

Minimal Cooperation and Group Roles¹

Katherine Ritchie
CUNY Graduate Center and City College of New York
kcritchie@gmail.com

forthcoming in *Minimal Cooperation and Shared Agency* (Springer, Ed. Anika Fiebich)

Abstract

Within philosophy, cooperation has primarily been analyzed in the context of theories of collective intentionality. These discussions have primarily focused on interactions between pairs or small groups of agents who know one another personally. Cooperative game theory has also been used to argue for a form of cooperation in large unorganized groups. Here I consider a form of minimal cooperation that can arise among members of potentially large organized groups (e.g., corporate teams, committees, and governmental bodies). I argue that members of organized groups can be minimally cooperative in virtue of playing roles in an organizational structure and having a common goal. The minimal form of cooperation considered here is not grounded in collective intentions involving symmetric mental states, special collective intentional modes, or joint commitments. More generally, I show how considering minimal cooperation in the context of organized groups provides an opportunity to reevaluate the extent to which the social world and social phenomena depend on internalist mental factors (e.g., intentions, beliefs) and externalist non-mental factors (e.g., laws, job descriptions). The view of minimal cooperation among members of organized groups developed here provides support for an externalist rather than internalist theory of at least one social phenomenon.

Keywords

cooperation; social groups; roles; externalism; collective intentionality

Within philosophy, cooperation has primarily been analyzed in the context of theories of collective intentionality. On these views, cooperation is taken to involve complex often symmetric mental attitudes. These theories center their discussions on small groups usually involving pairs of agents who interact with mutual knowledge of (some of) the others' intentions, beliefs, ends, and commitments (Bratman, Gilbert, Searle, Tuomela). An alternative account of cooperation relies on game theory. On these accounts, the focus is on cooperation in large unorganized groups, like participants in political demonstrations or flash mobs. Paternotte, a proponent of a using cooperative game theory to analyze minimal cooperation, takes collective intentionality accounts to center on "egregious cooperation" and argues that an analysis could be "weakened without making cooperation disappear" (2014, 48). Here I focus on a third sort of group—organized groups—and

¹ This paper has significantly improved through conversations and comments on previous versions. In particular I thank Derek Anderson, Kent Bach, Barrett Emerick, Jeremy Evans, Nick Leonard, Rebecca Mason, Matt Moss, Matthew Rachar, Jules Salome, Tuomo Tuusala, the members of an audience at a colloquium at the University of San Francisco, and an anonymous commentator.

develop a novel account of minimal cooperation that relies on neither collective intentionality nor cooperative game theory.

Organized groups—like corporate teams, committees, and legislative bodies—need not be composed of pairs or small groups of agents with mutual knowledge of collective beliefs, intentions, goals, and so on. They are not mere coalitions of agents, as in cooperative game theory. Instead, they involve individuals and organizational structures with roles for members to play. I argue that members of organized groups can be minimally cooperative in virtue of playing roles in an organizational structure and having a common goal. The view departs from the analyses of cooperation that begin with collective intentionality and those that begin with game theory. The minimal form of cooperation is not grounded in complex mental states/representations or in beneficial payoffs. We will see that roles individuals play in organizational structures are central to account for the form of minimal cooperation developed here.

Considering cooperation between members in organized groups will also prove illuminating in assessing the extent to which theories of social interactions, and perhaps social ontology more generally, should be construed in internalist or externalist terms. Call views of social phenomena and entities centered on mental representation, mental states, or mental contents **internalist views of the social**.² Call views that rely on features that are external to mental states or mental contents—for instance documents, laws, or material resources—**externalist views of the social**.³ When analyzing cooperative interaction between two agents working in close proximity, it is plausible that both agents have representations of the other agent and of (some of) her beliefs, desires, and intentions. That is, an account that is largely internalist is plausible. However, once we move to a scale at which there are multiple agents interacting in organizations that involve defined roles and action-types, the situation is different. Agents need not personally know all the people with whom they cooperate or have complex representations of others' mental states, or so I will argue. It is useful to have well defined roles for senators and for members of a corporate team as they require one to do less mental work to understand what others believe, intend, and so on. It is costly to represent others' mental states. It is far easier to represent roles which are simplified in ways that

² Even if one is a mental externalist, one could be an internalist about the social. On such a view one holds that social entities, facts, collective intentionality, etc. depend on mental states or mental contents that are themselves dependent on external features of the natural or social environment.

