

Beyond Individualism

Stephanie Collins

This is the final draft of a chapter whose definitive version will be published in:

Theron Pummer and Hilary Greaves (eds), *Effective Altruism: Philosophical Issues*. Oxford University Press.

Introduction

In this chapter, I argue that group contexts pose two problems for individuals who are engaged in the project of benefiting others as much as possible. The first problem is that collective agents—as well as individual ones—can engage in morally worthy projects and (more importantly) can bear moral duties. When a collective bears a moral duty, the collective's members have 'membership duties.' For citizens of powerful democratic states, membership duties threaten to be numerous, varied, and weighty. They may leave little space for the project of benefiting others as much as possible. The second problem is that individuals who aim to benefit others as much as possible have reasons to signal their willingness to coordinate. These reasons imply that the actions an individual ought to take are likely not actions that ensure the marginal difference the individual makes is as high as possible. The combined result of these two problems is that the project of individually and marginally benefiting others as much as possible might reasonably take a back seat in an individual's day-to-day reasoning about what she morally ought to do.

Membership Duties

Some groups can—and can have duties to—engage in morally worthy projects. I will call such groups 'collectives.' Roughly, a collective is a group of individuals with a group-level moral decision-making procedure—a procedure that can take in reasons (including moral reasons),

deliberates upon them, and outputs decisions and instructions for enacting those decisions—that is ‘operationally distinct’ from the procedures held respectively by members. Its procedure is operationally distinct in that (i) the reasons it takes in differ from the reasons any of its members take in when deciding for themselves and (ii) its method for processing those reasons is different from the method of any one member when deciding for herself (Collins ms, ch. 6; French 1984; Rovane 1998; List and Pettit 2011). As an example, a collective might take the moral beliefs of *all* members and process these using a *majoritarian* method, thereby using a distinctively group-level set of inputs, and deliberative procedure, to arrive at the moral beliefs of the collective. Members are unlikely to use these inputs, processed in this way, when deliberating upon the moral beliefs they will hold themselves.

This can lead to distinctively collective-level moral agency—and, thereby, distinctively collective-level actions and duties—in the following way. When a collective operates its procedure over time, it acquires various beliefs (including moral beliefs), preferences, intentions, and so on. The collective then faces rational pressure to maintain coherence amongst these. Responding to this pressure, the collective might come to hold a moral belief that few—or even none—of its members holds (Rovane 1998; List and Pettit 2011). This belief might rationally compel the collective to decide to do something right or wrong—again, even if no individuals agree with, or control for, this collective decision (Pettit 2007). When a collective decides to act (including performing actions that are right or wrong), it will distribute roles to members jointly sufficient for enacting the decision. When members act within and because of their roles to successfully enact the decision, these actions by members are actions of the collective (French 1984). In this way, collectives can perform actions—including committing wrongs and producing goods—that are ontologically distinct from the actions of the aggregate of their members (for discussion of this ontological distinctness, see List and Spiekermann 2013).

To be clear, the idea here is not that we must posit collective actions and duties because sometimes individuals ‘make no difference’ or ‘make an imperceptible difference’ to a group-level outcome. For the vast majority of collectives’ actions, members will make some small difference to what the collective does—even if just by ensuring that the collective’s action occurs in one specific way rather than another. To illustrate, when I give a lecture, this partly constitutes the university’s action of providing lectures. But my individual action *does* make a difference to the exact way in which my university provides lectures—I can ensure the university provides (at least some) lectures that are interactive, or not, for example. So individual members can make perceptible differences to collectives’ actions, and yet those actions can be distinctively collective: my university’s action of providing lectures was the result of an intention produced by a decision-making procedure that is distinct from the procedure of any of its members. This is what makes it a collective’s action. Conversely, there are cases in which individuals (seem to) make *no* difference to an outcome in which they are (broadly speaking) involved, yet there is no group-level decision-making procedure, and so no collective agent, and so no collective actions or duties (such cases are described by, e.g., Parfit 1984, ch. 3; Kagan 2011). What I will say about membership duties has no bearing on the latter cases (though what I will say in the next section—on coordination—will).

That is: a group is a collective agent if and only if it has a group-level decision-making procedure. Sometimes members make a difference to the collective’s actions; sometimes they do not; in either case, the existence of a group-level decision-making procedure is what determines that there exists a collective agent with its own actions and duties. The collective is not independent or free-floating; it is physically constituted by members (Richie 2013; Hess forthcoming). But *if* there is a group-level decision-making procedure that has distributed roles to members sufficient for enacting the group’s decision, then the control over that outcome is wielded by the collective itself (Strand 2012).

