

Cooperation: With or without Shared Intentions*

Jules Salomone-Sehr

McGill University, Centre de Recherche en Éthique

jules.salomone@mcgill.ca

*Final Draft. Please cite the published version from *Ethics*.*

This paper articulates our everyday notion of cooperation. First, I topple an orthodoxy of shared agency theory by arguing that shared intentions to φ are neither necessary, nor sufficient for φ to be cooperative. I refute the necessity claim by providing examples of shared intention-free cooperation (in institutional contexts and beyond). I refute the sufficiency claim by observing that coercion and exploitation need not preclude shared intentions but do preclude cooperation. These arguments, in turn, lead to my positive proposal. People cooperate, I argue, just in case their activities are coordinated in ways that do not undermine any participant's agency.

1. Introduction

The notion of cooperation has received sustained discussion and led to formal elaboration in a variety of often intersecting bodies of literature: evolutionary biology and psychology,¹ decision theory, game theory, and evolutionary game theory,² as well as moral and political philosophy.³ A more informal, everyday sense of cooperation has also made crucial appearances in other areas of philosophy where shared agency is not the main focus of theoretical attention. In epistemology, Edward Craig⁴ has argued that as humans have little choice but to rely on each other's word, they have come to form the concept of knowledge in order to refer to what good

Cooperation: With or Without Shared Intentions

informants provide through ongoing epistemic cooperation. In philosophy of language, Paul Grice has famously argued that implicatures occur when a speaker assumed to be cooperative flouts the maxims which, if observed, ensure in general the observance of what he has termed the Cooperative Principle.⁵

In this paper, I attempt to articulate this informal, everyday sense of cooperation.⁶ In order to do so, I suggest taking a detour via the philosophical literature on shared agency. Contributors to that literature have focused a lot of their theoretical attention to the sharing of intentions, and have suggested that cooperation be, if not reduced to activities resulting from shared intentions,⁷ at least reduced to a sub-class of such activities.⁸ My wager is that engaging critically with this claim and inquiring into how cooperation relates to and, more importantly, differs from the phenomena that have taken center stage in philosophical theories of shared agency will deliver an adequate account of our everyday concept of cooperation.

This paper has a negative and a positive part. I first argue that shared intentions to φ are neither necessary nor sufficient for φ to count as cooperative. Put differently, we may well be able to act together in a cooperative manner even when we do not share relevant intentions about our activity; and we may well fail to cooperate even when we act together in ways that satisfy some intention we share. If this is right, then the study of cooperative activities and the study of activities resulting from shared intentions should be kept, to some extent, separate.

These claims may at first seem paradoxical because, on the surface, the concept of cooperative activity and the concept of activities resulting from shared intentions appear closely related. To use Michael Bratman's example,⁹ suppose you and I share the intention to paint our house. If I deliberately buy a paint color you said you did not want, then I may undermine your sharing the intention to paint the house, especially if other similarly priced paint colors were

Cooperation: With or Without Shared Intentions

available and you knew that they were. Examples of that sort superficially suggest that uncooperative behavior stands in the way of the sharing of intentions; as a result, they might tempt us to conclude that collective actions resulting from shared intentions all require cooperation. But as I subsequently argue, it turns out that parties to an activity resulting from intentions they share might fail to act cooperatively.

Conversely, many shared agency theorists believe that people do not truly act together unless they share relevant intentions about their activities.¹⁰ As a result, they are led to conclude that cooperative activities all require shared intentions.¹¹ But as I subsequently argue, this is a mistake. In fact, parties to a projected activity need not share intentions for this activity to be cooperative.

The arguments developed in the negative part of the paper pave the way to my positive claim: cooperation consists in coordinated activities whose parties do not undermine the others' agency, for instance through coercion or exploitation. One important message conveyed by this claim is that under our everyday sense of cooperation, certain power relations prevent coordinated activities, whether resulting from shared intentions or not, from counting as cooperative. Another important message is that our everyday sense of cooperation does not require that participants be as tightly aligned as shared agency theorists have believed in mistakenly regarding shared intentions to be necessary for cooperation. Accordingly, a benefit of my positive claim is that it captures the whole continuum of cooperative activities, and not just the narrow set of cases on which the literature has predominantly focused.

The paper proceeds as follows. I first explain the significance of my denial of what I call the *necessity claim*, namely the claim that shared intentions are necessary for the performance of cooperative activities (§2). After arguing for some minimal conditions on shared intentions (§3),

Cooperation: With or Without Shared Intentions

I then make my argument against the necessity claim (§4). In particular, I develop the counterexample that role occupants of structured organizations need not share intentions in order to cooperate. This argument, I hope, might contribute to rectifying the relative neglect of institutional role-based cooperation in the shared agency literature.¹² After that, I argue for my denial of the *sufficiency* claim, namely the claim that collective actions that satisfy intentions we share are always cooperative (§5). I subsequently explain the significance of that denial (§6). I then take stock and articulate the (I take it, everyday) conception of cooperation that my arguments suggest (§7).

2. Pluralism about the Collectivity Condition

The inquiry into shared intentions to ϕ is one of the main games in town for shared agency theorists. There is a line of reasoning which, whether explicitly endorsed or unreflectively assumed, might account for this state of affairs.

Let me first give some background. Shared agency theorists understand one of their central tasks to consist in explaining how collective actions relate to, and differ from, individual actions. An important aspect of that task is to articulate the principled reasons we might have to draw a distinction between merely parallel individual actions (e.g. individual investors each selling their shares of some company), possibly resulting in non-additive outcomes (e.g. the company becoming insolvent), and cases of genuine collective actions (e.g. investors *concertedly* selling en masse their shares of some company in a boycott to protest the company's human rights violations). In examples of the first kind, the consensus is that individuals' agential capacities fail to be connected in ways required for their aggregate behaviors to count as a

Cooperation: With or Without Shared Intentions

genuine collective action. Let us call *Collectivity Condition* the condition which, when met by a collection of agents, ensures that their resulting action is genuinely collective (synonymously, shared). To summarize the foregoing, shared agency theorists seek, among other things, to articulate the way—or ways—in which that *Collectivity Condition* may be satisfied.

Let me now spell out the line of reasoning that might account for the sustained discussion that shared intentions have received in the literature.¹³ It is widely agreed that, at the individual level, a behavior counts as an action just in case the intentions of the individual behind this behavior are relevantly involved in its production.¹⁴ On this view, behaviors that owe nothing to the intentions of their author (e.g. a sudden cough or sneeze) fail to count as actions. They are *mere* behaviors that happen to their authors without involving their agential powers. From this popular view about individual actions and intentions, one might be tempted to draw by analogy the conclusion that at the collective level too, a collection of behaviors constitute a collective action just in case shared intentions of the individuals behind such behaviors are relevantly involved in their production.¹⁵ In so arguing, these shared agency theorists in effect commit themselves to a kind of monism about the *Collectivity Condition*: on their view, there is exactly one way through which the *Collectivity Condition* might be satisfied—namely, acting on shared intentions.

There is quite a bit that can be disputed in that argument by analogy (to wit, the initial premise about individual agency, the analogy itself between individual and collective actions, the conclusion). In fact, some shared agency theorists have already taken issue with it. To give an example, Sara Rachel Chant has shown that though there may well be an analogy between individual and collective actions, it does not warrant the conclusion that shared intentions are

Cooperation: With or Without Shared Intentions

required, in some relevant sense, for some collection of actions to count as genuinely collective.¹⁶

In section §4, I show that shared intentions are not necessary for cooperation (that is, I reject what I earlier called the necessity claim). Given that any cooperative activity is collective, this denial is significant insofar as it amounts to a rebuttal of the conclusion of the previous argument, and more specially its implication that monism about the Collectivity Condition holds true. To be sure, I do not mean to deny that shared intentions successfully acted on result in genuinely collective actions. Instead, my denial of the necessity claim negates the idea that shared intentions pave the only available route to cooperative, *hence collective*, endeavors. In so arguing, I in effect offer a defense of pluralism about the Collectivity Condition.¹⁷

Enough with the background. Let me now proceed to my defense of the denial of the necessity claim. In other words, let me show that shared intentions are not required for the performance of cooperative activities.

3. Shared Intentions and Closeness of Content

In the arguments I offer against the necessity and sufficiency claims (respectively, in §4 and §5), I deliberately remain quite vague about what I mean by cooperation and I rely for the most part on intuitions. In doing so, my hope is that the arguments of these two sections sketch the contours of a conception of cooperation that I more fully articulate subsequently (§7). By contrast, I must specify what I mean by shared intentions to φ .¹⁸

At least to some extent, shared *intentions* resemble individual intentions. What are intentions? From a role-functional perspective, intentions are intentional states, mental or

Cooperation: With or Without Shared Intentions

otherwise,¹⁹ that typically conclude an episode of practical reasoning (whether fully explicit or somewhat subpersonal) and that commit one to a course of action. The distinctive force of this commitment stems from the cluster of rational requirements that intentions bring in their train: absent special circumstances, it is required, when intending to φ , (roughly) to choose means likely to result in successful φ -ing, to avoid forming new intentions that turn out impossible with φ -ing, to refrain from giving up one's intention to φ when faced with a surmountable challenge, etc. Though intentions must commit their possessors to some course of action, what that course of action is need not be painstakingly specified: after all, the content of our intentions at the initial planning stages of some project often only makes reference to some vague goal and a vague plan to attain it. When an intention is unspecific, however, means-end coherence will require that further deliberation specify further, in a timely fashion, both the goal and the corresponding plan. A consequence of this (that I will make use of later) is that though the content of intentions need not be finely specified, it cannot be so unspecific that it does not constrain further specifications—for otherwise we would be committed to so little that we would not count as intending anything.

Now that these uncontroversial features of intentions are on the table, I will assume that such features help us characterize shared intentions. Indeed, if they are to be intentions in more than a merely metaphorical way, shared intentions must possess roughly the same role-functional profile as that of individual intentions.

