

The Feasibility of Collectives' Actions¹

Holly Lawford-Smith

[This is a preprint of an article whose final and definitive form will be published in the *Australasian Journal of Philosophy*, 2012; the *Australasian Journal of Philosophy* is available online at: <http://www.tandf.co.uk/journals/>].

Does 'ought' imply 'can' for collectives' obligations? In this paper I want to establish two things. The first, what a collective obligation means for members of the collective. The second, how collective ability can be ascertained. I argue that there are four general kinds of obligation, which devolve from collectives to members in different ways, and give an account of the distribution of obligation from collectives to members for each of these kinds. One implication of understanding collective obligation and ability in the proposed way is that a *group* can be blameworthy for failing to realize its obligation even when no *member* of the group is blameworthy for failing to realize her obligation.

Keywords: collective obligation, collective ability, feasibility, ought, can.

1. Obligation and inability

Moral philosophers standardly take inability to be sufficient for denying an alleged obligation. This is the basic idea behind the principle that 'ought implies can' [Collingridge 1977; Pigden 1990; Sapontzis 1991; Sinnott-Armstrong 1984; Ward Smith 1961; see also Vranas 2007]. But little has been said about how that principle is supposed to work for alleged *collective* obligations. Moral philosophers can be excused for their silence; after all, they deal in individuals and actions. While political philosophers also deal in individuals and actions, their standard subject matter is *collectives* and *outcomes*. So political philosophers must confront this question.

Does 'ought' imply 'can' for collectives? Or, to put it another way, should we assess whether the requirements of a theory are feasible for a collective agent in the same way we assess whether they are feasible for an individual agent? Answering that question requires answering two component questions. First, what does a collective obligation *mean* for members of the collective? Once we know that, we can ask

¹ I am grateful to audiences at the Australian National University for helpful discussion on this paper, and to Bob Goodin, Wolfgang Schwarz, Rekha Nath, Stephanie Collins, John Maier, Christian Barry, John Cusbert and two anonymous reviewers for detailed comments and criticism.

whether the collective has the ability to realize the obligation. Second, and relatedly, how are we to establish collective ability? Are collective abilities established independently of the abilities of their members, or are they determined by them? Before we can say that 'ought implies can' for collectives, we need to know what 'ought' and 'can' *mean* for collectives. There are various literatures in the neighbourhood of these two questions, such as those dealing with collective action, shared intentions and shared cooperative activity, and collective *responsibility*, but no treatment of exactly those questions. I deal with collective obligations in §2, and collective abilities in §3 (there is some overlap between the two, because if we think 'ought' does imply 'can' for collectives, then obligation *assumes* ability). In §4 I discuss one implication of understanding collective obligation and ability in the way proposed, and in §5 I conclude.

2. Obligations for collectives

Imagine that Kewa, Tom, Mark, and Jonno are the four members of a furniture removal company. The company has a business agreement stating that each member is an equal shareholder in its earnings, and an equal participant in its undertakings. That explicit agreement is cemented by the implicit agreement stemming from the men's long friendship. It happens that one day, in the process of shifting furniture out of an apartment, the men stumble in the stairwell, dropping the heavy piano² they are carrying and in the process hurting a small child. The child is trapped, and if not soon released, will die.³

Many, perhaps all, moral theories would require the company to lift the piano off the child. It is imperative that redress is provided for harms one has caused; it is imperative to provide aid wherever there is suffering and it wouldn't cost one too much to do so; there is more good in the world in which the piano is lifted; lifting the piano is what a good person would do.

² Raimo Tuomela and Kaarlo Miller introduce lifting a piano as a typical collective action problem in [Tuomela and Miller 1988].

³ Here the duty to lift the piano is most plausibly contribution-based (the company dropped the piano, directly causing harm, and they have correlative duties to redress that harm). The obligation is clear and strong, which allows a charitable exploration of how collective obligations distribute to members. It would be interesting, though, to think about whether the same conclusions would follow from a weaker case, for example where the duty to lift the piano is most plausibly assistance-based (e.g. the company who dropped the piano have fled the scene, and another furniture removal company who happen to be clearing out a neighbouring apartment then have the opportunity to lift the piano). My inclination is to think that the *way* the collective obligation distributes will not change, but the *strengths* of the distributed obligations will vary across contribution- and assistance-based cases, correlating to the respective strengths of assistance-based versus contribution-based duties (most people think the latter are much more stringent, see e.g. [Pogge 2004]). It is also worth noting that assistance-based obligations are often very different from contribution-based obligations, so the ability to generalize from the latter to the former may be limited.