Much more could be said about these issues. For now, I will assume that collectives—and their projects and actions, and their duties to engage in projects and perform actions—are entities in our moral universe. And I will assume that collectives can have the full complement of moral duties: duties to keep promises, to avoid causing harm, to return benefits received from injustice, to rectify past harms, to help those particular others whose needs most depend on their actions, and so on. (On collectives' abilities to bear this range of duties, see Collins and Lawford-Smith 2016.) Most of these duties have the following features: they entail claim-rights, they are dischargeable (rather than 'imperfect'), they are enforceable, and they are derived from 'the right' as constraints on action, rather than derived from 'the good' as the goal of action. On many ethical theories, these features give these duties priority over the project of benefitting others-at-large as much as possible. I will not assume that that priority is absolute—but I will assume that it exists.

Finally, I will assume that many powerful groups in our world—states, multinational corporations, intergovernmental organisations, and so on—are collectives that bear this range of duties. When we consider such powerful collectives, it is clear that their duties are weighty. After all, their promises tend to be backed by more formal commitments; they have a dangerously high propensity to cause harm; any benefits they have received from injustice tend to be large; their ability to assist those in need runs deep; and the harms they have caused tend to be devastating. And they interact with more persons than most individuals do, simply due to the scale of their operations. So, powerful collectives' duties tend to be weightier (in virtue of more fully satisfying various scalar duty-generating criteria), and more numerous (in virtue of interacting with more persons), than the analogous duties held by individuals.

This is stark when considering states. Because states' power is so pervasive, they acquire numerous duties. These duties have various grounds: the extreme harms states have caused via their imposition of oppression, colonialism, and present-day trade policies; the huge

benefits they have received via the same impositions; the weighty promises they have made via treaties; and (in the case of affluent states) their strong ability to assist those whose need depends on their actions. These features of states are the grounds of duties—held by the state—that are owed to whichever specific entities (individual or collective) the state has harmed, or has benefitted from the plight of, or has made promises to, or has the ability to assist, or so on. Additionally, states have final authority over a territory. This produces an unusual kind of duty: a duty owed to occupants of that territory, to use democratic procedures to create, enforce, and apply just laws within that territory, where these laws bestow benefits (most abstractly, rights) upon the territory’s occupants. This duty is grounded in the state’s being in a unique position to confer these benefits upon those within its territory, in virtue of it having final de facto authority over that territory.

What do all these duties—held by states—imply for states’ members? A lot. Although collectives are ontologically distinct from their members, this is in roughly way a solar system is ontologically distinct from the planets that constitute it.¹ There is not the collective, over there, with its duties, and the members, over here, with their duties.² Some facts about the whole have implications for the constituent parts, even though the whole has properties that

¹ Hess uses this analogy (Hess forthcoming).

² Such stark separateness is implied by Jeff McMahan when he asserts that “I am neither a community nor a state. I can determine only what I do, not what my community or state will do.” (McMahan 2016, 97) His first sentence is true; his second is false: I’ll go on to suggest that McMahan’s actions within-and-because-of his role in a collective determine (via constituting) *aspects* of actions of that collective, even if the collective acts differently at the overall level.

the parts do not. In particular, I have argued elsewhere that if a collective has a duty to see to it that X, then:

(1) Each member has a duty to use their role, as appropriate, to put inputs into the group's decision-making procedure with a view to it being the case that (the procedure distributes roles to members in a way that): if enough members used their roles with a view to seeing to it that X,³ then that would be sufficient for X in a high proportion of likely futures. These are 'X-sufficient' roles.

(2) Once X-sufficient roles are distributed, then each member has a duty to use their role, as appropriate, with a view to seeing to it that X. (Collins 2017, 40-43)

When a member 'uses their role with a view to seeing to it that X,' they act *within* (i.e., consistently with) and *because of* their role so as to increase the likelihood of X. Importantly, (1) and (2) might mean using the more fundamental aspects of one's role (e.g., one's right to free speech) to challenge other, less fundamental, aspects of one's role (e.g., that one pays taxes that serve unjust purposes). That is: a member's actions in discharging her membership duties are active and contestatory, not passive and obedient. Of course, in some collectives, one's role might not include any scope for pushing the collective towards distributing other roles. In such a collective, the most one can do is to use one's role—that is, act consistently with and because of one's role—with a view to seeing to it that X. If a member's role positively excludes actions (role-distributing or otherwise) with a view to X, then the 'as appropriate' clause is not met, and the member does her membership duty simply by checking that there is nothing she can do towards X within her role. (She might then have a duty to act upon the collective 'from the outside'—but this would not be a *membership* duty, i.e., a duty held qua member that contributes towards the collective's fulfilment of *its* duty.)