Additionally, if a shared intention is to be *shared* in, again, more than a merely metaphorical way, either of the following must be true:

Cooperation: With or Without Shared Intentions

- Supposing it makes sense to attribute a shared intention *collectively* to an entire collection of agents (that is, supposing we can meaningfully attribute intentions not just to individuals who compose a collective, but to the collective itself), then that shared intention must have sufficiently determinate propositional content, for otherwise it would fail to commit us to anything;
- Supposing it makes sense to attribute shared intentions *distributively* to each member of a collection of agents, then individual so-called “we-intentions”,²⁰ that is the individual intentions of the participants to some joint endeavor, must have narrowly overlapping propositional contents.²¹

To summarize the foregoing, if a collection of agents share an intention, then there is a goal (albeit relatively vague) and/or plan (albeit relatively unspecific) that sharers of that intention can all be roughly said to intend—in some way or other, whether distributively or collectively or both.²² Let us call *Closeness of Content* this minimal condition that any adequate account of shared intentions should impose on the content of such intentions.

Major accounts of shared intentions satisfy Closeness of Content. For instance, Bratman’s account satisfies the distributive version of Closeness of Content. On his view, if you and I share the intention to paint our house together, then both of us individually have the intention that we paint our house together; thus, the content of my intention narrowly overlaps with the content of yours. Margaret Gilbert’s account too satisfies Closeness of Content, but unlike Bratman’s, hers meets the collective version of that requirement.²³ On her view, our intention to go on a walk is analyzed in terms of a joint commitment to intend as a body to go on that walk. Thus, using her form of words, the intention ascribed to the *plural subject* we thereby form has determinate propositional content, namely that we go on that walk together.

Cooperation: With or Without Shared Intentions

It is commonly assumed that cooperative activity must involve shared intentions, where these intentions satisfy Closeness of Content; but, as I shall argue in the following section, this is a mistake. There can be cooperation where the parties to the cooperative activity do not have any such shared intentions. In other words, parties to a cooperative activity need not commit themselves, not even roughly, to achieving the same goal or following the same plan.

4. Denial of the Necessity Claim

Consider institutional-role based cooperation in structured organizations (e.g. a trade union, army, or firm). To make things more concrete, consider a rather big firm with a complex production process, many different roles, many hierarchical layers, and a rather fine-grained division of labor. I claim that occupants of such roles may successfully cooperate even when they do not share intentions in any of the above senses.²⁴

First off, role occupants do not have to all individually intend to maximize the profits of the firm, secure its market shares or what have you, to be engaged in cooperation. Intuitively, it seems that they can be content with fulfilling the various obligations and responsibilities carried by their specific roles in order to cooperate; and it also seems that fulfilling such obligations and responsibilities need not require that role occupants commit themselves either to the achievement of some shared goal or the completion of some shared plan. After all, I may well be a good employee and cooperate successfully with my coworkers on the assembly line, as well as with my boss, even if I could not care less about the profits of the firm that hires me, and even if I did not know what role my individual contributions play in the firm's complex production process.

Cooperation: With or Without Shared Intentions

This shows that distributive ascriptions of shared intentions, that is ascriptions of we-intentions with narrowly overlapping contents, are not necessary for cooperation.²⁵

But here is a worry: the intuition that role occupants of the firm cooperate might be driven by the thought that they share, in the collective sense, some overarching intention (for instance, the intention to maximize profits), even if, as I just argued, they do not have individual intentions with narrowly overlapping contents. In other words, one might think that the collective formed by all the role occupants of the firm possesses some intention, and if that were true, their cooperation might well just be a consequence of this intention they share in the collective sense. At best, the big firm would then only offer a case of cooperation without intentions shared in the *distributive* sense, and so I would still have to find (if any) an example of cooperation where intentions are neither distributively, *nor collectively*, shared.

Let me dispel this worry; that is, let me show that when pushed out the door under their distributive guise, shared intentions do not come back through the window under their collective guise.

I grant that the activities of all role occupants of the firm may pursue collective ends. Indeed, the role occupants' combined activities surely serve some common purpose or other. For instance, their activities might maximize the profit of the firm, as well as result in the sales of the firm's products. Moreover, the common purposes served by the employees' activities can plausibly figure in functional explanations of what the employees do and how they do it. For instance, the employees of the firm work staggered shifts *because* staggered shifts increase overall productivity, hence contribute to profit maximization. However, ends and purposes differ from *intentions*. Though maximizing profit may be one of the ends of the firm, the firm's employees need not be committed to profit maximization—in the way that they would be if

Cooperation: With or Without Shared Intentions

indeed they shared the intention to maximize profit. To recap, even if the employees' activities served a common purpose, it would not follow that the employees shared some intention or other, either in the collective or in the distributive sense.²⁶

Now that the distinction between ends and intentions is on the table, could it make sense to attribute intentions, and not just ends, to a collection of agents as such even when there is no intention that agents of that collection distributively share? And if it does, does cooperation require intentions shared in this collective sense when intentions shared in the distributive sense go missing?

Some prominent shared agency theorists have taken seriously collective attributions of shared intentions.²⁷ Christian List and Philip Pettit on the one hand,²⁸ and Margaret Gilbert on the other,²⁹ have attempted to give a precise content to this idea.³⁰ Part of their view is that a correct collective attribution of some shared intention to some collection of agents need not go hand in hand with the possession, by each of these agents, of we-intentions with narrowly overlapping contents.³¹ To use one of Gilbert's examples,³² though together we share the intention to take a walk, it may well be the case that none of us individually intends to take that walk.³³

My goal here is not to assess the merits of either Gilbert's or List and Pettit's views. Instead, I will show that even in the absence of intentions shared in the distributive sense, neither Gilbertian, nor Listian and Pettitian shared intentions are required for cooperation. Going back to my example, I claim that role occupants of the large firm under scrutiny may cooperate even when they do not possess intentions shared collectively either à la List and Pettit, or à la Gilbert.

Two obstacles stand in the way of a correct attribution of a Gilbertian shared intention to role occupants of the firm.³⁴ First, these role occupants may not have created the roles they play

Cooperation: With or Without Shared Intentions

and the corresponding responsibilities and obligations they take on, and this runs against Gilbert's view that joint commitments must be created by their parties.³⁵ Second, as argued earlier, role occupants, when cooperating, need not act in light of shared goals but instead in light of the responsibilities and obligations carried by their roles, and such responsibilities and obligations need not make reference to shared goals. This in turn runs against Gilbert's view that successful fulfillment of a shared intention requires that parties to that intention be motivated by it.³⁶

Turning now to List and Pettit, they argue that groups of agents may literally acquire minds of their own. This happens when such groups have procedures through which their members pool (either mechanistically³⁷ or flexibly³⁸) their intentional states, both doxastic and motivational. The resulting pool of intentional states may differ from the states entertained by any of the members of the group, which justifies regarding the states of this resulting pool as those of the group itself. To illustrate, a hiring committee might form a group-level intention to hire some candidate or other based on a set of selection criteria the members of the committee voted on, as well as based on the majority view about how each application they received met these selection criteria. Under certain circumstances, the resulting group-level intention to hire John, the candidate that has scored the best on most criteria, might not be had by any member of that hiring committee.³⁹

With this rough and ready summary of List and Pettit's view in mind, let me argue that group-level intentions are not required for cooperative activities. Going back to the foregoing example, I claim that role occupants, content with fulfilling their role-based responsibilities and obligations, might cooperate even though the firm for which they work do not possess group-level intentions. Think of a (rather old-fashioned) firm where the director makes the decisions

Cooperation: With or Without Shared Intentions

alone. He⁴⁰ may well gather relevant information about the decision-making context thanks to his employees' reports, but what happens in his firm is ultimately responsive to his preferences, projects, and more generally, his motivational states only. In this scenario, it would not make much sense to characterize the director's intentions as group-level intentions; if anything, such intentions are exclusively *his*. But there seems nothing to stand in the way of saying that all role occupants of that firm (including perhaps the old-fashioned director) are all cooperating with one another: provided that no one's agency is undermined,⁴¹ their activities are cooperative.

In addition to showing that intentions shared collectively à la List and Pettit are not necessary for cooperation, the case of the old-fashioned firm suggests that collectively shared intentions *tout court* are not required for cooperation. Indeed, no plausible account of intentions shared in the collective sense should predict not only that intentions of that kind are present in the old-fashioned firm, but also that cooperation in that scenario depends on such intentions.

My argument has so far made use of examples of institutional role-based cooperation. But I believe it can be extended to more informal and interpersonal cases of cooperation. Consider Dominique and Claude. Claude is at his local gym, climbing the gym wall. Dominique, a friend of his, sees him up there and gets excited: she was hoping to find someone with whom she could practice her climbing instructing skills. After greeting each other, Dominique starts giving instructions to Claude about the holds he should grab next, and Claude takes Dominique's instructions into account by adjusting his moves accordingly. In so doing, they certainly cooperate—provided neither Dominique nor Claude undermines the other's agency.⁴² Now, do they intend the achievement of the same goal and/or the completion of the same plan? Not necessarily.

Cooperation: With or Without Shared Intentions

First, they need not share intentions in the distributive sense. To be sure, given that Dominique is supplying instructions to Claude, it is natural to suppose that she intends the concatenation⁴³ of her actions (i.e. her oral instructions to Claude) and Claude's own (i.e. his adherence to such instructions). By contrast, this concatenation of actions need not figure in Claude's intentions. Perhaps Claude is not really in the mood to interact with anyone, let alone receive instructions. As a result, though Claude might *expect* Dominique to keep on offering her instructions, he might not *intend* her to do so. That said, as he might be no more in the mood to upset Dominique by overtly asking her to stop indicating which holds he should grab next (or worse, by ignoring her instructions), he might not mind observing such instructions, knowing that Dominique will not be long anyway before she turns her attention to her own workout. In this example, it is therefore not true that Dominique and Claude share an intention whose content represents the concatenation of their activities. Although they both enact a plan in which one person is the teacher, guiding the other person, the student, through the process of climbing the wall,⁴⁴ this plan need not form any part of Claude's intentions.⁴⁵

Don't Dominique and Claude at least distributively share some intention or other about the specific holds that Claude grabs on the climbing wall? Perhaps, but not necessarily. Consider this specific scenario. Putting to the test an instructional trick she heard about, Dominique tells Claude to grab holds which she thinks Claude cannot yet reach; in doing so, her goal is to draw his attention to nearby holds which she intends him to see and be able to use only at the next attempt. However, out of a desire to have an effective workout, Claude makes special efforts and manages to grab the holds Dominique considers out of reach for him.