Would those moral theories require the *company* to lift the piano, though? Why not say that Kewa, Tom, Mark and Jonno dropped the piano, or that Jonno stumbled and caused the others to drop the piano? Then the duties would fall to the members, not the collective. I assume here that it is *the company* that dropped the piano, because this paper is about how collective obligations devolve to their members. Some people prefer to resist collective obligations, but I suspect one explanation of that is they don't think collectives are the *kinds of things* that can realize obligations. But that is precisely the question at issue. A convincing story about when collectives have the ability to realize their obligations will help to weaken resistance to collective obligations. For the sake of argument, then, assume that the *company* dropped the piano and is therefore obliged to lift it. What does that collective obligation mean for the four men who make up the collective?⁴

One possibility is that it doesn't mean anything. We might just say that the group has an obligation, and the individuals don't. Frank Jackson and Robert Pargetter suggest at one point in their [1986] that we should deny distribution over conjunction for 'oughts' (see also discussion in [Lewis 1973: 79-80]). The cases they talk about involve more and less ideal circumstances, but we can apply the same discussion to groups. The structure of their idea is as follows. It ought to be that *A* & *B*. If 'ought' distributes over conjunction then it follows that it ought to be that *A*, and it ought to be that *B*. But now imagine, given that it *won't* be that *A*, it ought not to be that *B*. This creates a contradiction: it ought to be that *B*, and it ought not to be that *B*. Or, to flesh out that structure with an example, it ought to be that the procrastinating professor accept the request to review a book, and actually review the book. If 'ought' distributes over conjunction, then it ought to be that the professor accept the request. But given that he won't review the book, it ought *not* be that he accepts. The world in which he accepts the request but doesn't fulfill it is worse than the world in which, knowing he won't fulfill it, he refuses the request (then someone else can review it). This creates a contradiction: he ought to accept, and he ought not to accept.

⁴ Before we can figure out *how* collective obligations devolve to members of collectives, we must establish *who* the members of collectives are. Often this is obvious: the members of the furniture-removal company are Kewa, Tom, Mark and Jonno. But who are the members of 'Australia'? Maybe Australian citizens; maybe members of parliament; maybe only the Prime Minister. We need a good model of the membership of some collectives before we devolve obligations to those members, otherwise we risk getting the story about the distribution of obligation wrong. Furthermore, we need a good theory of what counts as a collective. To count as a collective, certain organizational structures must be in place, such as a formal or informal decision-making procedure. The fact that we can draw an arbitrary line around some set of persons isn't enough to make them a collective, and there cannot be collective obligations where there are not collectives (which is just to make the simple point that non-existent agents cannot have obligations). There isn't space for me to develop a theory of collectives here, but as a starting point see e.g. [French 1979; Wendt 2004; Pettit 2007].

The solution Jackson and Pargetter suggest is a partitioning of option sets. Relative to the set of options including the actions 'accept and review' and 'accept and don't review', the Professor should accept and review; relative to the set of options including the actions 'accept and don't review' and 'refuse', the Professor should refuse. What he should do depends on what options we consider to be available to him [Jackson and Pargetter 1986: 254]. That looks like a neat solution, and at least in ordinary language it avoids the contradiction. It is not that *all things considered* the professor ought to accept and review, and also refuse. It's only that relative to one set of options he ought to accept and review, and relative to another he ought to refuse. (It's not clear that this is such a neat solution for the logicians, however, because it's not clear that deontic logic can handle an obligation being *relative* to a particular set of options).⁵

One reason to reject Jackson and Pargetter's solution is that the procrastinating professor does not have more than one option set available to him. He has just one, and what he ought to do depends on what is in it. The fact that he won't write the review does not suffice to limit his option set in a way that changes what he ought to do, i.e. it does not make another option set the right one to look at. What he ought to do is the best of what he can do. If the best action in his option set is accepting the request and writing the review, then that is what he should do.

But it's not clear whether we can say the same thing for all pairs of ideal and non-ideal obligation statements. Assume that it's bad for people to carry weapons. And now suppose the following is true: it ought to be that women are never assaulted, and do not carry defensive weapons. If 'ought' distributes over conjunction then it ought to be that women do not carry defensive weapons. But now suppose that the following is also true: given that women are assaulted, they ought to carry defensive weapons. That creates a contradiction: it ought to be that women carry defensive weapons, and it ought to be that they do not carry defensive weapons. When we think about the way the world ought ideally to be, in comparison with the way it ought to be *given* some of the ways it actually is, plenty of contradictions of that form will arise. I'm not convinced that it will always work to tell a story about relativized option sets. That is one reason to think that Jackson and Pargetter's first inclination, to deny distribution over conjunction for 'oughts', was correct.

Their discussion was about different things that a person allegedly ought to

⁵ To be more precise, examples like the procrastinating professor and the samaritan paradox show that deontic logic cannot be monadic. David Lewis, for example, has suggested moving to dyadic deontic logic instead of admitting that obligation cannot be modeled within standard modal logic [Lewis 2000: 5-19]. Anyhow, the issue of whether ought distributes over conjunction in ordinary language (whether 'ought *a* & *b*' means 'ought *a* & ought *b*') is a separate issue from how the logicians handle 'ought' formally.

do. But we can apply it to different parts of things that a collective ought to do. Let's try parsing the statement 'the company is obliged to lift the piano off the child' as 'Kewa and Tom and Mark and Jonno are obliged to lift the piano off the trapped child'. If 'ought' does distribute over conjunction, then it will be true that Kewa is obliged to lift the piano off the trapped child. But supposing neither Tom, nor Mark, nor Jonno will help him, and supposing he can't lift the piano alone, he can't be obliged to lift it. That is because he cannot be obliged to do what he cannot do.⁶ Distributing 'ought' over conjunction seems to get things wrong in that example. That is a reason to resist distributing. And this paves the way for collective obligations to exist when member obligations do not. It might be true that the company ought to lift the piano, and not true that Kewa ought to. That aside, if I can tell a good story about *how* collective obligation distributes from the collective to its members, then I don't *have to* accept that ought fails to distribute over conjunction as a way to distinguish collective obligation from individual obligation (or reconcile collective obligation with individual obligation).