³ I return below to cases in which members know that X will not in fact happen.

Call (1) and (2) ‘membership duties.’ In a state, X might be ‘creating, enforcing, and applying just laws,’ ‘apologising for wrongs done to the aboriginal population,’ ‘negotiating with other states to close international tax loopholes,’ ‘increasing and redirecting the foreign aid budget,’ and so on. If an individual is a member of a state,⁴ and if that state has such a duty, then the individual has a membership duty.

What does this imply for the project of helping others as much as possible? It implies that our duties as members of powerful collectives (most conspicuously, states) threaten to consume the time, energy, and money that we might have spent on that project.⁵ States’ duties—particularly their duties grounded in past wrongdoing—are weighty, varied, and numerous. So—I will now suggest—citizens’ membership duties are likewise weighty, varied, and numerous. This is so, I will suggest, *even if* a citizen’s doing her membership duty will make a tiny difference, in expectation, to whether her state ends up ‘seeing to it that X’ in some realisation or other. (Though, to reiterate, my point is not that members make *no* difference, in expectation, to whether—or certainly to in what realisation—collectives do their duties. It is just that such difference-making is not a necessary condition on membership duties.)

⁴ I assume liberal-democratic citizens are members of states in the relevant sense. On this, see Collins and Lawford-Smith (ms).

⁵ Using a different framework, Ashford (forthcoming) roughly argues that we should *both* act as members of powerful collectives *and* use some resources to benefit others as much as we marginally, individually, and directly can with those resources. As stated above, I’m assuming that directed, dischargeable, enforceable, constraint-based duties—including where those duties are held by collectives—are more important than projects directed at benefitting others-at-large.

To see the weight of membership duties, consider a member of a democratic state. She does her membership duty by voting for a morally good election candidate—and using her right to free speech to encourage her compatriots to do the same—despite the overwhelming evidence that the candidate’s evil rival will win. Her collective’s duty is ‘seeing to it that a morally good candidate is elected’ and she has acted within and because of her role with a view to that, by way of voting and campaigning. When we consider merely the expected difference her action will make to whether the collective does its duty in some realisation or other, the value of her action looks low indeed (Gelman, Silver, and Edlin 2012).

But her action should not be assessed solely on the basis of its expected difference-making. The *content* of her duty is to act so as to make X more likely, but that duty’s *strength* is not determined simply by the expected difference her action will make to X. Instead, the strength of her duty is also partly determined (and increased) by the fact that it is a pro tanto duty-fulfilling feature *of the state* to have constituent parts that do their membership duties, regardless of the likelihood that those membership duties will lead to the state doing its overall duty.

This might sound strange, so let me draw an analogy between a state and an individual moral agent. I assume an individual’s attitudes towards others can have intrinsic moral value. Yet sometimes an individual has the right attitudes, without performing the right action that accords with those attitudes. The classic example of such an individual is Huckleberry Finn in Mark Twain’s novel *The Adventures of Huckleberry Finn* (Twain 1884; for philosophical discussion, Bennett 1974; Arpaly 2002). Huck respects his friend Jim, who is a slave. Huck has real human sympathy for Jim. Yet, at various points in the novel, Huck acts contrary to this respect and sympathy. This is because he also views Jim as an object and the rightful property of his former owner. Whatever else we say about Huck, the following is true: Huck is morally right in a way (he has respect and sympathy for Jim) *and* Huck is morally wrong in a way

(Huck views Jim as an object). Overall, the ‘right aspects’ get put aside, and ‘wrong aspects’ lead Huck to do wrong (at least at several points in the novel). His right attitudes are overridden (and sometimes disregarded) in his deliberations. In certain of his acts, his rightful attitudes are not reflected at all. But even as he does moral wrong, we want to praise those overridden or disregarded aspects of him that are morally right—that is, we want to praise his attitudes of respect and sympathy, and judge them to have value, *even though* they have made no difference to what Huck does.