It follows from the foregoing that Dominique and Claude do not have we-intentions with narrowly overlapping contents; and yet they are surely cooperating nonetheless.

Cooperation: With or Without Shared Intentions

One might perhaps object that in this example, there is some admittedly vague goal that both Dominique and Claude distributively intend to achieve, and which expressions like “having fun,” “doing something together,” or “continuing whatever it is that they started doing” go some way towards capturing. If that is true, so the objection goes, then the purported example of institutional role-free cooperation without shared intentions fails.

In light of Claude’s preference to be left alone, and the resulting fact that it is not his intention to be instructed by Dominique, I believe it is already obvious that Claude does not share with Dominique an intention with any of these vaguely described contents. But perhaps not everyone shares my belief, so here is an additional argument why it is a mistake to ascribe a shared intention with any such unspecific contents to Dominique and Claude:

- Either the propositional contents expressed by the aforementioned phrases (“having fun,” “doing something together,” etc.) are so vague that they neither significantly constrain further practical deliberations, nor relevantly guide Dominique and Claude’s activities. As shared intentions with contents of this very vague sort would provide virtually no orientation for Dominique and Claude’s practical energies, they turn out explanatorily idle.
- Or contents captured by expressions like “having fun,” “doing something together,” etc., are specific enough to constrain further specifications and practical deliberations, as well as guide resulting activities, but only because of (e.g.) the specific kind of fun that Dominique and Claude respectively have in mind. In that case, Dominique and Claude’s intentions fail to have narrowly overlapping contents, for Claude’s fun (if that is any fun at all) resides in refraining from upsetting Dominique whereas Dominique’s fun lies in her practicing her instructing skills.

Cooperation: With or Without Shared Intentions

To put it differently, either the content of Dominique's intention to (e.g.) have fun with Claude is too open-ended to be explanatorily indispensable; or that content would already be too narrowly specified for Dominique's intention to be effectively shared in the distributive sense with Claude's.⁴⁶

To recap so far, Dominique and Claude do not share intentions in the distributive sense. Do they share intentions in the collective sense? Gilbert for one would, perhaps, redescribe Dominique and Claude as manifesting readiness to jointly commit to intend, as a body, to do whatever it is that they end up doing. But the redescription of Dominique and Claude's cooperative interaction as involving an intention shared in the collective sense is neither here nor there. Indeed, Dominique and Claude seem too loosely connected to have incurred a joint commitment to some course of action, especially a joint commitment of the kind that Gilbert theorizes. Suppose Claude went down the climbing wall, called it a day without seeking Dominique's authorization for him to do so, and took leave politely. If Dominique and Claude were both parties to a Gilbertian joint commitment, Dominique would have a right to rebuke Claude, even despite Claude's polite leave-taking.⁴⁷ But given the specifications of the case, if Dominique indeed rebuked Claude, Claude would surely be justified in responding to the rebuke with raised eyebrows. More generally, Dominique and Claude's interaction is importantly analogous to the earlier case of the old-fashioned firm: as my discussion suggests, the plan Dominique and Claude enact figures exclusively in Dominique's intention—just as the plan that workers of the old-fashioned firm enact figures exclusively in the content of the director's intention. It is therefore a mistake to regard Dominique and Claude as sharing any intention in the collective sense.⁴⁸

Cooperation: With or Without Shared Intentions

So far, I have raised and ultimately refuted the objection that Dominique and Claude's interaction involves, despite my claims to the contrary, some intention or other shared either distributively or collectively. Alternatively, one might dispute my claim that Dominique and Claude are cooperating, and thereby challenge my view that their interaction exemplifies cooperation without shared intentions. But this line of objection strikes me as implausible. After all, Dominique and Claude's respective activities do not mesh accidentally: again, they form part of a single plan where Dominique, the would-be climbing instructor, guides Claude as he climbs the gym wall. Granted, Dominique and Claude's cooperation requires less psychological alignment than dominant accounts of shared agency demand: to reiterate, their plan does not figure in Claude's intentions. I believe, however, that this example, and for that matter the other examples of this section, show that these dominant accounts have gone wrong: they have failed to capture attenuated, yet genuine, cases of cooperation for having mistakenly regarded shared intentions to be necessary for cooperation.⁴⁹

This concludes my case against the necessity claim.

5. Denial of the Sufficiency Claim

I now critically engage with the sufficiency claim, i.e. the claim that any activity resulting from shared intentions is cooperative. In his treatment of collective intentionality, John Searle⁵⁰ sounds like he simply accepts the sufficiency claim uncritically when he writes that: "The notion of a we-intention, of collective intentionality, implies the notion of *cooperation*."⁵¹ His considered view, however, expresses a more qualified endorsement of the sufficiency claim⁵² which, if my reconstruction is right, goes like this:

Cooperation: With or Without Shared Intentions

- Activities may be either cooperative or uncooperative depending on the descriptions one uses to pick them out;⁵³
- Consider some activity; though it might be uncooperative under some description, if there is some other description that figures in the content of intentions shared by parties to that activity, then that activity is cooperative under that very description.

To illustrate, suppose that you and I play chess, and I end up winning. Context permitting, I might talk about our playing chess using the following description: the activity that eventuated in your humiliating defeat. It is pretty clear that under that description, our activity was not cooperative. We might still have shared intentions about our game of chess. If that was the case, then according to this version of the sufficiency claim, whatever description of that activity figured in our shared intention (e.g. “our abiding by the rules of chess for our moves to count as chess moves”) would be a description under which our activity was cooperative. This version of the sufficiency claim is the one I will now refute. Put differently, I am about to show that the description of an activity that figures in the content of a shared intention might well be a description under which that activity is uncooperative.⁵⁴

Consider the following scenario. Azealia and Bryn, members of the same soccer team, can score a goal together using any of the following three profiles of strategies: either they distribute their efforts equally and, as a result, share the admiration of the spectators; or Azealia (symmetrically, Bryn), intending to outshine her teammate, contributes significantly more to the effort than Bryn (symmetrically, Azealia), and as a result secures the lion’s share of the praise. All three strategy profiles are equally likely to secure a goal for Azealia and Bryn’s team, provided that neither Azealia nor Bryn change their mind once one of these profiles is underway, on pain of not scoring. By contrast, if both players attempt to outshine one another, they will not

Cooperation: With or Without Shared Intentions

reach the level of coordination required to score, an outcome both least prefer. When they manage to score a goal, both players prefer courses of action that result in their receiving the biggest share of their spectators' admiration.⁵⁵

Suppose Azealia does not attempt to outshine Bryn, and Bryn, taking advantage of Azealia, proceeds to outshine her. Suppose in addition that Bryn's intended action can be captured by the following characterization: "scoring by way of my outshining Azealia." Nothing precludes the possibility that Azealia forms an intention whose content narrowly overlaps with the content of Bryn's and which the expression "scoring by way of my being outshone by Bryn" surely goes a long way towards capturing. In fact, Azealia might have her reasons for including her being outshone in the content of her intention: she knows herself, and in particular she knows that if the way she conceives of the action she commits herself to performing does not highlight her being irreversibly outshone by Bryn, she will tend to lose sight of that fact; as a result, her unchecked competitiveness will lead her to attempt to in turn outshine Bryn, and because of that, no goal will be scored. If the foregoing is plausible, we have the crucial building blocks of an intention shared by Azealia and Bryn, whichever your preferred account of shared intentions might be.⁵⁶

So, it appears that Azealia and Bryn's activity might be the outcome of an intention they share and whose content describes their activity along the following lines: "scoring by way of Bryn outshining Azealia." But quite intuitively, this activity is not cooperative under such a description, for cooperation is usually paradigmatically contrasted with cases where one's cooperativeness is taken advantage of by co-participants.

Is it the case, however, that when an activity results from some shared intention, there is always at least one way of reformulating the content of that shared intention in a way that we share

Cooperation: With or Without Shared Intentions

intentions with that reformulated content, and that reformulated content describes our activity in ways that draw our attention away from any uncooperativeness? For if this were the case, at least some—admittedly diluted—version of the sufficiency claim would end up vindicated. Under that diluted version, of all the descriptions of our activity that figured in intentions we shared, at least some would be descriptions under which our activity would turn out cooperative. Going back to Azealia and Bryn, if both share the intention to score by way of Bryn outshining Azealia, they also surely share the intention to score together, and that description of their joint activity does not carry on its sleeve any indication of uncooperativeness.

When cast in these terms, the issue is little more than terminological. Thus, the vindication of this diluted version of the sufficiency claim would offer at best a Pyrrhic victory to its advocates. Besides, it is worth noting that there seems to be one usage of the term “cooperation” and its cognates in ordinary language that runs counter to even this diluted version of the sufficiency claim.⁵⁷ Suppose that prior to the soccer game, Bryn coerced Azealia into accepting to be outshone—for instance by using intimidation and threats. Azealia and Bryn may still share intentions, but insofar as their activity of goal-scoring originates in coercion, the claim that that activity is cooperative, under that usage of “cooperative,” now sounds inadequate regardless of the description one uses to pick that activity out, hence including under any descriptions that figure in the content of their shared intentions—even those descriptions that do not wear any indication of uncooperativeness on their sleeve.⁵⁸

This concludes my case against the sufficiency claim.