One way the collective obligation obviously *doesn't* distribute is in the same form. The company is obliged to lift the piano off the child, but the piano is too heavy for Kewa to lift alone. So, being charitable, it can't be that the collective obligation distributes to each member such that he is obliged to *lift the piano off the child*.⁷ The distribution must be more sophisticated.

The more plausible and familiar suggestion is that members are obliged to *do a part*, or *take a share* in doing what the group is obliged to do. Assuming that it takes all four members of the furniture-removal company to lift the piano, the distribution might obligate each member to take a quarter of the weight of the piano. The members' obligations *add up* to the group obligation.

Of course, distribution into shares doesn't always mean distribution into *equal* shares. This is clear when we consider that one member of the company might be a lot bigger than another, or one a lot smaller. It is plausible that obligation is relative to capacity. If Kewa is a lot stronger than Jonno, then Kewa should probably shoulder more of the weight of the piano than Jonno. This is just as we think that

⁶ Obligations are complex and usually plural. Perhaps Kewa's obligation is a long disjunction of successively less ideal things he could do to help the trapped child. We might think of the first of these as the 'primary obligation' and the successors as 'secondary obligations' (which kick in when it becomes apparent that the primary obligation cannot be fulfilled), or we might think of the whole disjunction as a complex primary obligation – nothing important to this discussion seems to hang on which.

⁷ This might sometimes be the right story about distribution e.g. in overdetermination cases, where any member could fulfill the group's obligation by acting unilaterally. For example, imagine that the group of philosophers attending the Tuesday seminar is obliged to provide the speaker with a glass of water. Any one attendee might fulfill the group's obligation by getting the speaker a glass of water. But such cases are relatively rare.

the poor are not obliged to donate as much to charity as the rich. So let's say that for any distribution, the size of the share is relative to the capacity of the member. An obligation upon a collective (here a furniture removal company) translates into a capacity-relative obligation upon the members of the collective (here Kewa, Tom, Mark and Jonno) to do a *part* of what must be done.

Now imagine Kewa knows that the other members are squeamish about children suffering, and predicts they will flee the scene upon realizing a child has been trapped. What is Kewa's obligation if that transpires? Knowing that there is nothing he can do to shift the piano by himself, is he under any kind of obligation at all? (Remember we are concerned with his obligations *as a member of the group*, not the obligations that bear on him directly as an individual). Whatever his other obligations, he cannot have an obligation to lift a capacity-relative share of the piano.

It might be tempting to think that that is because he cannot take a capacity-relative share, in light of the fact that the others won't. That is to suggest that the obligation exists *prima facie*, but is dissolved as a matter of his inability. But this should be resisted. Rather we should say that Kewa did not have a categorical obligation to take a capacity-relative share in lifting the piano in the first place. His obligation, as a member of the company, was to take a capacity-relative share in lifting the piano *given a belief* that the other members would do the same.⁸ In the literature concerned with providing the necessary conditions for group action, theorists include a condition of common or mutual belief along the same lines. For example, Michael Bratman includes 'it is common knowledge between us that...' [Bratman 1992: 338], Philip Pettit and David Schweikard include 'each believe in common that...' [Pettit and Schweikard 2006: 21-24], and Raimo Tuomela includes that an agent 'believes that there is (or will be) a mutual belief among the participating members...' [Tuomela 1991: 263].

What other obligations might Kewa have, in virtue of his membership in the group? We can borrow from the legal notion of joint and several liability in supposing that the members of the company have not only a fair share of the responsibility for any action of the company's, but also a *full share* in the event that the others default. The analogy doesn't work perfectly for the piano example, because I've already said that one person can't lift the piano alone. But a 'full share'

⁸ Some might prefer to say that his obligation is conditional upon what the others *do* rather than what he *believes* they will do, and that his beliefs only provide an excuse. I resist this formulation here because I think obligation must remain practical: a person cannot have an obligation to do something he (reasonably) doesn't know about, e.g. an agent cannot have an obligation to rescue a child drowning outside in his swimming pool, if he has no way of knowing there is a child outside and anywhere near his swimming pool unsupervised.

here might be interpreted to mean 'taking responsibility for persuading the others to do their share' or for 'being the primary coordinator of the collective action'. The notion of 'joint' responsibility suggests that the story doesn't end at Kewa being obliged to do his part, conditional on believing the others will. The members of the company might have further duties to (try to) *bring it about that* the company does what it is obliged to do. Then Kewa will also be obliged to (try to) *get* the others to do their parts. Later in the paper, when I talk about a 'capacity-relative share', I mean to include both *doing* a part of the collective action, and persuading / convincing / motivating other members in the way necessary to *being able to do* one's part.

