to find a completely satisfactory theory of imperfect duties – the match is close enough to call the proposal I've offered a theory of beneficence as an imperfect duty. (If you doubt this, then I have at least offered a proposal which would play much of the role usually assigned to an imperfect duty of beneficence.)

Supposing, then, that we can make sense of an imperfect duty of beneficence in this way, can the proposal be expanded? Could we say that imperfect duties are in general duties held by groups? Could we say that the kind of “fuzziness” and latitude which has seemed to characterize imperfect duties is really just the indeterminacy which can result when a group fails to settle on proper plans for fulfilling its obligations? My tentative answer to these questions is a very qualified “yes”. There are certain other intuitively imperfect duties which fit the analysis here very well. For example, consider global climate change: it's plausibly true that my marginal contribution to climate change is negligible, but it's certainly true that I'm a member of many groups whose contributions are significant. The duty to reduce CO₂ emissions will therefore generally be held by large groups, and my individual obligations will be determined by how the group divides up its responsibility (e.g. through carbon allowances or taxation).

In the case of other duties which have been taken to be imperfect, a group level analysis may fit, although the case isn't nearly so clear cut. Take the duty to develop one's talents. While Kant's official position is that this is an imperfect duty owed to oneself, I believe it is more plausible to think of it as a duty owed to others, in virtue of the common pursuit of a goal. In fact, Kant himself (in a somewhat anomalous passage) expresses that thought:

For, quite apart from the need to maintain himself, which in itself cannot establish a duty [of self-perfection], a human being has a duty to himself to be a useful member of the world.³⁹

If there are certain things that we – a basketball team, choir, fire department, army platoon, or nation – ought to accomplish, then I may have a duty to develop the talents that will help us to achieve those ends.

³⁹ *Metaphysics of Morals*, Ak. 6:446. A similar argument can be read into Kant's discussion of the duty under the Formula of Universal Law (*Groundwork*, Ak. 4:423). See e.g. Robert Johnson, 'Kant's Moral Philosophy', *Stanford Encyclopedia of Philosophy* (2008), §5.

Finally, there are some classically imperfect duties, such as gratitude and mercy, that don't seem plausibly analyzable as duties held by groups. Whether or not this constitutes an objection to construing imperfect duties as group obligations will depend on whether or not we think these things really are *duties* – I suspect they aren't – and (if they are) on how significant the similarities are between those duties and duties like beneficence. Even if we think that there is a duty of mercy, say, that carries with it some kind of latitude, we might find that the latitude is of a different type, or has a different justification, than the latitude had by the duty of beneficence. That would show that the class of classically imperfect duties does not constitute a fully unified set. Given the remarkable lack of philosophical consensus on imperfect duties, though, perhaps that shouldn't come as a surprise.⁴⁰

⁴⁰ I thank Tom Beauchamp, Christine Korsgaard, Craig Nishimoto, Kristi Olson, Tim Scanlon, Lucas Stanczyk, two anonymous referees for this Journal, the members of Harvard's Workshop in Moral and Political Philosophy, and audiences at Carleton College, the College of William & Mary, the University of Nebraska, Claremont McKenna College, and Georgetown University for helpful comments on earlier drafts of this paper.