Cooperation: With or Without Shared Intentions

6. Significance of the Denial of the Sufficiency Claim

We saw in §4 that the sharing of intentions is not the unique route to cooperative, hence collective activities. The point that emerges now is that the collective actions to which shared intentions give rise may involve all sorts of power relations (coercion, exploitation, oppression, etc.) that block such actions from being cooperative. This, in turn, highlights the most serious risk run by advocates of the sufficiency claim: they will have to ignore cases where people share intentions and yet fall short of sharing their agency in any way we would recognize as cooperative.

More specifically, advocates of the sufficiency claim will be conceptually unable to acknowledge the fact that the powerful can sometimes utilize shared intentions as part of their exploitative tricks. To illustrate, think how managers may encourage employees to care about the profits of their firm even if higher profits may not benefit them in any substantial positive way—and may even harm them if the pursuit of such higher profits requires strenuous effort. Even without going that far, advocates of the sufficiency claim will fail to do justice to the practical conflicts experienced by those who, like Azealia, are being used and yet involved in endeavors that satisfy intentions they share with others. To summarize, shared intentions are certainly a salient way of sharing agency in a manner likely to be cooperative; but advocates of the sufficiency claim fail to recognize that this agency-sharing process is no guarantee of genuine cooperation.

7. Cooperation

At this point, one might object that the two claims that I defended so far in this paper, namely the claim that cooperation need not require shared intentions (§4), and the claim that activities satisfying intentions we share need not be cooperative (§5), might not involve the same conception of cooperation. As a result, so the objection goes, my attempt to sketch an account of cooperation through my denial of the necessity and the sufficiency claims fails, as there may be more than one sense of cooperation at work in the intuitions I have implicitly relied upon regarding what counts as cooperative.

However, I believe that the senses of cooperation involved in my previous arguments need not differ. In fact, these arguments help clarify a familiar everyday conception of cooperation that has played a prominent role in some philosophical theories and arguments, for instance in Grice's account of speaker-meaning⁵⁹ and in Craig's account of testimonial exchanges between what he has called good informants.⁶⁰ In this section, I first articulate this everyday conception of cooperation (§7.1 and 7.2).⁶¹ I then respond to the objection that this conception of cooperation is too weak (§7.3).

7.1. *Coordination Condition*

First, this everyday conception is one according to which cooperative activities require coordination between a plurality of agents. Engaging in joint endeavors with others and exerting some amount of control over our causal powers when they are already intertwined is surely something we desire for a variety of reasons, prudential, ethical, or otherwise. One of the salient manners through which we gain control over such causal powers is by achieving coordination,⁶²

Cooperation: With or Without Shared Intentions

that is, by making it the case that our activities interlock in the service of a common purpose (or a common set of purposes).

This purposive interlock of individual activity is what the Coordination Condition deems essential to cooperation. More specifically, the Coordination Condition asserts that a collection of agents is engaged in cooperation only if their respective activities are coordinated, that is, only if such activities interlock in the service of a common purpose (or set of purposes).

Grice thought that cooperative efforts involved “a common purpose or set of purposes, or at least a mutually accepted direction.”⁶³ The conception of coordination here sketched, together with my view that coordination is necessary for cooperation, capture his insight. It is worth stressing that this conception of coordination differs from a weak conception on which all that it takes for individual behaviors or activities to be coordinated is that they be merely reciprocally adjusted to one another. To compare these two conceptions, consider Eli and Karim, two boxers locked in a clinch in the middle of a match. The conception of coordination involved in my Coordination Condition does not classify their clinch as a coordinated action. For though Eli’s moves are reciprocally adjusted to Karim’s (and so are weakly coordinated to them), they are meant to frustrate Karim’s ends—and vice versa. The entanglement of their arms and shoulders is the mere byproduct of conflicting ends and, thus, lacks a common purpose. Notice, however, that the mere pursuit of one and the same purpose is not quite sufficient for coordination: you and I might pick up plastic bottles on the same beach, hence aim at the same purpose (i.e. that the beach be clean), and yet we might be cleaning that beach in a wholly uncoordinated fashion. On the conception of coordination that I am invoking, our activities must *interlock*, that is, be relevantly interdependent, for them to be coordinated.

Cooperation: With or Without Shared Intentions

The foregoing conception of coordination differs from—and, I believe, is superior to—game theorists' own conception of coordination. In game theory, a collection of agents is traditionally said to coordinate when their activities form a combination such that any party's unilateral deviance from that combination is in some way costly for herself and for others. A combination of actions with that property is what David Lewis calls a coordination equilibrium.⁶⁴ However, though costly deviance might explain the viability of some coordinating schemes, it does not capture what makes a collection of activities count as coordinated in the first place. For one thing, individuals might coordinate even despite individual incentives pulling apart from coordination, hence, even when the scheme on which they coordinate is not an equilibrium. More importantly, what makes some collection of actions coordinative in the first place does not lie so much in costly unilateral deviance as in the purposive mesh that the actions of that collection combine to form.

How can collections of activities be made to interlock in the service of a common purpose, hence count as candidate instances of cooperative activities? In the same way that the formation and the acting on intentions help individuals coordinate their temporally extended agency, that is, enable them to bring activities performed at different times under a common purpose (or common set of purposes), agents may come to share intentions so as to coordinate their activities and thereby bring their practical energies under a common purpose (or common set of purposes).⁶⁵ But as shown by my denial of the necessity claim,⁶⁶ cooperative, hence coordinated, activities do not require shared intentions. There must therefore be other mechanisms than just shared intentions through which collections of activities interlock in the service of a common purpose.⁶⁷

A full account of these alternative coordination mechanisms is beyond the scope of this paper. Nevertheless, the examples developed in §4 offer valuable insight about what such

Cooperation: With or Without Shared Intentions

mechanisms might be. In particular, they suggest that institutionally defined roles in structured organizations as well as more informal roles, but also practices of issuing and acting on instructions, commands, and requests constitute coordination processes that channel our combined agential powers and put them in the service of common purposes.

Let us pay one last visit to the big firm from §4. The employees working on the assembly line have surely been instructed to perform certain tasks in a certain way—and so they do. Now, even if they do not care about, or are not aware of, the purposes the supervisor had in mind when they issued their commands, it is nevertheless the case that the employees' contributions serve the supervisor's purposes. Assuming that the supervisor meant their supervisees' contributions to be part of a single project, such contributions interlock in the service of a common purpose (namely, to achieve the supervisor's project). To summarize, when a collection of activities done by distinct agents satisfy instructions, commands, or requests issued with a view to achieving a single project, such activities are coordinated—even in the absence of shared intentions.^{68, 69}

Let us now pay one last visit to Dominique and Claude's local gym. As argued previously, Dominique and Claude both play their part in a plan where the would-be teacher guides the reluctant student through the process of climbing the wall. Their activities, thus, arguably pursue the purposes that teaching and learning serve (e.g. the dissemination of knowledge), and these purposes, in turn, provide a blueprint for their interaction (e.g. Dominique is expected to instruct, Claude is expected to practice, etc.). And yet, Claude is not personally committed to any of that: he half-heartedly sticks to observing Dominique's advice about which hold to grab next, while looking forward to the moment when Dominique will eventually turn her attention to her own workout. In short, a collection of activities that satisfy interlocking roles are coordinated—even in the absence of shared intentions.⁷⁰

Cooperation: With or Without Shared Intentions

Different coordination mechanisms impose epistemic requirements of varying stringency. On most accounts, parties to shared intentions to φ must have common knowledge that they all intend to φ .⁷¹ By contrast, role-based coordination arguably might only require that role occupants have correct beliefs about what their respective roles entail.⁷² As a result, role-based coordination does not require that parties' psychologies register the fact that they are indeed in the business of coordination. As it turns out, role occupants might sometimes have reason to suspend judgment whether coordination is underway. In one example, Ludwig invites us to imagine nuclear missile operators at distinct locations following their respective roles in a sophisticated missile launching procedure after their country has been hit by a preemptive strike.⁷³ Communication is now impossible and, given the attack they have suffered, each of these operators has some reason to doubt that the other operators have survived. As a result, they all seem justified in suspending judgment whether the missile launching procedure, which requires the contribution of all, is taking place. Assuming the operators have survived, however, they surely count as engaged in role-based coordination when launching the counterstrike, even while they are far from satisfying any version of the common knowledge condition.⁷⁴

Similarly, different coordination mechanisms come with normative requirements of varying stringency. Parties to a coordinated activity may be tightly normatively entangled, for instance when mutually bound by a joint commitment à la Gilbert: under such circumstances, they frame the performance of their respective parts as an obligation which, when flouted, warrants criticism and rebuke. Alternatively, participants may decide in advance for a 'no strings attached' scenario where doing one's part and remaining involved in the resulting collective endeavor are conceived of as beyond the call of duty.⁷⁵

Cooperation: With or Without Shared Intentions

7.2. *Agency Condition*

Second, as shown by my critical discussion of the sufficiency claim (§5-6), this everyday conception of cooperation is one according to which certain power relations prevent coordinated activities from counting as cooperative. Azealia and Bryn's case suggests that taking advantage of others or exploiting them, possibly through deception, conflicts with the pursuit of cooperative activities. Similarly, coercion stands in the way of cooperation. To take Bratman's example,⁷⁶ Japanese troops and British prisoners of war might share the intention to build a bridge over the River Kwai, but the Japanese troops' coercive power prevents the resulting activity to count as cooperative.

These observations argue in favor of the Agency Condition, that is, the claim that a collection of agents are φ -ing cooperatively only if no agent of that collection has undermined, or attempted to undermine, their co-participant's agency as part of the execution of φ (e.g. by bypassing it, as in the case of coercion,⁷⁷ or by hijacking it, as in the case of exploitation).^{78, 79}

The Agency Condition captures Craig's insight when he claims, in his account of epistemic cooperation, that teamwork requires that knowers treat each other not merely as sources of information, but as informants. In order to explain this distinction further, he writes: "What I have in mind is the special flavour of situations in which human beings treat each other as subjects with a common purpose, rather than as objects from which services, in this case true belief, can be extracted."⁸⁰ The Agency Condition generalizes Craig's insight by ensuring that parties to genuinely cooperative ventures do not treat co-participants as mere props from which agential energy can be extracted.