In a related discussion, Virginia Held argues that even random collections of individuals might be blameworthy for failing to *constitute* the kind of collective that would have the capacity to undertake some morally desirable collective action [Held 1970: 479].⁹ In that spirit, four strangers who happen to be passing in the stairwell where the piano is dropped (after the members of the furniture-removal company have fled the scene in panic) might be blameworthy if they do not attempt to *form* the kind of collective that could lift the piano. Presumably they each accept that the child ought to be rescued, and see that in order to be rescued, they must coordinate with each other in lifting the piano.

There are a few difficulties with this proposal. One is that 'collectivizing' is plausibly itself a collective action, because persons cannot 'collectivize' alone. Assuming the analysis given above,¹⁰ an agent will only have an obligation to collectivize if he has the reasonable belief that others will do so too. But that problem iterates, always depending upon an individual's beliefs about what others will do. An individual might have an obligation to bring it about that he can bring it about that he can bring about that... he can play his part in the collective action. The iteration stops only at the point at which the individual believes that others will do their shares.

The other difficulty is that non-existent agents cannot have obligations. So if the piano case involved a group of random strangers, the *reason* that each stranger ought to form a collective with the others couldn't be that *the (future) collective is obliged to lift the piano*. It can only be because it *would be good if* the piano were lifted, and one way to make realizing that good more likely is to form a collective with the relevant

⁹ I am grateful to an anonymous reviewer on this point. For an extended discussion on the duty individuals have to collectivize in order to achieve morally desirable outcomes, see [Collins manuscript].

¹⁰ Given the worry in the next paragraph, it's not clear that the obligation to collectivize is belief-dependent. The story I've been telling here is about how group obligations distribute to members. But non-existent groups cannot have obligations, so there's no story to tell about how they distribute (which would be back in time, as well as into parts) to would-be members. Thus the story about the obligation to collectivize needs a different treatment than I have given here.

others. But this is a paper about the distribution of collective obligation, so cases where there is no obligation and only a 'would be good if', however interesting, are outside its scope. Where there is a group, the members have certain obligations in virtue of their membership; where there is no group, the obligations members have bear on them *as individuals*, perhaps even because of their *capacity* or *potential* to be members of groups, but not *as* members of groups.

If collective obligations distribute to members as obligations to do a part of a collective action, given a belief about what the others will do, then the furniture company's obligation to lift the piano and free the trapped child distributes to Kewa, Tom, Mark and Jonno as an obligation upon each to do a part of lifting the piano, conditional upon a belief that the others will do a part too. This story seems to get things roughly right. Kewa doesn't believe that the others will do their parts, so he isn't obliged to lift a share of the piano. That is good, because his trying to lift it alone would be futile. Rather it is better that he fulfill a secondary obligation, perhaps looking for others to help him.¹¹

But let me try to be more precise about what a belief-dependent obligation must look like. All I have said so far is that Kewa is obliged to do his part only if he believes the others will do theirs. That is because for the particular task involved, all the members are needed for the collective action to be produced. But what kind of belief is sufficient? And what exactly is the logical structure of the obligation?

In response to the first question, we shouldn't allow just any old belief. What if Kewa stubbornly refuses to believe the others will do their shares, despite strong evidence to the contrary? Imagine that Tom, Mark, and Jonno immediately upon dropping the piano bend down to pick it up again, and they simply pause at that point waiting for Kewa to take his corner. If Kewa fails to do his share at this point, we should surely say it's because he fails to fulfill his obligation, not because his obligation was only to do his share *given* a belief that others would do their shares, and he didn't have that belief. Thus it seems that the belief must be *reasonable* (which is a placeholder for the idea that it should be sensitive to the available evidence).

In response to the second question, Kewa's obligation involves a conditional. He ought to: do his share *if* he has a reasonable belief that others will do their shares. 'If *a* then *b*' is logically equivalent to 'not-*a* or *b*'. So Kewa is obliged *either* to do his share, *or* not to have the reasonable belief that others will do their share. One way to fail to have this belief is to have an *unreasonable* belief that others will do their share. But having that belief seems like an undesirable way for him to be able to fulfill his

¹¹ Some might think that secondary obligations arise only when the original obligation comes from *contribution* rather than assistance. However, failing to fulfill the original obligation, no matter its source, surely does nothing but reveal further obligations – the question is only whether these are accurately called 'secondary'. (See also footnote 6)).

obligation. So the belief-dependence must be formulated in a way that doesn't allow such escape routes. One way to do that is to formulate the conditional negatively. It is obligatory for Kewa that, *unless* he reasonably believes that the others won't do their shares, he does his share. If he does not have the reasonable belief that the others will *defect* (not do their shares), then he must do his share. 'If not-*a* then *b*' is logically equivalent to '*a* or *b*'. If a person does not have the reasonable belief that others won't do their shares then he must do his share; he must *either* have the reasonable belief, *or* do his share.

So far we have only looked at an example in which *all* members of the collective must take some part of the collective action if the collective action is to be produced. The obligation upon the collective to produce that action translates into an obligation upon each member to do a share of the action, conditional upon not having a reasonable belief that the others would not also do their shares. But there are others kinds of cases.