My suggestion is that we view the *members* of a collective analogously to the *aspects* of an individual.⁶ That is: a member who does her membership duty performs an action that constitutes a ‘right aspect’ of her collective, regardless of whether her doing her membership duty is disregarded at the level of her collective’s deliberations, and even if it is not reflected in what her collective does overall. When a member acts within and because of her role with a view to her collective doing its duty, this is a right aspect of the collective. This is just as it is a right aspect *of Huck* that he has respectful and sympathetic attitudes to Jim, regardless of what he does overall. Of course, that aspect would have *even more* value if it were (or had higher expectation of being) efficacious at the level of what Huck does overall. But such causal efficacy does not exhaust the value of Huck’s attitude: that attitude is still of value even once it has become fully determined that it will not affect Huck’s actions. I hope this is obvious in the case of Huck. If I am right that collective agents and individual agents should be viewed analogously in these matters, then it is also true of a collective such as a state. A member doing their membership duty is like Huck having respectful and sympathetic attitudes: morally right, even if inefficacious.

⁶ I expand on this, and on the Huck Finn analogy, in Collins forthcoming.

How, then, is a membership duty's importance to be assessed? The importance of a given membership duty is a function of two things: first, its expected efficacy at inducing the collective to discharge its overall duty; second, the importance of that overall collective duty (which will, in turn, be partially a function of what type of duty it is: to repair harm done, return benefits received, keep promises, etc). The second of these reflects the value of an entity's having aspects (in the collective case, members) that accord with its duty, even if that entity will not do its duty overall. The bigger and more politically significant a collective is, the more weighty, stringent, and demanding its duties will be. So, the more weighty, stringent, and demanding the relevant membership duties will be (holding expected efficacy constant).⁷

My claim is not that membership duties *always* override an individual's non-membership duties. If much good—or much right—can be done by ignoring one's membership duties, then that might be what an individual all-things-considered ought to do (Berkey 2018). But if the collective is socially and politically significant—if it is, for example, a state or multinational corporation—then one's membership duty will be weighty enough to *compete* with the good one can do on one's own. By 'compete', I mean that the membership duty and the duty to do good must be balanced against one another in one's deliberations about what to do on a particular occasion, and that neither should be persistently favoured across many occasions. A more precise statement of the relative values would fail to do justice to the

⁷ This conception of what collectives' duties imply for members' duties is similar to Woodard's (2011) proposal, but my argument relies on collective agency (and its similarities to individual agency) in a way Woodard rejects. The role of collective agency in generating reasons for members to unilaterally do their parts is likewise not considered by Dietz (2016, 971-5), who therefore concludes (in my terminology) that membership duties exist only if the collective will do its overall duty if I do my membership duty.

distinct (perhaps incommensurable) nature of the values. Suffice to say, sometimes, you ought to take to the streets to voice your views qua citizen (even if this will not change your state’s behaviour), or take to the picket line to voice your views qua employee (even if your employer will not listen)—because these are the membership duties that are entailed by your state’s or employer’s duties that it owed to specifiable entities—rather than spending that time benefiting others-at-large as much as possible.

Coordination Duties

There is a second problem that group contexts pose for individuals who are engaged in the project of benefitting others as much as possible. This second problem is internal to that project: that is, whereas membership duties suggested we have reasons to do something other than pursue that project, the present problem persists if one assumes that the project should be pursued. The problem is best introduced via a simple example, depicted below. (The example derives from Rousseau 2004, 29-30; it is discussed extensively by Brian Skyrms (e.g. Skyrms 2001); a similar example is discussed by Dietz (forthcoming) in the context of effective altruism, but Dietz posits group reasons to solve the problem—as we shall see, my solution eschews group-level reasons in groups that are not agents).

	Hunt stags	Hunt rabbits
Hunt stags	5 , 5 (10)	0 , 3 (3)
Hunt rabbits	3 , 0 (3)	3 , 3 (6)

In the example, two individuals are out hunting. Each can choose to hunt either stags or rabbits. It will take two of them to kill a stag, but each can kill a few rabbits on their own.

So, if each chooses to hunt stags, then each will receive five units of wellbeing. Each individual who hunts rabbits will receive three units of wellbeing. If one of them hunts stags on their own, then that individual will receive no wellbeing. The total value produced by the two agents together appears in parentheses in the matrix. I will assume both aim to benefit others as much as possible, and both weigh others' wellbeing equally to their own. So this total value is what they each care about.

There are two Nash equilibria: two situations in which neither has a reason to change their strategy once they learn the other's strategy (assuming that the other will not change their strategy). These are, first, the situation in which both hunt stags and, second, that in which both hunt rabbits. So what you should do depends on what you think the other will do: if you think the other will hunt stags, you should hunt stags; if you think the other will hunt rabbits, you should hunt rabbits. This is not helpful if you have no beliefs about what the other person will do—especially if you know that the other person is rational and has no beliefs about your potential actions (in which case, they will get stuck with the same two conditionals). (Bacharach 2006, 44; Gold and Sugden 2007; Tuomela 2013, ch. 7).