The Agency Condition also captures the grain of truth contained in the game-theoretic conception of cooperation in terms of mutually beneficial collective activity.⁸¹ Uncoerced non-

Cooperation: With or Without Shared Intentions

exploitative activities are indeed very often activities from which parties mutually benefit. That said, it is worth noting that under a relatively thin conception of benefit, exploitation and coercion may well sometimes be mutually beneficial. To take an example with a Marxist undercurrent, under certain circumstances (which Marx considered typical of capitalism), exploited workers are surely better off than members of the reserve army of labor (that is, the unemployed). Unlike the game-theoretic conception of cooperation, mine does not count the mutually beneficial yet exploitative activities of capitalist firms as cooperative.⁸²

In §5, I showed that shared intentions may well result in coordinated actions that do not satisfy the Agency Condition. Put differently, I established that the sharing of intentions is a coordination process that does not guarantee cooperation. At this point, it is nonetheless still an open question whether some other coordination mechanism automatically leads to cooperative coordination. Might there be, then, a coordination process, or possibly some combination of such processes, that always lead to collective activities where no party's agency is threatened, or otherwise undermined? In other words, might we have good reason to uphold some version of the sufficiency claim that would advert not just to shared intentions, but also to other coordination mechanisms and, possibly, their combinations?

I do not think so. Section §5 showed that all sorts of power relations may prevent activities resulting from shared intentions from counting as cooperative. I now suggest that the point applies more broadly to any coordination mechanism as well as any combination of such mechanisms. Put differently, we arguably have reasons to deny the truth of any version of the sufficiency claim.

A full defense of this denial would require, first, the identification of all the coordination mechanisms there are, and, second, for any such coordination mechanism (and any combination

Cooperation: With or Without Shared Intentions

of them), an example where some participant flouts the Agency Condition, hence prevents the resulting coordination from counting as cooperative. A defense of this sort obviously lies beyond the scope of this paper.

But I still believe that we have enough evidence to conclude inductively in favor of a blanket denial of any version of the sufficiency claim. We know too well that not all institutional role-based joint activities are cooperative: examples of exploitative firms are all too abundant. We know too well that our practices of issuing and acting on instructions, commands, and requests can fall short of cooperation: abuses of the authority sometimes required to issue such commands and requests are far too common. More generally, it seems that no coordination mechanism is immune to exploitation, coercion, and the like.⁸³ If I am right that there is no infallible recipe for cooperation, then cooperation is more of an art than the rigid application of some formula.

7.3. *An Overly Weak Conception of Cooperation?*

Even despite the Agency Condition, some might think that my account of cooperation is overly weak. For though this condition requires that parties to cooperative activities do not undermine co-participants' agency, the Coordination Condition might be argued to leave open that overly diluted interactions count as instances of cooperation. Consider one such putative counterexample, namely commuters who adjust the speed and direction of their respective trajectories as they walk along in a crowded subway station. One might fear that the Coordination Condition does not block commuters' moves from being classified as a case of cooperation.⁸⁴ If this objection were on point, then the two conditions fleshed out in the previous

Cooperation: With or Without Shared Intentions

subsections (§7.1 and 7.2), though perhaps individually necessary for cooperation, would fail to be jointly sufficient.

However, I think it is a virtue of my account that it regards such commuters as cooperating (if their moves satisfy in addition the Agency Condition). Though they need not share any intention, commuters of this example clearly do something together: they are avoiding collision and thereby sharing a crowded public space.⁸⁵ In fact, employees of the public transport authority in charge of supervising rush hour transit might thank commuters who proceed in an orderly manner for their cooperation.

In so arguing, am I casting my net too wide and focusing attention away from the core phenomena that shared agency theorists have attempted to explain? I do not think so. In fact, I believe that my view helps us to refocus on the very subject matter of shared agency theory by drawing attention not just to activities that satisfy intentions we share, but to all collective activities.⁸⁶ The necessity claim, i.e. the claim that shared intentions are necessary for cooperation, has led shared agency theorists astray in encouraging them to concentrate exclusively on an overly narrow set of cases where parties to collective actions are tightly knitted. As a result, attenuated instances of cooperation have gone largely unaddressed.⁸⁷ By contrast, my Coordination Condition only requires that participants' activities interlock in the service of a common purpose. Unlike necessity claim theorists, I am therefore able to capture several kinds of relatively diluted cooperative activities, namely, cognitively unsophisticated cooperation (e.g. commuters' coordinated trajectories), institutional role-based cooperation in large structured organizations⁸⁸ (e.g. workers' interconnected tasks in big firms), as well as cooperative activities in which some agents participate reluctantly (e.g. Dominique and Claude's interaction as described in section §4).

Cooperation: With or Without Shared Intentions

Alternatively, one might think that the Agency Condition is to be blamed for the overly weak conception of cooperation that this paper articulates. Advocates of this line of reasoning might perhaps regard the Coordination Condition as sufficient for engaging in collective activity. However, they contend that the Agency Condition fails in adequately drawing the line between cooperative and uncooperative collective activities. In addition to not coercing, exploiting, and deceiving co-participants, it might be suggested that genuine cooperation requires, for instance, a willingness to help others if need be.⁸⁹

I grant that a disposition for mutual help is surely desirable in a great many cooperative activities. In fact, in structured organizations where cooperation is the outcome of well-designed institutional roles, it might be a judicious idea to set up and enforce role-based obligations to assist co-participants.⁹⁰ However, I believe it is a mistake to regard a willingness for mutual help as a necessary feature of cooperative enterprises. Consider a singing competition in which contestants are at some point asked to sing in duets but judged nonetheless individually for their ability to be in unison. Given that singers of such duets ultimately compete against each other, they surely have no desire, and besides, arguably no reason, to go out of their way to assist each other. Nevertheless, as they sing together, they are surely cooperating (provided of course that none of them has flouted the Agency Condition).⁹¹

My conception of cooperation is weaker than the conceptions found in the shared agency literature. However, in light of the foregoing, this is not a flaw, but a strength: dominant conceptions were unjustifiably restrictive in the first place.

8. Conclusion

I have argued that shared intentions are neither necessary nor sufficient for cooperation. I have then articulated the everyday conception of cooperation that the denial of the necessity and sufficiency claims bring in sharper focus.

I have first explained the significance of the denial of the necessity claim: if that denial is correct, then contrary to what most shared agency theorists have assumed, shared intentions do not pave the only available route to genuinely collective activities. I have then denied the necessity claim. As I have argued, the interpersonal coordination required for cooperation may be for instance achieved by the satisfaction of the responsibilities and obligations carried by institutional roles, and the satisfaction of such responsibilities and obligations need not involve intentions shared either distributively by each role occupant or collectively by all role occupants taken together.

I have then refuted the sufficiency claim: intentions to ϕ may be shared by a collection of agents even when the agency of some members of that collection is undermined by other members. A consequence of this is that though intentions constitute a rather salient mechanism regularly used to share our agency in a cooperative fashion, this mechanism is not fool-proof: people might well fail to cooperate even when their activities satisfy intentions they share.

I have finally argued for the positive conception of cooperation that the arguments of this paper suggest. On that conception, a collection of agents cooperate if and only if their activities are coordinated and, when in the business of coordination, the agency of no agent of that collection is undermined by some other agent of that collection. This conception succeeds in including certain intuitively cooperative activities unfortunately excluded by other accounts,

Cooperation: With or Without Shared Intentions

while excluding cases of coordinated activity where some agents have diminished their co-participants' agency.

NOTES

* I have incurred many debts of gratitude at various stages of development of the ideas presented in this paper. Thanks to Sacha Bourgeois-Gironde, Charles Mills, Elisabeth Pacherie, David Papineau, Matthew Rachar, and Katherine Ritchie for their insightful comments. I presented an earlier version of this paper at the 2018 Social Ontology Conference (Tufts University), as well as at a Colloquium at the 2019 Pacific Meeting of the American Philosophical Association (Vancouver). I thank attendees for their questions and challenges, in particular, Caroline Arruda for her wonderful written comments, as well as Gunnar Björnsson, Stephanie Collins, Karen Jones, and Kirk Ludwig. This paper has also benefitted greatly from the invaluable comments of two anonymous referees and three editors at *Ethics*. My deepest thanks go to Jennifer Morton for her encouragements and precious advice, and to Miranda Fricker for her extremely generous support and extensive discussion of successive drafts.

¹ See (e.g.) Michael Tomasello, *Why We Cooperate* (Cambridge, MA: MIT Press, 2009), Samuel Bowles and Herbert Gintis, *A Cooperative Species: Human Reciprocity and Its Evolution* (Princeton: Princeton University Press, 2011), Philip Kitcher, *The Ethical Project* (Cambridge, MA: Harvard University Press, 2011), and Kim Sterelny, Richard Joyce, Brett Calcott, and Ben Fraser, *Cooperation and Its Evolution* (Cambridge, MA: MIT Press, 2013).

² See (e.g.) Robert Axelrod, *The Evolution of Cooperation*, rev. ed. (New York: Basic Books, 2009), Michael Bacharach, *Beyond Individual Choice: Teams and Frames in Game Theory* (Princeton: Princeton University Press, 2006), Ken Binmore, *Natural Justice* (Oxford: Oxford University Press, 2005), Hervé Moulin, *Axioms of Cooperative Decision Making* (Cambridge: Cambridge University Press, 1991), Brian Skyrms, *The Stag Hunt and the Evolution of Social*

Cooperation: With or Without Shared Intentions

Structure (Cambridge: Cambridge University Press, 2004), and *Evolution of the Social Contract*, 2nd ed. (Cambridge: Cambridge University Press, 2014).