For example, what if the company has *eight* members, and the piano can be lifted by four? What if more than four people would make lifting even easier? Or conversely, what if more than four people would make lifting much more difficult, by getting in each other's way? What if one person couldn't lift the piano alone, but could push it enough that the child would suffer a bit less? And what if for every member who pushed at it, the child would suffer less and less? Which of these kinds of situation we're looking at matters a lot for what story we tell about how collective obligation distributes to members. There are four basic types of situation, each needing a slightly different treatment.

The first, where every member of the company can make things a bit better for the child by pushing at the piano, I shall call *incremental good*. The more a member contributes, the better; and the more members that contribute, the better. These are the easiest cases to deal with, because the collective obligation distributes in a way that's categorical. Each member of the company should do a share of what the company is obliged to do, regardless of his beliefs about what everyone else will do, because the more that do a share (and the greater the share), the better.¹²

¹² I am assuming here that the collective good is fixed, and the individual contributions are incremental advances towards it. The situation is different when the collective good is itself an incremental good. For example, there is presumably a collective obligation upon Australia to lower its carbon emissions, and the more it can lower them the better. There it doesn't make sense to talk about an individual Australian's 'share'. It is not that there is some fixed outcome that is divided between the number of members in the group, so that when each does his share (or more) it is good, and when more people do their share, it is good. It can't be, because there's no fixed outcome to divide up. So when the collective good is itself incremental, it makes more sense to say that the distributed obligations are capacity-relative contributions to the collective *pursuing* the desired (or a desirable) outcome.

The second I call *joint necessity*. The piano case as described at the beginning of this section is an example: for the collective outcome to be produced, it is necessary that every member of the collective act. That is why any member's action is conditional upon the others' action: no member can realize the *collective* outcome alone.

There are two others. One I shall call *threshold good*, and the other *threshold good with harm*. When there are eight members of the group, the fact that it takes only four to lift the piano makes for threshold good. Any four members taking a share of the action is sufficient for meeting the threshold and producing the collective action, so there are a further four members who aren't strictly required to do anything.

When there are eight members of the group, and any more than four trying to lift the piano would be more of a hindrance than a help, we have a threshold good with harm case. The collective outcome is produced when the threshold is met *but not exceeded*. The only case that doesn't require the distributed obligation to be conditional upon belief is *incremental good*. The distribution there works as follows:

(1) *Incremental good*. When a collective has an obligation to ϕ , every individual member of the collective has an obligation to take a capacity-relative share in fulfilling (pursuing) the obligation.

Otherwise, we need to conditionalize upon beliefs, and in slightly different ways:

(2) *Joint necessity*. When a collective has an obligation to ϕ , every individual member of the collective has an obligation to take a capacity-relative share in fulfilling the obligation, unless she has the reasonable belief that at least one other member of the collective will not take a capacity-relative share in fulfilling the obligation.

(3) *Threshold good*. When a collective has an obligation to ϕ , every individual member of the collective has an obligation to take a capacity-relative share in fulfilling the obligation, unless she has the reasonable belief that sufficiently many other members of the collective will take a capacity-relative share such that the collective obligation will be fulfilled.

Sometimes cases of this third type will be such that all that matters is the threshold being met. But other times it might ease the burden on the others, or make things easier, if the same burden is shared between more members (this is not to be confused with when *it would be better if* more members contributed, which would be an example of *incremental good*). For example, it takes four people to lift the piano. But imagine that a fifth member of the original company, who happens to be passing by on his day off, sees the predicament that the company is in. He recognizes that helping them would ease the burden on each – even though it is not strictly necessary, and would produce the same outcome as when the four members lift the

piano alone. Nonetheless, we might still want to say that the fifth member ought to contribute. So we might want to add a caveat to (3): '...and it would not lessen the burden on those taking a capacity-relative share for her to contribute in addition').

(4) *Threshold good with harm*. When a collective has an obligation to ϕ , every individual member of the collective has an obligation to take a capacity-relative share in fulfilling the obligation, unless she has a reasonable belief that sufficiently many other members of the collective will take a capacity-relative share in fulfilling the obligation so that her own contribution would be detrimental to the collective obligation being fulfilled.

Just to reiterate, (2) – (4) say that an individual member of a collective has an obligation to a conditional. It is obligatory that, unless she has the relevant belief, she contribute. This is different from saying that if she has the relevant belief, then she is obliged to contribute. The obligation ranges over the conditional, not just its consequent (it has wide-scope, not narrow scope). In (3), the idea is that the agent ought to contribute unless she's sure sufficiently many others will; in (4), the idea is that the agent contribute unless she thinks so many others will that her own contribution would be harmful. The reason the obligation has to distribute conditional upon beliefs is that there is no action that every member of the collective must actually perform.¹³ Some must act, but others must refrain from acting (e.g. to avoid harm). Members of collectives can *fulfill* the individual obligations that devolve to them by being sensitive to what others will do. (It is the fact that the obligation is wide-scope that allows the obligation to be *fulfilled*. If the obligation were narrow-scope, we would rather say that it is *dissolved*. See the discussion in §4 below).