³ The distinction is not exclusive as views might include both mental and non-mental features. A view might be more or less internalist/externalist. See Torrenco (2017) for a recent defense of externalism about institutions. I discuss his view further in §IV.

agents—in their full psychologically complicated glory—are not. Playing a role is central to the more externalist account of minimal cooperation among members of organized groups developed here.

The article is structured as follows. I begin (§I) by considering the nature of organized groups. In particular, I consider the ways organized groups are structured to involve roles that are defined in ways that are interdependent. Then (§II) I consider three prominent views of collective intentionality and how they falter when extended to large organized groups. Since members of organized groups can be minimally cooperative and their cooperative behavior is not well captured by extending a theory of collective intentionality, we need to look elsewhere to understand this form of minimal cooperation. Next, (§III) I argue for a view of minimal cooperation that relies on roles group organizational structures impose and a common goal. Then I consider internalist and externalist views of the social and whether we should expect a unified account of cooperative activity (§IV). Most theorists hold largely internalist theories of social entities and phenomena according to which mental states and representations ground or constitute the social world and social phenomena. I show how the view of minimal cooperation among individuals in organized social groups supports externalism about at least some social phenomenon. Finally (§V) I draw concluding remarks.

§I. Organized Groups and Group Roles

Before an analysis of cooperation by members of organized groups can be undertaken, we need a better understanding of the nature of organized groups. In Ritchie (2013, 2015, forthcoming) I argued that groups like corporate teams, committees, clubs, the U.S. Senate, and British Parliament are entities with organizational structures. Organized groups are not merely structures. That is, they are not identical to structures. Rather they are entities that have members that are structured or organized in a particular way; they are structured wholes.

The structures of organized groups consist of roles and relations between them. Roles are defined in terms of relations to other roles, tasks that role-players are allowed or required to carry out, and in some cases specific features a role-player must have. Relations between roles might be hierarchical or non-hierarchical. Relations that involve deference and power are hierarchical. For instance, a role might allow a role-player to give orders to individuals playing other roles. Relations of seconding a motion or reporting on a project involve relations between group members that are non-hierarchical. Relations between roles also capture the ways playing a role depends on other roles being played. In discussing organizations Miller (2014) states that “there is a degree of

interdependence among ... roles” and that “the performance of the constitutive tasks of one role cannot be undertaken, or cannot be undertaken except with great difficulty, unless the tasks constitutive of some other role” are undertaken. Roles are largely defined in terms of relations; they are interdependent.

Roles are also defined in terms of powers, norms, and responsibilities. These might involve relations to other role players. For instance, some roles involve the power to give directives to other role holders. Norms about how to properly interact might also be part of what defines a role. Roles might also involve powers, norms, or responsibilities pertaining to particular tasks. For instance, some role might involve being responsible for the group’s media presence, while another gives power and responsibility to a role player to control fundraising efforts.

Finally group organizations might place particular restrictions on who can play a role. For instance, one might be required to be born in a particular country, be at least 35, or meet other external requirements like being appointed or elected by non-group members.⁴ A role may also allow for multiple role-players or only allow for a single role-player. This could also be specified by a requirement on the role itself.

Consider a simple example of a committee with roles for President, Vice President, Secretary, Treasurer, and Member without a leadership role. The role of the president might involve calling meetings to order, delegating tasks to members of the committee that do not have leadership roles, initiating votes, and so on. The president role might be specified so that only someone with a particular nationality can play the role. The Member role plausibly allows for more than one role-player, while other roles might allow for only a single role-player.

The Senate, a corporate team, and other organized groups are not just organizational structures. The structure of a group captures its functional organization, but an organized group itself is an entity that *is* structured or that *has* a structure. While a precise metaphysical account of organized groups is not needed for our purposes here, being clear that organized groups are not identical to structures is important for several reasons. First, the same organizational structure might be had by more than one group. For instance, there might be many corporate teams or many committees with roles defined in the same way. If groups were identical to structures, there could

⁴ The conception of roles relied on here bears similarities to the way Koslicki defines places in structures. She states that structures are “entities which make available positions or places for other objects to occupy, provided that these occupants satisfy the type restrictions imposed by the structure on the positions in question” and which impose on the objects “a particular configuration or arrangement” (2008, 235-6). Here I focus on more than just configuration or arrangement, but the focus on relations and type restrictions fits with part of the picture I am offering.

not be two groups with the same structure. Groups should not be identified merely in virtue of having the same organizational structure. Members of a group matter for group identity conditions.