³ See (e.g.) Donald H. Regan, *Utilitarianism and Co-Operation* (Oxford: Oxford University Press, 1980), David P. Gauthier, *Morals by Agreement* (Oxford: Oxford University Press, 1986), Martin Hollis, *Trust Within Reason* (Cambridge: Cambridge University Press, 1998), Elizabeth Anderson, “Symposium on Amartya Sen’s Philosophy: Unstrapping the Straitjacket of ‘Preference’: A Comment on Amartya Sen’s Contributions to Philosophy and Economics,” *Economics and Philosophy* 17, 1 (2001): 21-38, John Rawls, *A Theory of Justice*, rev. ed. (Cambridge, MA: Harvard University Press, 2009), and Philip Pettit, *The Birth of Ethics: Reconstructing the Role and Nature of Morality* (Oxford: Oxford University Press, 2018).

⁴ Edward Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Oxford University Press, 1990).

⁵ “Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.” In Paul Grice, *Studies in the Way of Words* (Cambridge, MA: Harvard University Press, 1989), 26.

⁶ For another attempt at conceptual clarification, see Geoffrey Brennan and Geoffrey Sayre-McCord, “On ‘Cooperation,’” *Analyse & Kritik* 40, 1 (2018): 107-30.

⁷ Kirk Ludwig, for instance, contends that cooperative activities, under a thin conception of cooperation, boil down to activities resulting from shared intentions. See Kirk Ludwig, “What Is Minimally Cooperative Behavior?” in *Minimal Cooperation and Shared Agency*, ed. Anika Fiebich, 9-39 (New York: Springer). To be fair, however, he also argues that cooperation, under a richer conception, requires more than shared intentions (*ibid.*, 13-14).

Cooperation: With or Without Shared Intentions

⁸ See (e.g.) Michael Bratman, *Shared Agency: A Planning Theory of Acting Together* (Oxford: Oxford University Press, 2014). To use his terminology, Bratman suggests that shared cooperative activities (that were the focus of his seminal 1992 paper) are a sub-class of shared intentional activities (Bratman, *Shared Agency*, 38, 87, 102, 121). This would seem to indicate that for Bratman, cooperation requires shared intentions. See ft. 15 below for an important qualification.

⁹ See Michael Bratman, “Shared Cooperative Activity,” *The Philosophical Review* 101, 2 (1992): 327-41, as well as *Shared Agency*.

¹⁰ See (e.g.) John Searle, “Collective Intentions and Actions,” in *Intentions in Communication*, ed. Philip R. Cohen, Jerry Morgan, and Martha Pollack, 401-415 (Cambridge, MA: MIT Press, 1990), p. 420, Christopher Kutz, “Acting Together,” *Philosophical and Phenomenological Research* 61, 1 (2000): 1-31, p. 3-10, Abraham Sesshu Roth, “Shared Agency and Contralateral Commitments,” *The Philosophical Review* 113, 3 (2004): 359-410, p. 361, Philip Pettit and David Schweikard, “Joint Actions and Group Agents,” *Philosophy of the Social Sciences* 36, 1 (2006): 18-39, p. 23-24, Facundo M. Alonso, “Shared Intention, Reliance, and Interpersonal Obligations,” *Ethics* 119, 3 (2009): 444-75, p. 445, and Margaret Gilbert, *Joint Commitment: How We Make the Social World* (Oxford: Oxford University Press, 2013)—though see *ibid.*, 12 for a qualification.

¹¹ For an explicit endorsement of this claim, see (e.g.) Cédric Paternotte, “Minimal Cooperation,” *Philosophy of the Social Sciences* 44, 1 (2014): 45-73, p. 64, and Ludwig, “What Is Minimally Cooperative Behavior?” p. 18.

¹² For inspiring exceptions, see Scott J. Shapiro, *Legality* (Cambridge, MA: Harvard University Press, 2011), “Massively Shared Agency,” in *Rational and Social Agency: The Philosophy of*

Cooperation: With or Without Shared Intentions

Michael Bratman, ed. Manuel Vargas and Gideon Yaffe, 257-293 (Oxford: Oxford University Press, 2014), Ludwig, *From Plural to Institutional Agency: Collective Action II* (Oxford: Oxford University Press, 2017), and Katherine Ritchie, “Minimal Cooperation and Group Roles,” in *Minimal Cooperation and Shared Agency*, ed. Anika Fiebich, 93-110 (New York: Springer).

¹³ For a similar diagnosis, see Sara Rachel Chant, “Collective Action and Agency,” in *The Routledge Handbook of Collective Intentionality*, ed. Marija Jankovic and Kirk Ludwig, 13-24 (Abingdon: Routledge, 2017).

¹⁴ See (e.g.) Donald Davidson, “Actions, Reasons, and Causes,” *Journal of Philosophy* 60, 23 (1963): 685-700, and “Agency,” in *Agent, Action, and Reason*, ed. Robert Binkley, Richard Bronaugh, and Ausonio Marras, 3-25 (Toronto: University of Toronto Press). In using above the admittedly vague locution ‘relevantly involved’, my intention is to avoid getting bogged down in a discussion of rather thorny issues of action theory. To be more specific:

- I said ‘relevantly’ in order to leave a placeholder for the theoretical moves that would be required to accommodate unintentional actions;
- I said ‘involved’ in order to sidestep the debate between Anscombian and advocates of the causal theory of action regarding the explanation of the contrast between mere behaviors and genuine actions.

Given my purposes, these matters need not worry us.

On a related note, the claim that the involvement of intentions distinguishes actions from mere behaviors should not be read as entailing that intentions are involved in the same uniform way in the production of *all* actions. One might think, for instance, that in contrast with highly reflective and deliberate actions, the connection between intentions and merely habitual actions is, in some

Cooperation: With or Without Shared Intentions

sense, attenuated. Nothing that I am saying in the main text precludes this possibility. I thank an associate editor for inviting me to clarify this point.

¹⁵ For a similar line of reasoning, see Raimo Tuomela, *The Philosophy of Sociality: The Shared Point of View* (Oxford: Oxford University Press, 2007), 108-109, and Bratman, *Shared Agency*, 9-10. This passage of Bratman's book would seem to be evidence of his commitment to what I earlier called the necessity claim (see ft. 8 above for further evidence). Now, in all fairness, Bratman primarily aims to offer *sufficient* conditions for what he calls modest sociality (*ibid.*, 36, 105), that is, small-scale shared agency between adults with no (or minimal) institutional structure, no asymmetry of authority, and where the participants remain the same over time. He also concedes that modest sociality is multiply realizable (*ibid.*, 169, endnote 86). As a result, he seems open to the idea that there might be small-scale shared activities where the agency-sharing mechanism is *not* a shared intention. (I develop one such case at some length below in the second half of §4.) I thank an anonymous referee for helping me to realize that Bratman, after all, might well not be a necessity claim theorist.

¹⁶ More specifically, she argues that some collective behaviors that are collectively intentional under *no* description may yet count as collective actions. See Sara Rachel Chant, "Unintentional Collective Action," *Philosophical Explorations* 10, 3 (2007): 245-256. Ludwig too holds a view that conflicts with the conclusion of the foregoing argument by analogy. Roughly, he argues that an action is collective just in case there is an event of which each member of some collection of individuals is an agent. On his view, global warming might be considered the outcome of a collective action since it might be regarded as the outcome of an event of which each of us is an agent. See Ludwig, *From Individual to Plural Agency: Collective Action I* (Oxford: Oxford University Press, 2016), 138-144.

Cooperation: With or Without Shared Intentions

¹⁷ Though my denial of the necessity claim entails that shared intentions are not necessary for cooperative, hence collective, behavior, I do not believe this denial shows that shared intentions are not necessary for *collectively intentional* behavior. The question whether shared intentions are necessary—and/or sufficient—for collectively intentional behavior is, indeed, left open by the arguments developed in this paper. On a related note, a full account of what it would take for behaviors to count as collectively intentional is beyond the scope of this paper.

¹⁸ Throughout the paper, I consider ‘shared intentions to φ ’ and ‘shared intentions that we φ ’ to be synonymous locutions. In doing so, I bracket a terminological nuance to which some contributors to the shared agency literature attribute a relatively important role (e.g. Bratman, *Shared Agency*, 63-64).

¹⁹ David Velleman is open to the possibility that intentions need not be mental states (and for that matter, need not be states at all, but rather representational *acts*). He makes great use of this possibility in his account of shared intentions. See David Velleman, “How to Share an Intention,” chap. 9 in *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000).

²⁰ Wilfrid Sellars coined this term which Raimo Tuomela and Kaarlo Miller later on commandeered for different theoretical purposes. See Wilfrid Sellars, “Imperatives, intentions, and the logic of ‘Ought’,” in *Morality and the language of conduct*, ed. Hector-Neri Castañeda and George Nakhnikian, 159-218 (Detroit: Wayne State University Press, 1963), “Objectivity, Intersubjectivity and the Moral Point of View,” chap. 7 in *Science and Metaphysics: Variations on Kantian Themes* (London: Routledge), and Raimo Tuomela and Kaarlo Miller, “We-Intentions,” *Philosophical Studies* 53, 3 (1988): 367-389. For my purposes, I define we-intentions as the individual intentions of parties to a shared activity. This usage does not commit

Cooperation: With or Without Shared Intentions

me to the truth of any of the specific accounts of we-intentions found in the literature. In particular, it does not commit me to taking onboard Searle's controversial view that we-intentions constitute a sui generis type of intentions. See Searle, "Collective Intentions and Actions."

²¹ I said "narrowly overlapping propositional contents" instead of "identical contents" in part to spare myself potential worries about the semantics of indexicals (e.g. Can indexicals be eliminated from propositional attitudes without altering the content of such attitudes?), as well as related worries about the alleged distinctive role of the first-person concept in propositional attitudes (especially in the content of intentions). In order to stave off these distracting issues, my suggestion is to remain neutral on sameness of content and, instead, to call attention to uncontroversial deep similarities between contents. To illustrate, I will assume, then, that the content of *your* intention that you and I φ , and the content of *my* intention that you and I φ , narrowly overlap.