A problem occurs when the members of a collective conspire to free themselves of their obligation to do a part in the collective action by *agreeing* that they will each not do their parts.¹⁴ In virtue of agreeing to this conspiracy, each member of the group comes to have very good evidence for the belief that the others will not contribute. Each member has a belief that respects the available evidence, so the belief counts as reasonable. According to (2) - (4) the members are then free of any obligation to do a part in producing the collective action. But what is really going wrong here? Surely members are blameworthy for conspiring to escape their obligations. But insofar as they *had* conspired, each is surely not blameworthy for not doing a part in producing the collective action, which they would believe with a high degree of confidence to be futile. Conspiracy shows that (2) – (4) give the right

¹³ This is true unless forming a reasonable belief counts as an action, in which case we should distinguish between physical acts and mental acts. There's no physical act that every member must actually perform.

¹⁴ I am grateful to Bob Goodin for this suggestion. See his [Goodin forthcoming].

answer.

Uncertainty matters in distributing collective obligations. The cases discussed so far feature synchronic decision-making, where a member of a collective must decide on the basis of the evidence available to her whether the others will do their shares, and thus whether she will do hers. But decisions are not always like this. With diachronic decision-making the agent sees that others are contributing, and decides in light of that fact to contribute also. Sometimes there is communication and coordination, so that members can decide together how they will act, and who will take which burden (often members will have to perform quite different actions from one another, e.g. the manifold different tasks involved in a large engineering company constructing a new nickel mine). In situations where a lot of information is available about others' intentions, it will be easier to produce the collective action; in situations with much less information, it will be more difficult. But this paper is not about when collective action is easier or harder to produce. It is about what a collective obligation means for members of the collective (and what collective ability means in general).

In summary, notice what has happened to the obligation in the original case. The furniture removal company was obliged to lift the piano off the trapped child. That *means* each member is obliged to take a capacity-relative share in fulfilling the collective obligation (and remember that this includes an obligation to try to bring it about that the collective acts, which might involve persuading or convincing the other members to do their shares), unless he has a reasonable belief that others will not do their shares.¹⁵ Members' obligations to act are conditional upon their beliefs in all types of cases, with the exception of *incremental good*.

Notice that if any member has the relevant reasonable belief, and does not act *for that reason*, then she will have *fulfilled her obligation*. This is important. It is a bit odd to talk about an obligation being 'fulfilled' when the outcome in question doesn't actually result. But fulfilling an obligation just means doing what one is obliged to do.

¹⁵ It might seem that the obligation to persuade or motivate the other members to do their share should kick in when the member does not believe others will do their share, i.e. as a further duty, rather than part of the distributed collective obligation. There are two intuitions here that are closely related but should be distinguished. The first is that doing one's capacity-relative share does not exhaust the distribution of collective obligations. Members must take not only joint but *several* responsibility for the collective action being produced, which entails doing what they can to bring the collective action about (e.g. persuading others). The second is that if the member reasonably believes the others *will not* do their shares, there's still more she should do, like try to get others to take the members' place. The crucial distinction is that the former obligations distribute from the collective to the members, while the latter obligations bear upon individual members *as individuals*. Individuals finding themselves in situations of collective non-compliance may have duties to pursue the good which involve encouraging others to collectivize, but these are not duties that bear on them *in virtue of their membership in the collective*.

Because the obligation has wide-scope, ranging over a conditional, members fulfill it by *either* having the relevant belief *or* doing a capacity-relative share. An individual can only fail to do what she is obliged to do by not having the relevant belief *and* not acting. The group, on the other hand, fulfills its obligation by actually *doing* what it is obliged to do, in the piano example lifting the piano off the child. In §4 I consider an interesting implication of this. But first, collective *ability*.

3. Ability for collective agents

How do we establish, for example, whether it's true or false that the German military had the ability to overthrow Hitler in a military coup? How do we ascertain group ability, and does group ability come apart from member ability? The German military was physically close to Hitler in a way that few other groups were (compare with the students, the blue-collar workers, and so on). It had plenty of weapons, and strategic training. If any group was to succeed in taking down Hitler, it would be the military. This claim seems plausible, but how do we show that it is true?

The proposal in this section is that group ability is determined with reference to group members' abilities. With obligation, the distribution was downwards. Groups have an obligation, which divides into parts in its bearing on the parts of the group, the members. But with abilities, the distribution is upwards. Members have abilities, and these are aggregated to determine a collective ability.

So if we want to figure out whether the German military had available to it an action that would likely have resulted in a successful military coup against Hitler, we look at whether the individual soldiers had available to them actions that could have aggregated to form a successful military coup against Hitler. Did they? A coup requires intense strategizing and planning. It needs leaders and supporters. But think about the conditions under which individual soldiers in the German military were operating. Loyalty to Hitler was extremely fierce. The penalty for treason was severe – in all likelihood death. There were spies and informants everywhere. This means that no soldier could have (without high risk of death) started planning and strategizing in the way required to initiate a successful coup. If you don't know who you can trust, and the chances are that you can't trust *many people*, the risks of trusting anyone are too high. Furthermore, soldiers swore individual oaths of allegiance to Hitler himself, which were regarded as extremely serious. This would have made it difficult for individuals to even *conceive of* conspiring against Hitler. Thus closer inspection of the claim that the German military could have overthrown Hitler reveals that it is probably false. The military had the ability to overthrow Hitler if the soldiers making up the military each had the ability to do their parts in overthrowing Hitler. But they didn't; the conditions prevented it.