Of course, we usually have beliefs about what other people will do. And if your credence that the other person will choose to hunt stag is higher than about 0.41, then your choosing to hunt stag will maximise expected goodness.⁸ And it might seem that your credence that the other person will choose to hunt stags will definitely be higher than 0.41, because, as Parfit puts it, “the unique best outcome is clearly salient. It is the obvious place to meet.” (1988, 13) So, Parfit thinks, people will naturally have a high enough credence that the other person will choose to hunt stags. So each will reason their way into hunting stags, via aiming to maximise expected value.

⁸ I thank Hilary Greaves for pressing this.

Unfortunately, it is sometimes reasonable for each individual to have a high credence that the other will hunt rabbits. In this case, expected value theory will tell each to hunt rabbits. And yet, there is an intuitive sense in which something has gone wrong in such a case: they should have coordinated on hunting stags. Lest this sound like something that would never happen, consider how Stag Hunt is analogous to the real-life project of helping others as much as possible. Roughly, we might replace ‘hunt stags’ with ‘pursue systemic change’ (as discussed by Gabriel and McElwee (this volume)) and replace ‘hunt rabbits’ with ‘pursue marginal change.’⁹ All those who are concerned to help others could act with a view to systemic changes in international institutions—pushing their states and intergovernmental organisations for reform of international processes, norms, and laws. Alternatively, each could donate some percentage of their income to charities that are seeking to make incremental improvements to the lives of the world’s worst-off. If all pursued incremental improvements, then this would do some good. But it is plausible that this would not do as much good as if all acted responsively to one another with a view to systemic change. Finally, there is the outcome where some act responsively to one another with a view to systemic change, and some contribute to incremental improvements. Here, the efforts of the systemic changers will be futile—and the efforts of the incremental improvers will do some good, but less than half the good that would have been done if all had engaged in systemic change. If an individual reasonably has a high credence

⁹ I am not the first to suggest an analogy between Stag Hunt and the project of helping others as much as possible. Benjamin Todd (2016) has worried that those within the effective altruism social movement might end up ‘hunting rabbits’ via piecemeal interventions, instead of pursuing “large scale changes.” But Todd finds this unlikely, since “[s]tag hunt situations arise rarely in real life, because if both groups communicate, then they’ll both go for stag.” The problem of communication is the one I highlight below.

that others will pursue incremental change, then that individual should also pursue incremental change. And yet, something has intuitively gone wrong—the world has fallen short of an attainable ideal—if all those who want to help others as much as possible pursue incremental change, as a result of them each having reasonably having a high credence that the others will do likewise.

So what should each do?¹⁰ A natural answer is that they should communicate with the others, so that everyone comes to believe that everyone else will pursue systemic change, so that everyone can reason their own way into doing likewise. This answer is not available in the usual set-up of Stag Hunt, where it is assumed the players cannot communicate. But it does seem available in the real-world case of individuals who are concerned to help others as much as possible. This is roughly the answer of Donald Regan’s “cooperative utilitarianism.” In brief, according to Regan, each individual should take the following five steps:

- (1) Be willing to take part in the joint attempt to produce the best possible consequences by coordinating with whoever is willing to coordinate, whoever they turn out to be (where it could be that you are the only one);
- (2) Determine who else is a ‘co-operator,’ i.e., such that they: (a) have taken step (1); and (b) correctly identify who else has taken step (1);

¹⁰ Dietz (2016) argues that in cases like this, *the group* should pursue systemic change (hunt stags) even while *the individuals* should pursue marginal change. Parfit (1988, 7-9) suggests likewise. I’m sceptical that non-collective groups can bear reasons. And the group I’m concerned with in this section—namely, that constituted by all people who are concerned to help others as much as possible—is not a collective. I also don’t think this solution captures what’s intuitively gone wrong. So I’m cutting straight to the question of what each hunter should do, where this doesn’t derive from what the group should do.

- (3) Ascertain what the behaviour will be of the non-co-operators (i.e., those who do not meet conditions (a) and (b));
- (4) Identify the best pattern of behaviour for yourself and the other co-operators, given the behaviour of the non-co-operators that was ascertained at step (3);
- (5) do your part in the pattern identified at (4).¹¹

This gives the following nice result: if you believe others will pursue incremental change no matter what, then you should not identify them as co-operators; if you are willing to pursue incremental change, then others should identify you as a co-operator. So, what each should do depends on what others have been identified as co-operators.