²² I owe the realization of the usefulness of the distinction between distributive and collective attributions to Ludwig, *Collective Action I*.

²³ See Gilbert, *Joint Commitment*.

²⁴ The argument I offer generalizes Ritchie's (see Ritchie, "Minimal Cooperation and Group Roles").

²⁵ In the shared agency literature, Shapiro (in *Legality* and "Massively Shared Agency") is one of the few philosophers whose views are in tension with the necessity claim. Specifically, he has argued that members of a group engaged in shared agency need not possess what he terms shared plural intentions, that is, intentions that we φ shared in the distributive sense (*ibid.*, 275-276, 282). In so arguing, he denies that intentions shared à la Bratman are necessary for shared

Cooperation: With or Without Shared Intentions

agency. In his own account of shared agency, however, Shapiro still requires that participants be committed to a shared plan (*ibid.*, 280) to count, strictly speaking, as having acted together with others (*ibid.*, 283). The example above establishes that even this requirement is overly restrictive: satisfying the obligations and responsibilities attached to one's role, hence doing one's bit in cooperation with others, need not entail any commitment to the plan of which these obligations and responsibilities are a proper part.

It is also worth noting that my argument against the necessity claim is more general and systematic than Shapiro's. His, by contrast, mostly focuses on Bratmanian shared intentions—though he also briefly discusses Gilbert's and Kutz's accounts (Shapiro, "Massively Shared Agency," 276-277). I thank an anonymous referee for asking me to clarify the respects in which my arguments and views differ from Shapiro's.

²⁶ In my positive account of cooperation, I make use of this idea that our activities might pursue a common purpose (synonymously, a collective end) even while we do not share the intention (either collectively or distributively) to pursue that purpose. See §7.1 as well as ft. 69. A background assumption of all this is the uncontroversial claim that an activity might pursue some purpose even when that purpose does not figure in an intention of the agent of that activity (as when the purpose of my driving below the speed limit is to reduce the risk of accidents on the road, even though my personal intention in doing so is only to avoid getting a speeding ticket). To sloganize, *purposes can be external*. Similarly, then, a collection of activities by distinct agents might pursue a common purpose even when that common purpose does not figure in an intention that these agents share. I thank an anonymous referee for requesting clarification on this point.

Cooperation: With or Without Shared Intentions

²⁷ Some other shared agency theorists, by contrast, have taken issue with such collective attributions. Now, notice that if, like Ludwig (*Collective Action I*, 184-186), you do not countenance collective attributions of intentions, then my denial of the necessity claim is complete. For then neither intentions shared distributively (as shown before), nor intentions shared collectively (an absurdity, on your view), are required for cooperative activities.

²⁸ See Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford: Oxford University Press, 2011). It is worth noting that List and Pettit's account of collective attributions of intentions (and other mental states, for that matter) is not primarily intended to advance our understanding of collective activity but to develop a systematic theory of temporally extended group agency. Their target phenomenon is, thus, quite different from shared agency theorists'. I thank an anonymous referee for inviting me to clarify this point.

²⁹ See Gilbert, *Joint Commitment*.

³⁰ Velleman ("How to Share an Intention") has done that too, but unlike the authors just mentioned, it would not make sense, on his view, to collectively attribute an intention to a collection of agents if, distributively, such agents could not be said to all have intentions with narrowly overlapping contents.

³¹ Carol Rovane too holds a similar view. See Carol Rovane, *The Bounds of Agency: An Essay in Revisionary Metaphysics* (Princeton: Princeton University Press, 1998).

³² See Gilbert, "Walking Together: A Paradigmatic Social Phenomenon," *Midwest Studies in Philosophy* 15, 1 (1990): 1-14.

³³ That Gilbert believes shared intentions to be shared in the collective sense is, I take it, a standard interpretation of her view. Yet, some of Gilbert's readers disagree with it. Recently,

Cooperation: With or Without Shared Intentions

Ludwig has argued that for Gilbert, it is not true that people bound by a joint commitment form a plural subject (Ludwig, *Collective Action I*, 261-264). Rather they act *as if* they formed one. To be sure, Gilbert does claim that a collection of individuals share an intention to φ if and only if they are jointly committed to intending as a body to φ (Gilbert, *Joint Commitment*, 114). But if Ludwig's interpretation were correct, this would be shorthand for being jointly committed to acting as if they formed a single body that had the intention to φ . On Ludwig's interpretation, therefore, collective ascriptions of intentions would, strictly speaking, form no part of Gilbert's view. This interesting interpretation, however, ignores Gilbert's commitment to what she calls the *disjunction criterion*, that is, the claim that "an adequate account of shared intention is such that it is not necessarily the case that for every shared intention, on that account, there be correlative personal intentions of the individual parties" (Gilbert, *Joint Commitment*, 105). As the disjunction criterion makes it clear, Gilbert believes that a collection of agents might share an intention even while the agents of that collection do not have individual intentions with overlapping contents. In such scenarios, what, then, is the bearer of the shared intention if not the collection of agents itself?

³⁴ See Ritchie, "Minimal Cooperation and Group Roles."

³⁵ See Gilbert, *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society* (Oxford: Oxford University Press, 2006), 138, 168, and *Joint Commitment*, 40, 117.

³⁶ See Gilbert, *Political Obligation*, 146.

³⁷ E.g. via majority rule.

³⁸ List and Pettit, *Group Agency*, 60-63.

³⁹ This example illustrates what List and Pettit call a premise-based procedure (*Group Agency*, 56).

Cooperation: With or Without Shared Intentions

⁴⁰ In this version of the example of the big firm, I use the masculine pronoun to refer to the firm's director only to match the old-fashioned character of the management style—about which in no way do I feel nostalgia.

⁴¹ I here adumbrate the conception of cooperation I articulate further in §7.

⁴² Again, this anticipates my full account of cooperation.

⁴³ I am here borrowing Bratman's term (*Shared Agency*, 92-93).

⁴⁴ I thank an anonymous referee for suggesting this apt language to capture the plan that undergirds Dominique and Claude's interaction.

⁴⁵ The presentation of this case has greatly benefitted from the anonymous referees' comments.

⁴⁶ To be sure, I do not deny that some shared activities originate in shared intentions whose content is initially relatively vague (though not too vague, for otherwise such intentions would commit oneself to virtually nothing). After all, our Friday evening spent at the movies might owe much to a prior shared intention whose initially vague content was that we would do something together over the weekend. To enable coordination in action, however, initially vague intentions of that sort must be specified further as the time of action gets closer—first, by deciding to go to the movies as opposed to an art show, then by picking the movie theater, etc. Shared intentions with a vague content can therefore figure in an explanation of what a collection of agents have done together, but typically only under the condition that such intentions have led, at some point in the planning process, to more specific intentions. I thank an anonymous referee for asking me to clarify this point. Now, notice that in the example developed above, Dominique and Claude did not share an initially vague intention that they understood as having to be refined and specified later. (For that matter, neither did they subsequently form shared intentions with a more specific content.) Therefore, it is hard to see how the shared intentions with vague content that

Cooperation: With or Without Shared Intentions

the objector to my view is ready to ascribe distributively to Dominique and Claude could figure in a plausible explanation of Dominique and Claude's coordinated actions. In §7.1, I sketch a more persuasive picture of the coordination mechanism involved in Dominique and Claude's interaction—one that does not involve shared intentions. This alternative picture will add support to my claim that the vague intentions invoked by the objector are indeed explanatorily idle.

⁴⁷ On leave-taking, see Javier Gomez-Lavin and Matthew Rachar, "Normativity in joint action," *Mind and Language* 34, 1 (2019): 97-120.

⁴⁸ The alternative explanation of Dominique and Claude's coordination in action that I sketch in §7.1 will provide further evidence that the collective ascription of some shared intention or other to Dominique and Claude would have no explanatory power regarding the coordination of their activities.

⁴⁹ I expand on this line of objection, as well as my response to it, in §7.3.

⁵⁰ In a similar vein, see Ludwig, "What Is Minimally Cooperative Behavior?" Ludwig believes that as soon as we satisfy intentions that we share, we minimally cooperate, even if (e.g.) you coerced me into doing my part.

⁵¹ Searle, "Collective Intentions and Actions," 406 (his emphasis).

⁵² *Ibid.*, 413-414.

⁵³ As a result, there would not be such a thing as the class of cooperative activities.

⁵⁴ There is another version of the sufficiency claim which I take Bratman to have successfully refuted. This other version asserts that if φ satisfies intentions we share, then there cannot be *any* description of φ under which φ is uncooperative. This assertion is false: as argued by Bratman (*Shared Agency*, 102), Japanese troops and their prisoners of war might intend to build a bridge over the River Kwai even though the former coerced the latter into taking part in the construction

Cooperation: With or Without Shared Intentions

effort. However successful, Bratman's argument does not refute the version of the sufficiency claim on which I here focus: for all he has argued, it might still be that the description of the building of the bridge that figures in the content of the Japanese troops and the prisoners' shared intention is a description on which what they do together shows no sign of uncooperativeness (just like "our abiding by the rules of chess for our moves to count as chess moves," in the example developed above, showed no sign of uncooperativeness). To refute the version of the sufficiency claim on which I focus, it will not be enough to offer a case of shared φ -ing resulting from a shared intention where φ is uncooperative under some description or other. What we want, instead, is that the very description of φ that figures in the content of that shared intention is a description that wears uncooperativeness on its sleeve.

⁵⁵ The decision problem Azealia and Bryn face is an instance of game theorists' Hawk-Dove.

⁵⁶ In particular, Azealia and Bryn's shared intentions might well satisfy the strictures of Bratman's account. It is worth stressing that as stated, this case leaves fully open the possibility that both players intend their teammate's intentional agency to be involved in the achievement of the joint move. Put differently, the case might well satisfy Bratman's requirement that Azealia (respectively, Bryn) intend the achievement of the joint move in part by way of Bryn's (respectively, Azealia's) intention (Bratman, *Shared Agency*, 50). Therefore, Azealia and Bryn's joint move need not be construed as structurally akin to Bratman's mafia case, where members of competing gangs intend to go to New York by way of throwing the other into the trunk of some car, hence by making use of means that completely bypass the other's intentional agency.