This answer depends upon a certain understanding of what suffices for inability. It is one thing to say that an action *cannot* be done by an agent, it is another to say that it is one she *should not* do, perhaps because it would be foolhardy. We could say that a collective action is not ruled out if the parts of it are not ruled out for any member of the collective; and the parts of it are not ruled out for any member of the collective if the member has an action available to her that *could* produce her doing a part. That would mean only the soldiers' forcible prevention from doing their parts would suffice to genuine collective inability. But that seems much too strong. The soldiers are prevented for all practical purposes. If one tried to begin planning the coup, he would soon enough confide in an informant, and the price of that would be death. This is true for any soldier. Most people would say the soldiers didn't really have the option of planning a coup, even though it is true that there's something they could have – very recklessly – done.

If that slightly weaker understanding of the inability members is the right one, then the soldiers did not have the constituent parts of the military's ability to pull off a coup. If the members didn't have the abilities then the collective didn't have the ability, and so the military can't have had an obligation to overthrow Hitler.

However, if we *mistakenly* decide that the collective has the ability, and then we distribute the collective obligation to members of the collective, it shouldn't be surprising that individuals have the abilities to fulfill their distributed obligations, despite lacking the ability to do a constituent part of the collective action. The military is not able to overthrow Hitler. But if we mistakenly decide that it is, we will say that the soldiers have an obligation to do a part of what is required to overthrow Hitler, unless they believe the others won't (or sufficiently many others will; or their own contribution would be detrimental; or if overthrowing Hitler were an incremental good situation; depending). The soldiers have the ability to fulfill *that* obligation, because it requires them either to contribute (which we have already established they cannot do) *or* to have the right kind of belief. They have the ability to have the right kind of belief, so they are in a position to fulfill their obligations.

Let's think about a situation in which a collective *does* have the ability to fulfill its obligation. Presumably the furniture removal company has the ability to lift the piano off the trapped child. We establish that ability by identifying the required action, and asking whether the members of the group each have the ability to do the relevant parts of it. In the original example, this is straightforward, because the members *all* have to do something, and they all have to do the *same* thing. They each have the ability to take a capacity-relative share of the piano. If they do this, the piano will be lifted.

If the members were all extraordinarily weak, so that it was beyond them to

lift the piano (setting aside the problem of how they could have come to drop it), or if the company had only two members, and it would take four to lift it, then we should say the company lacks the ability to lift the piano.

Groups have abilities because members have abilities. But individuals don't have abilities because groups have them. The fact that Germany had the ability to beat England in the 2010 World Cup didn't give Bastian Schweinsteiger the ability to beat England. Rather the fact that Schweinsteiger and his teammates had their respective abilities is what *gave* Germany the ability to beat England. Nor do groups *automatically* have the abilities that members have. Keweenaw has the ability to do a triple somersault, but that's no reason to say the *company* has that ability.

4. Asymmetry in blameworthiness

Collective obligation, and collective ability, do not come apart from individual members' obligations, and individual members' abilities. If a collective has an obligation, its members have a (distributed) obligation. And if a collective has the ability to fulfill its obligation, its members have the ability to do their constituent parts. Collective blameworthiness however *does* come apart from individual members' blameworthiness. The furniture company has the *ability* to lift the piano, because the four members who make up the company each have the ability to do the relevant parts of lifting the piano. And the company is *obliged* to lift the piano, because it dropped it in the first place (and for the other reasons discussed). There is no inability of the kind sufficient to reject obligation. Thus if the company fails to lift the piano, *it is blameworthy*. It is obliged to lift the piano, and it has the ability to lift the piano, so if it doesn't, then it fails to do what it is obliged to do.

But that is not necessarily true for the members of the company. As we have seen, the collective obligation distributes to them in the form that they are obliged to take a capacity-relative share in lifting the piano unless they have a reasonable belief that the others will not do their shares. So if any member has good reason to believe that at least one of the others will fail to take a share, he will not be obliged to do his own share. The collective action can fail without any member being to blame. So long as the beliefs members have are reasonable, it is reasonable for them to act upon them. The fact that members can fulfill their obligations by reasonably believing that others won't do their share shows that the collective can be blameworthy for failure without the members being blameworthy for failure.

Can this be right? Surely if an obligation distributes, when the distributed parts of it are done the whole thing will be done. To give a simple example, if I have a cake, and cut it into eight pieces and give them to eight different people to eat, there shouldn't be any cake left. How can it be that each member of a collective

satisfies her distributed obligation, and yet the group does not satisfy its? Strange as this may seem, it looks to be the right answer. Making the distribution such that satisfying it would *guarantee* the collective action would require making the shares categorical. But that would have the bizarre result that even when a member of the group *knows* his contribution will be futile, and maybe even counterproductive given that he could be doing something else, he'll still be obliged to do it.