Second, organizational structures themselves might be abstract entities. Many philosophers take relations to be abstracta. If organizational structures are complexes of relations they would also be abstract entities. Yet, groups seem to be causally relevant and have spatial and temporal locations; they do not seem to be abstracta. Groups are, therefore, not just organizational structures. They are entities that have structures or that are structured.

An organized group comes to be when some people play the roles required by a group organizational structure. When some person, *a*, stops playing a role *r* and another person, *b*, starts playing *r* the group's membership changes. While *a* was a member of the group, she no longer is. More formally, we can define membership in an organized group as follows:

Organized Social Group Membership: Some things, X, are the members of a group with an organizational structure S at time t and world w if, and only if, together X play the roles in S (i.e., X are related or normatively bound in the ways required by S).⁵

The definition of membership allows for organized groups to change members and captures how membership relates to a group's organizational structure.⁶

§II. Collective Intentionality and Cooperation

Analyses of cooperation have been addressed largely in the context of collective intentionality. One might think the explanation as to why discussions of cooperation have focused on collective intentions is obvious—cooperation involves collectively intending to ϕ with others. So, the thought goes, an account of collective intention will also be an account of cooperation. I will ultimately challenge the idea that all cooperation requires collective intentionality. To see why, we first need to consider how collective intentionality is understood.

Definitions of cooperation usually involve “a set of individual intentions that have certain origins and enjoy certain relations, all of which is common knowledge (that is, public or transparent)

⁵ This definition is as in Ritchie (forthcoming), but departs from those in Ritchie (2013, 2015). Note that the predicate ‘together X play the roles in S’ is a collective predicate. That is, it applies to some things X without applying to each of the individuals (contrast with ‘is tall’ in ‘the players are tall’, compare to ‘surround the building’).

⁶ Organized groups might also persist through changes in their organizational structure. Here I will not develop a theory of *how much* structural change a group can sustain.

among agents” (Paternotte 2014, 47). While sharing a common core, views vary according to whether they involve primitive irreducible *we*-intentions, the ways they take individual intentions and plans to be involved, and whether plural subjects are relied upon. They also vary in being causal or normative explanations.

Bratman, Gilbert, and Searle offer three prominent accounts of collective intentionality and cooperation.⁷ Bratman (1992, 1999, 2014) argues for a reductive account of what he calls ‘shared cooperative activity’. He analyzes “we intend to J” in terms of individual attitudes as follows:

We intend to J if and only if

1. (a) I intend that we J and (b) you intend that we J.
2. I intend that we J in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b; you intend that we J in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b.
3. 1 and 2 are common knowledge between us. (1999: 121)

The account reduces collective intentionality to individual intentions with subplans that are consistent and can be jointly satisfied (i.e., that “mesh”) and common knowledge of one another’s intentions.

Gilbert (1989, 2006) argues for an account based on joint commitments and plural subjects. She takes a plural subject to be formed via a joint commitment to do something as a body. Joint commitments involve obligations to the others with whom one formed the commitment to carry out the activity to which they committed. They also give one the right to rebuke other parties in the joint commitment if they fail to act in a way that is appropriate to their goal obtaining. Gilbert holds that forming a joint commitment requires all of the parties bound by it be involved in its creation (2006: 135, 138, 168). She states “[t]he parties to a joint commitment ... are those who comprise both its creator and its subject” (135). Going further she states that it is a “condition that all of the parties must be involved in the creation of a joint commitment” (ibid.).⁸ Further, prior to forming a joint commitment all involved must express “readiness for joint commitment” which must be common knowledge among them (2006: 138). In her final analysis Gilbert states that “[t]wo or more people are acting together (doing something together) if and only if: (1) they are jointly committed to espousing as a body the appropriate goal; (2) they are fulfilling the behavioral conditions

⁷ For an account relying on a so-called *we*-mode see Tuomela (2006). Miller (2001) gives an account of collective ends or goals that does not rely on collective intentions.