⁵⁷ Bratman, *Shared Agency*, 37-39.

Cooperation: With or Without Shared Intentions

⁵⁸ Whether we should adopt the foregoing usage of the term “cooperation” is a question I do not have to decide given my purposes in this paper. Bratman commits himself to this usage—albeit relatively tentatively (Bratman, *Shared Agency*, 38).

⁵⁹ See Grice, *Studies in the Way of Words*.

⁶⁰ See Craig, *Knowledge and the State of Nature*. I do not mean to suggest that the sense of cooperation I articulate in this section is identical to the sense of cooperation involved in Grice and/or Craig’s views. In fact, a systematic comparison of my conception to theirs would require more exegesis than I can handle in this paper. I still believe that my analysis does justice to some of Grice’s and Craig’s insights—as I explain respectively in §7.1 and in §7.2. I thank an anonymous referee for inviting me to clarify this point.

⁶¹ Both subsections have considerably benefitted from the anonymous referees’ invaluable comments and requests for clarification.

⁶² Another salient, albeit ethically suspect, manner through which we gain control over our causal powers is brute force.

⁶³ Grice, *Studies in the Way of Words*, 26.

⁶⁴ See David Lewis, *Convention: A Philosophical Study* (Hoboken, NJ: Wiley-Blackwell, 2002), 14.

⁶⁵ See Bratman, *Intention, Plans, and Practical Reason* (Cambridge, MA: Harvard University Press, 1987), 28.

⁶⁶ The necessity claim I argued against stated that shared intentions were necessary for cooperation. There is another, broader, necessity claim in the vicinity of the first that states that shared intentions are necessary for coordination. But this broader necessity claim is obviously

Cooperation: With or Without Shared Intentions

false, just because the narrower is false (see §4) together with the fact that all cooperative activities are coordinated.

⁶⁷ Using the vocabulary introduced earlier (§2), there must be other mechanisms than just shared intentions through which collections of activities satisfy the Collectivity Condition. This alternative wording implicitly relies on the intuitive claim that coordinated activities, that is, activities that interlock in the service of a common purpose, are collective.

⁶⁸ This point does not hinge on the fact that the firm is a structured organization whose members play formal roles. For even instructions that I give informally with a view to achieving a single project surely bring the activities that conform to such instructions under the purpose (or set of purposes) that I have in mind.

⁶⁹ That the interlock of employees' contributions might serve a common purpose even though that purpose does not figure in the content of an intention the employees share is an instance of my previous claim that purposes can be external. See ft. 26. I should stress, then, that 'common purpose', as I am using the locution, is a bit of a term of art. It might be that, in common parlance, attributing a common purpose to a collection of activities done by distinct agents suggests that that purpose figures in an intention shared by these agents. This is not, however, how I am using this locution. Given my theoretical agenda, a common purpose takes no more than one and the same purpose pursued by a collection of activities—whether or not that purpose figures in the content of a shared intention. I thank an associate editor for inviting me to clarify this point further. Notice also, in passing, that in the foregoing example, the supervisees' activities are interdependent, hence can be said to interlock, because they are steps of a single plan of action (namely, the supervisor's project).

Cooperation: With or Without Shared Intentions

⁷⁰ Notice that in this example, though Dominique issues instructions to Claude, she neither instructed, nor requested Claude to play the role of a student. This suggests that practices of issuing and acting on instructions, commands, and requests on the one hand, and interlocking roles on the other, are two *distinct* coordination mechanisms.

⁷¹ See (e.g.) Bratman, *Shared Agency* and Gilbert, *Joint Commitment*. There are some exceptions. See (e.g.) Ludwig, *Collective Action I*, 219-221.

⁷² I said ‘might’ because one could argue that even this minimal epistemic condition need not be met in certain cases of cooperation, for instance in cases where well-designed incentives nudge role occupants into satisfying the requirements of their roles even when they are unaware of these requirements, and perhaps also unaware of the very fact that they happen to play some role. People mindlessly following footprint stickers in subway tunnels, thereby ensuring unknowingly that the flow of commuters be smooth provide perhaps an example of such unaware cooperation. For similar examples, see Richard Thaler and Cass Sunstein, *Nudge: Improving Decisions about Health, Wealth, and Happiness* (New Haven, CT: Yale University Press, 2008).

⁷³ Ludwig, *Collective Action I*, 220-221.

⁷⁴ Although Shapiro is sympathetic with the idea that roles might enable coordination, his account includes a common knowledge condition (Shapiro, *Legality*, 138, and “Massively Shared Agency,” 277). If the foregoing remarks are on the right track (see also ft. 72), this common knowledge condition is unnecessarily restrictive.

⁷⁵ See Bratman, “Shared Intention and Mutual Obligation,” chap. 7 in *Faces of Intention: Selected Essays on Intention and Agency*, (Cambridge: Cambridge University Press, 1999), and *Shared Agency*, chap. 5.

⁷⁶ Bratman, *Shared Agency*, 102.

Cooperation: With or Without Shared Intentions

⁷⁷ Perhaps there are cases where coercion does not undermine agency, hence does not stand in the way of cooperation. Suppose we are in the business of robbing a jewelry, and noticing that, out of mere temptation and short-sightedness, you might well go for some low-hanging fruit (e.g. a couple of gold chains) that might distract you from a much bigger gain (e.g. myriads of flawless diamonds), I coerce you into sticking to the plan we have agreed upon (say, by waving my arm in your direction). My coercion, in this scenario, would help you to not give in to temptation and short-sightedness. Arguably, then, far from undermining your agency, this instance of coercion might even enhance it. If that scenario was plausible, then coercion would not necessarily, but merely *typically* preclude cooperation. Perhaps a claim of this sort is also true about exploitation. Be that as it may, my view, I should stress, is that it is the undermining of agency (whether effective or merely attempted) that stands in the way of cooperation.

⁷⁸ For those who think that whether φ is cooperative or not depends on the description one uses to pick φ out, the Agency Condition must be stated as follows: a collection of agents is engaged in cooperative φ -ing, under some description of φ , only if that description involves no reference to any agency-undermining behavior that an agent of that collection might have inflicted upon a co-participant as part of the execution of φ .

⁷⁹ There can be some fair amount of disagreement about whether a collection of coordinated activities (that is, a candidate instance of cooperation) satisfies or not the Agency Condition. People might reach different conclusions about whether an offer to take part in some collective activity was coercive, that is, whether one was made an offer *that can't be refused* (in the mafia sense of that locution). Similarly, whether I exploited you or not as part of our shared endeavor is something about which our intuitions might conflict. Insofar as coercion, exploitation, and the like are thick concepts (i.e. descriptive concepts that carry a normative load), and insofar as

Cooperation: With or Without Shared Intentions

coercion, exploitation, and the like stand in the way of cooperation (perhaps merely typically, see ft. 77), disagreements about whether coordinated activities count as cooperative might well have a normative inflection. I thank an anonymous referee for drawing my attention to this.

⁸⁰ Craig, *Knowledge and the State of Nature*, 36.

⁸¹ See Bowles and Gintis, *A Cooperative Species*, 2.

⁸² Notice, in passing, that one might flout the Agency Condition somewhat unintentionally, as this example suggests. After all, a business owner might have little choice but to exploit their employees in an economy where paying starvation wages is required to not go out of business. I thank an anonymous referee for inviting me to clarify this point.

⁸³ When I earlier argued against the sufficiency claim (§5), one might have worried that I chose all too easy a target. Accounts of shared activity found in the literature often involve not just shared intentions, but further ingredients (e.g. mutual obligations, mutual responsiveness, etc.), and it might be thought that the combination of all such ingredients (or at least some of them) ensures that the resulting shared activities are always cooperative. The argument against any version of the sufficiency claim I just sketched should allay these worries (helpfully brought to my attention by an anonymous referee).

⁸⁴ I thank an anonymous referee for pressing this objection and suggesting that I consider a similar example.

⁸⁵ Shared intentions will arguably not be necessary if these agents are weary, yet experienced, commuters, content with sticking to the more or less formal rules that provide for a smooth journey (e.g. following signs and arrows, standing on the right when using escalators, walking at a pace that makes one's trajectory easy to predict and to adjust to, etc.). More specifically, the contents of the individual intentions involved in satisfying such rules need not narrowly overlap,

Cooperation: With or Without Shared Intentions

especially if different commuters commit to different, yet compatible rules (e.g. when on an escalator, standers commit to standing on the right, while walkers commit to walking on the left). Granted, this shared intention-free strategy to navigating crowded subway stations might not be the most effective method to avoid collisions. It is nonetheless a plausible alternative to sharing intentions—especially for worn-out city dwellers.

⁸⁶ Or rather, to all *non-accidentally collective* activities. Car users all contribute to gas emissions. However, even if polluting the air we breathe counted as a collective activity (which some might perhaps dispute), this activity would only be accidentally collective and, as such, would not be the primary target of shared agency theory.

⁸⁷ Insofar as they attempt to capture these attenuated cases of collective actions, Kutz (“Acting Together”) and Shapiro (“Massively Shared Agency”) are exceptions. That said, my account requires even less than Kutz’s overlapping participatory intentions and Shapiro’s view that participants must be committed to a shared plan if they are to count as genuinely acting with others (*ibid.*, 280, 283).

⁸⁸ Cooperation of this sort is what Shapiro has called “massively shared agency” (*ibid.*).

⁸⁹ Here too, I thank an anonymous referee for pressing this objection.

⁹⁰ I owe this idea to an anonymous referee.

⁹¹ To be sure, we surely possess a concept of full cooperation in addition to the everyday concept of cooperation on which this paper has focused, and this concept of full cooperation certainly requires a willingness to help co-participants. See Ludwig (“What Is Minimally Cooperative Behavior?” 13-14) for an account of full cooperation.