For example, Kewa would be obliged to stay behind, pushing at the piano, even after the others had fled the scene. But surely he should look for others to help him, rather than do something he knows to be futile. Doing his share in the others' absence will not result in the collective action being produced. We only have the idea that the distributed obligations should be categorical because we want the collective action to be produced. But making them categorical doesn't do that. Only people always doing what they are obliged to do would do that. Making the obligations conditional upon beliefs at least avoids the outcome that one person does a share that is futile or counterproductive.

It is the fact that the obligation is wide-scope that allows the obligation to be *fulfilled*. If the obligation were narrow-scope, we would rather say that it is *dissolved*. An individual ought to do his share in a collective action unless he has a reasonable belief that others will not do their shares. On the wide-scope reading this is 'ought: act unless belief'; on the narrow-scope reading this is 'unless belief, ought act'. I have suggested that members' obligations have wide-scope. The *only way* they can fail to do what they ought is by not having the relevant reasonable belief, and not doing their own share. There are three other logical possibilities that allow a member to fulfill her obligation, namely having the belief and not doing her share, having the belief and nonetheless doing her share, and not having the belief and nonetheless doing her share. If the obligation had narrow-scope instead, only when the agent did not have the relevant belief would she have the obligation to do her share. The obligation would kick-in only when the antecedent of the conditional were true. And wherever the antecedent were false, the obligation would be *dissolved* rather than *fulfilled*. Still, collective blameworthiness and members' blameworthiness would come apart. Collectives would be blameworthy as before, wherever they ought and could but did not. But members would not be blameworthy, because the relevant antecedent being false would mean they *had no* obligation, and therefore couldn't be blameworthy for failing to fulfill it.

The upshot is that groups can be blameworthy when their members are not. (Of course, that doesn't mean they will *often* be such. No-fault failure to have the necessary beliefs doesn't happen all the time; more commonly if a group is blameworthy at least one of its members will be blameworthy too, as when Jonno

drops the piano and then flees).

5. Conclusion

Assessing whether theories' recommendations are feasible for collectives can be done in the same way as assessing whether theories' recommendations are feasible for individuals. Identifying the applicable case (joint necessity; incremental good; threshold good; threshold good with harm) allows us to say what obligations members have in virtue of their membership. Aggregating members' abilities to constitute group ability allows us to assess whether alleged group obligations meet 'can' constraints. An interesting upshot of the proposal defended here is that collectives can be blameworthy for failing to produce the obligatory collective action even when the members of the collective are not.

Charles Sturt University

References

- Bratman, Michael 1992. Shared Cooperative Activity, *The Philosophical Review* 101/2: 327-341.
- Collingridge, David 1977. 'Ought-implies-can' and Hume's rule, *Philosophy* 52/201: 348-351.
- Collins, Stephanie [manuscript]. Collectives' Duties and Collectivisation Duties.
- French, Peter 1979. The Corporation as a Moral Person, *American Philosophical Quarterly* 16/3: 207-215.
- Goodin, Robert [forthcoming]. Excused by the Unwillingness of Others? *Analysis*.
- Held, Virginia 1970. Can a Random Collection of Individuals be Morally Responsible? *The Journal of Philosophy* 67/14: 471-481.
- Jackson, Frank and Robert Pargetter 1986. Oughts, Options, and Actualism, *Philosophical Review* 95/2: 233-255.
- Lewis, David 1973. *Counterfactuals*, Harvard: Blackwell.
- Lewis, David 2000. *Papers in Ethics and Social Philosophy*, Cambridge: Cambridge University Press.
- Pettit, Philip 2007. Responsibility Incorporated, *Ethics* 117/2: 141-201.
- Pettit, Philip and David Schweikard 2006. Joint Actions and Group Agents, *Philosophy of the Social Sciences* 36/1: 18-39.
- Pigden, Charles 1990. Ought-implies-can: Erasmus Luther and R. M. Hare, *Sophia* 29/1, 2-30.
- Pogge, Thomas 2004. Assisting the Global Poor, in *The Ethics of Assistance* ed. Deen Chatterjee, Cambridge: Cambridge University Press.
- Sapontzis, Steve 1991. 'Ought' does imply 'can', *The Southern Journal of Philosophy* 29/3: 383-393.
- Sinnott-Armstrong, Walter 1984. 'Ought' Conversationally Implies 'Can', *Philosophical Review* 93/2: 249-261.
- Tuomela, Raimo 1991. We Will Do It: An Analysis of Group Intentions, *Philosophy and Phenomenological Research* 51/2, 249-277.
- Tuomela, Raimo and Kaarlo Miller 1988. We-Intentions, *Philosophical Studies* 53/3: 367-389
- Vranas, Peter 2007. I Ought Therefore I Can, *Philosophical Studies* 136/2:167-216.
- Ward Smith, James 1961. Impossibility and Morals, *Mind* LXX/279:362-375.
- Wendt, Alexander 2004. The state as a person in international theory, *Review of International Studies* 30: 289-316.