⁸ There are, however, points at which Gilbert appears to want to walk back from this strict requirement. For instance, she states that there are cases “where one or more people sign on to a joint commitment originally created by others” (2006: 134). Given the role creation is given in her account of plural subjects and joint commitments, the theory would need to be modified to allow for joint commitments to bind those who are not co-creators.

associated with the achievement of that goal; (3) their satisfaction of these conditions is motivated in each case by the existence of the joint commitment” (2006: 146). Gilbert’s account is non-reductive, as committing to do something as a body does not reduce to individual commitments or beliefs.

Finally, Searle (1990, 1995, 2010) argues for a view of collective intentionality that relies on primitive irreducible *we*-intentions, but which does not require a plural/group subject or meshing sub-plans. He argues that when engaged in cooperative collective action participants each have intentions of the form ‘we intend that we ϕ ’. According to Searle this involves a special collective or *we*-mode of intending that is not reducible to mere individual intentions in an *I*-mode. Searle’s account also involves common knowledge. He states that “[c]ooperation *implies* the existence of common knowledge or common belief,” but common knowledge “together with individual intentions to achieve a common goal is not by itself sufficient for cooperation” (2010, 49). Rather, he holds that “[c]ooperation requires the collective intention to cooperate,” which is an intention that *we* cooperate or that ‘we intend that we ϕ ’ (2010, 58).⁹

It is worth noting that the core cases on which Bratman, Gilbert, and Searle focus involve small groups of individuals. Gilbert states that “[i]n the example of acting together on which I focus there is no hierarchy, the people involved certainly know of one another and are in the midst of a relatively substantial personal interaction” (2006: 99).¹⁰ Similarly, after mentioning basketball teams and orchestras, Bratman states that “to keep things simple” he will focus only on “shared cooperative activities that involve only a pair of participating agents and are not the activities of complex institutions with structures of authority” (1992: 327). He suggests that we can “gain some insight” into shared intentionality by starting small and that “[p]erhaps our theory of small-scale shared agency can, with due adjustment and further additions, be extended to ... larger social organizations” (2014: 8). Starting with simple cases can be a good methodological strategy. Simple cases are easier to analyze; they have fewer “moving parts”, if you will. If one is concerned with understanding how cooperative or collective behavior comes to be, simple cases rather than highly complex institutional cases also seem like a more fruitful place to begin inquiry. However, even if

⁹ Searle also discusses what he calls ‘collective recognition’, a notion that is weaker than cooperation. He says that “collective recognition need not be a form of cooperation and thus does not require a collective intention to cooperate” (2010, 58). Here I am concerned with considering whether cooperation should require collective intentionality in the way theorists have often argued. So, I focus on Searle’s view of cooperation rather than collective recognition.

¹⁰ Although, Gilbert (2006) takes it to be possible for there to be large plural subjects which are impersonal (i.e., in which not all members know one another personally), anonymous (i.e., when members do not know that particular individuals exist at all), and hierarchical.

starting with simple cases is justified, it would be hasty to conclude that the same analysis of simple cases can always be scaled up to larger organized groups.

Groups like teams, committees, and legislative bodies are very different from two people going for a walk together or three people cooperating while preparing a meal. They involve organizations that have defined often hierarchical roles. Organized groups can vary in members across times and worlds and may be extremely large. They seem different from Gilbert's plural subjects which are formed whenever some individuals form a joint commitment to do something as a body. Moreover, individuals who are currently members of a group might not have been involved in forming the group, its organizational structure, or its goals. Recall that Gilbert requires that all parties bound by a joint commitment must be involved in creating it. In passing several times she states that individuals might join joint commitments, but her theory as espoused in (2006) would need to be modified to allow for this. In considering Gilbert from here, I will be considering her account as stated. On that account, members of many organized groups do not form plural subjects. Yet, she requires plural subjects in her account of acting together.

Members of some groups—like a large legislative body or a corporate team—might not know every other members. They also might not know which individuals are playing which roles or whether there are multiple individuals playing some roles. If a member, *a*, does not know that *b* is a member of the group, it is implausible to suppose that *a* has representations of *b*'s intentions, beliefs, or plans. Knowing that there are particular roles being played is not enough to enable one to have representations of mental states. After all, roles don't have mental states! Further, members of an organized group might not know what many roles involve. They might not even know what roles the group organization includes.