The pressing problem is how to identify co-operators—and how to get others to identify you as a co-operator. Regan does not solve this problem, saying it “simply does not matter, in theory. (In practice, of course, it may matter a great deal).” (1980, 152) But this does not mean philosophers can safely ignore the problem. After all, the line between theory and practice here is vague: it is unclear exactly in what sense the issue of communication is merely practical, even though the coordination problem as a whole is theoretical. One could say that the ‘theoretical solution’ is easy: the solution is ‘converge on hunting stags; everything else is merely practical.’ But presumably this solution would not satisfy Regan *qua* theoretical solution. So it is unclear why he rests content with not exploring the communicative issue.¹² And even if the problem of communication is purely practical, we should consider how to solve it.

¹¹ This paraphrases Regan (1980, 135-6); Regan’s full formulation is at (1980, 157-8), but the additional details don’t deal with the communicative issue I highlight below.

¹² Parfit agrees, describing Regan’s theory as a “partial failure” due to its failure to “wholly explain how the agents manage to cooperate.” (1988, 6)

Parfit also does not solve the communication problem. He gives three conditions that are jointly sufficient for an individual obligation to cooperate on the optimal outcome: “When (1) the members of some group would make the outcome better if enough of them acted in some way, and (2) they would make the outcome best if all of them acted in this way, and (3) *each of them both knows these facts and believes that enough of them will act in this way*, then (4) each of them ought to act in this way.” (Parfit 1984, 78-9, emphasis added) The question is how we can come to have the knowledge and beliefs in Parfit’s condition (3). When we find ourselves in situations like the Stag Hunt, we need to form the right beliefs about others—and, crucially, help others to form the right beliefs about us—before any of us can even satisfy the precondition (of having a sufficiently high credence that the others will cooperate) on having an obligation to act for the best overall outcome, let alone satisfy the obligation itself.

The problem is that those who are trying to help others as much as possible are a diverse bunch. The bunch includes trade unionists, activists, advocacy organisations, and those involved in local and national politics—as well as ‘incremental changers,’ that is, those who contribute to incremental change and who advocate that others should do the same. This raises the question of how, and whether, incremental changers are sending signals that encourage other potential co-operators to view incremental changers as co-operators or as non-co-operators; as stag hunters or as rabbit hunters. Numerous potential parties to systemic change are more likely to perceive incremental changers as *not* willing to take part in the joint attempt to help others as much as possible by coordinating with whoever is willing to coordinate. Simply put, if one flagrantly hunts rabbits, then one is probably not signalling willingness to hunt stags.¹³ This is particularly—but not only—so if one is pursuing marginal change through

¹³ Thus it’s inadequate to say, as Peter Singer does, that “effective altruists are already organizing together. Charities are themselves a form of coordination, enabling thousands of

such anti-systemic-change methods as earning-to-give while working in a job that perpetuates the very system that systemic change would overhaul. This is hunting rabbits in a way that scares away stags.¹⁴

To this, an incremental changer might object there is little consensus over which precise systemic change to pursue, how to pursue it, or how to signal that one is a co-operator. In the terms of the analogy: there is no agreement on where the stags are, how to hunt them, or how to let fellow hunters know that you would like to hunt them. Even if we concede this, it is no reason to give up on hunting stags. After all, the numbers in the matrix are net gains, so they already reflect the costs of various way of finding stags, planning the hunt, signalling willingness, and so on—where the costs of these diverse and numerous means are weighted by the probability that those means will be successful if the agent in question attempts to take them. The objector is effectively contending that ‘pursuing systemic change’ (where this includes the value of the pursuit, not just the value of the change once realised) would not be as valuable (relative to marginal change) as the analogy suggests. Such an objector would then need to turn to the arguments of Gabriel and McElwee (this volume), which I will refrain from reproducing.

donors to work together for a common goal, and beyond that, the effective altruism movement has several ‘meta-charities’ like The Life You Can Save, Giving What We Can and the Centre for Effective Altruism, which are doing their best to expand the movement or assess which charities are the most effective, and get more people involved in giving to effective charities.” (Singer 2016, 168) That is, “effectives altruists” (by which, from context, I suspect Singer means incremental changers) are engaging in co-ordinated mass rabbit hunting.

¹⁴ Earning-to-give in harmful jobs is advised against by, for example, the charity 80,000 Hours (<https://80000hours.org/articles/harmful-career/>)—but not for the reasons I give here.

A second objection to the Stag Hunt analogy is that—in the case of systemic change—hunting rabbits is a good way to hunt stags. For example, by funding anti-malarial bed nets, one improves people’s health. When people have good health, they are more able to co-operate for systemic change. Thus, incremental change boosts the prospects for systemic change. This is sometimes called a “flow-through effect” (Karnofsky 2013): the benefits “flow through” the direct recipients (bed net users) to indirect recipients (those whose rights are protected by the system that bed-net-users later campaign for). It is as if you are nursing a sick hunter by feeding them rabbits. This makes the sick hunter able to hunt stags in future. However, at the same time, the nurse is giving the sick hunter reason to identify the nurse as a rabbit hunter. This gives the sick hunter reason to be a rabbit-hunter themselves when the time comes, because (to reiterate) if everyone is hunting rabbits, then the best thing to do is to hunt rabbits.

How, then, does one signal one’s willingness to coordinate for systemic change? This is an empirical question, but the general answer is obvious: act directly upon the systems that need changing. This can be done via protests and demonstrations, membership and activism within political parties, involvement in trade unionism, and voting—alongside giving to those organisations that directly press for systemic change (advocacy organisations being prime examples).¹⁵ Ethical consumption—buying fair trade products, eating vegan, and so on—can also be understood as signalling one’s willingness to work with others to pursue systemic change (Lawford-Smith 2015). Signalling willingness will often mean acting for causes that others are already acting for, that is, causes that are not ‘neglected’ by others. It might also mean acting for causes that do not seem ‘tractable’ when considering just the marginal effect of one’s own individual contribution.

¹⁵ These suggestions are in the spirit of Srinivasan (2015) and Herzog (2016), though my argument for them differs from both.

Importantly, such signalling cannot be done by simply making a *one-off* public declaration. One cannot simply post on social media ‘I’ll push for system change if you all do!’, and then turn around to pursue marginal change until others post likewise. For one’s signal to be convincing, and for it be communicated to those beyond one’s immediate circle, it is better if the signal is persistent, consistent, and insistent. Of course, if one persistently, consistently, and insistent posts about one’s willingness on social media, then this may well do part of the trick. But even ongoing declarations of that kind will only get the signal across to a small audience—at least for those of us whose social media circle consists mostly of people we already know.

That said, in theory, one can signal one’s willingness to cooperate for systemic change while also pursuing (some types of) marginal change. One can be a conscientious citizen, consumer, union member, Amnesty International donor, *and* donor to charities that make incremental improvements in people’s lives. However, by focusing our immediate efforts on the project of marginally and individually benefitting others as much as possible, we run two risks: first, distracting ourselves from the importance of signalling for systemic change (since each of us only so much time, energy, and money); second, signalling those not engaged in this project that we are opposed to cooperation for systemic change—especially if we pursue incremental change in a vocal way.

Conclusion

The two problems I have explored combine to suggest that the project of ensuring that oneself—as a marginal individual—benefits others as much as possible might reasonably be crowded out of a person’s practical reasoning.

First, as members of states (and other collectives), we have duties to act within and because of our roles in the collective with a view to the collective responding to the duties that

apply to it. The duties that apply to collectives are not just duties to benefit others-at-large as much as possible, but are also duties not to do harm, to repay for harms done, to keep promises, and so on. Individuals' 'membership duties' gather weight not just from the member's expected effect on the collective's overall actions, but also from the fact that an individual's fulfilling their membership duty constitutes a right aspect of their collective. Second, an individual should not act so that they (as an individual) benefit others as much as possible, if this would distract them from—and send the wrong signals about—the greater change that a coordinated group could produce.

This suggests that political action is more important than might be implied by the simple idea of acting, as a marginal individual, to benefit others as much as possible. It suggests that we should attend to causes that are not neglected by others, and that might not be tractable by individual action. Perhaps we should engage in the project of (1) using evidence and reason to figure out to which others we have duties—where those duties might arise via our membership in collective agents—and (2) taking initial steps towards working together with others to ensure that we and others are responsive to those duties.

Acknowledgements

For helpful comments, I thank Brian Berkey, Hilary Greaves, Lisa Herzog, Brian McElwee, Avia Pasternak, Theron Pummer, Christian Wendelborn, and audiences at the 2016 meeting of the Association for Social and Political Philosophy (LSE) and a 2016 workshop on Effective Altruism (Oxford).

References

Arpaly, Nomy. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press, 2002.

- Ashford, Elizabeth. Forthcoming. "Severe Poverty as an Unjust Emergency."
- Bacharach, Michael. 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton University Press.
- Bennett, Jonathan. "The Conscience of Huckleberry Finn." *Philosophy* 49(188), 1974: 123–34.
- Berkey, Brian. 2018. "The Institutional Critique of Effective Altruism." *Utilitas* 30(2), 143–171.
- Collins, Stephanie. Forthcoming. "Who Does Wrong When an Organisation Does Wrong?"
- Collins, Stephanie. 2017. "Duties of Group Agents and Group Members." *Journal of Social Philosophy* 48(1), 38-57.
- Collins, Stephanie and Holly Lawford-Smith. 2016. "Collectives' Duties and Individuals' Duties: A Parity Argument." *Canadian Journal of Philosophy* 46(1), 38-58.
- Collins, Stephanie. Under contract. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.
- Collins, Stephanie and Holly Lawford-Smith. ms. "We the People."
- Dietz, Alexander. 2016. "What We Together Ought to Do." *Ethics* 126, 955-982.
- Dietz, Alexander. Forthcoming. "Effective Altruism and Collective Obligation." *Utilitas*.
- French, Peter. 1984. *Collective and Corporate Responsibility*. Columbia University Press.
- Gabriel, Iason and Brian McElwee. This volume.
- Gelman, Andrew, Nate Silver, and Aaron Edlin. 2012. "What is the Probability that Your Vote Will Make a Difference?" *Economic Inquiry* 50(2), 321-326.
- Gold, Natalie and Robert Sugden. 2007. "Collective Intentions and Team Agency." *Journal of Philosophy* 104(3), 109-137.

- Herzog, Lisa. 2016. "Can 'Effective Altruism' Really Change the World?" *OpenDemocracy.net*. <https://www.opendemocracy.net/transformation/lisa-herzog/can-effective-altruism-really-change-world#>
- Hess, Kendy. Forthcoming. "The Peculiar Unity of Corporate Agents."
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy and Public Affairs* 39(2), 105-141.
- Karnofsky, Holden. 2013. "Flow-through Effects." *The GiveWell Blog*. May 15. <https://blog.givewell.org/2013/05/15/flow-through-effects/>.
- Lawford-Smith, Holly. 2015. "Unethical Consumption and Obligations to Signal." *Ethics and International Affairs* 29(3), 315-330.
- List, Christian and Philip Pettit. 2011. *Group Agency*. Oxford University Press.
- List, Christian and Kai Spiekermann. 2013. 'Methodological Individualism and Holism in Political Science: A Reconciliation.' *American Political Science Review* 107(4), 629-643.
- McMahan, Jeff. 2016. "Philosophical Critiques of Effective Altruism." *The Philosopher's Magazine*, 2nd Quarter, 92-99.
- Parfit, Derek. 1984. *Reasons and Persons*. Clarendon Press.
- Parfit, Derek. 1988. "What We Together Do." ms.
- Pettit, Philip. 2007. "Responsibility Incorporated." *Ethics*
- Regan, Donald. 1980. *Utilitarianism and Cooperation*. Oxford University Press.
- Richie, Katherine. 2013. "What Are Groups?" *Philosophical Studies* 166, 257-272.
- Rousseau, Jean-Jacques. 2004. *Discourse on the Origin of Inequality*. Dover thrift edition. Ed. Greg Boroson. New York: Dover Publications.
- Rovane, Carol. 1998. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton University Press.

- Singer, Peter. 2016. "The Most Good You Can Do: A Response to the Commentaries." *Journal of Global Ethics* 12(2), 161-169.
- Skyrms, Brian. 2001. "The Stag Hunt." *Proceedings and Addresses of the American Philosophical Association* 75(2), 31-41.
- Srinivasan, Amia. 2015. "Stop the Robot Apocalypse." *London Review of Books* 37, 3-6
<http://www.lrb.co.uk/v37/n18/amia-srinivasan/stop-therobot-apocalypse>.
- Strand, Anders. 2012. "Group Agency, Responsibility, and Control." *Philosophy of the Social Sciences* 43(2),201-24.
- Todd, Benjamin. 2016. "The Value of Coordination." *80,000 Hours*. February 8.
<https://80000hours.org/2016/02/the-value-of-coordination/>
- Tuomela, Raimo. 2013. *Social Ontology: Collective Intentionality and Group Agents*. Oxford University Press.
- Twain, Mark. 1884. *The Adventures of Huckleberry Finn*.
- Woodard, Christopher. 2011. *Reasons, Patterns, and Cooperation*. New York: Routledge.