Without knowing who is playing a role the accounts given by Bratman and Searle appear not to apply. On Bratman's account individuals are required to have knowledge of others' intentions in order to engage in cooperative shared activity. Having a *we*-intention, as required in Searle's account of cooperation, might not be possible without knowing who is included in the group or plurality picked out by *we*. Even the condition of common knowledge or belief is challenged by the epistemic opacity about individual identities that can occur in large organizations. For instance, if *p* being common knowledge among a group *G* requires that each member of *G* have representations of

others knowing that p (as well as iterations of higher-order attitudes), then knowing who the members of G are is required for common knowledge.¹¹

Members of organized groups certainly seem to be able to cooperate. Members of teams, committees, and legislative bodies do complete goals, often in ways that are more efficient and effective than an individual or unorganized group of individuals could. Any account of cooperative activity that is meant to be general, must be able to account for such cases. The accounts just considered cannot, at least in their current forms, correctly accommodate *all* cases of cooperation. Moreover, it does not appear that they can be easily modified to handle the cases I focus on here. This is particularly so for Bratman and Searle's accounts. The methodological assumption of starting with small group one-off interactions and scaling up to large group long-term interaction should be abandoned.

In the next section I argue for a form of minimal cooperation that relies on roles in a group organizational structure. Collective intentionality as analyzed by Bratman, Gilbert, and Searle is not necessary for minimal cooperation among members of organized groups. Before turning to what I take is required for minimal cooperation in organized groups, what I am not arguing for needs to be clearly specified.

My aim is not to argue that the accounts of collective intentionality and cooperation just considered fail. One of the accounts, or another in the same spirit, might be correct for small groups working together on a specific joint action. For instance, a case of two people organizing a closet together might require symmetric attitudes involving representations of others' mental states. Moreover, Bratman and Gilbert explicitly state that these are their primary target cases.¹² So, my argument is not against their accounts given their particular target phenomena.

I am also not arguing that members of organized groups *never* cooperate in ways that should be analyzed in terms of joint commitments, representations of others' intentions, or *we*-intentions. Rather, I argue that there are cases of minimal cooperation among members of organized groups that do not involve collective intentions as spelled out in any of the accounts just considered. The accounts considered above fail to give necessary conditions for minimal cooperation. A theory

¹¹ Not all definitions of common knowledge require that agents know who all of the members of the set of group are. Gilbert (2006) argues that 'population common knowledge' can hold even in cases in which group members do not know one another, but do know that there is a relevant population. So Gilbert's account does not succumb in this way. However, given that she does spell out joint commitments and plural subjects in a way that requires creation by all those committed, her account still falters. See also Paternotte (2014) for discussion of common knowledge in conditions of anonymity and further citations.

¹² Although Gilbert (2006, especially Ch. 8) argues that the account can be extended to large social groups like societies. I return to her account below.

relying on collective intentionality might be correct for cases involving extensive or what Paternotte calls “egregious” cooperation, but minimal cooperation among individuals playing roles in group organizations does not place such stringent mental requirements on agents.

Finally, I am not claiming that the creation of organized groups or organized group structures is independent of collective intentionality. It might be the case that an account of collective intentionality like those just considered is often part of the explanation for how organized groups or a group structures *come to be*. For instance, some individuals might form a *we*-intention to create a certain group or a group with particular defined roles. Tomasello suggests that cooperative interaction in humans requires that “participants coordinate their roles—their plans and sub-plans of action, including helping the other in her role as needed—which are interdependent” (2009: 61). These minimal roles that are formed in one-off joint actions are plausibly the precursor to the reified institutionalized roles that humans create in organized group structures. I am not arguing for a view of how organized groups or group structures come about. As far as I am concerned here, it is an open question whether a group coming to be requires collective intentions. That is a question in social ontology, rather than a question about the conditions required for cooperation. Instead I am arguing that even if complicated symmetric attitudes, the formation of a joint commitment, or special *we*-intentions are needed to bring about an organizational structure or a group, these are not necessary for members of organized groups to minimally cooperate.

To put the point another way, consider diachronic and synchronic accounts of cooperation in organized groups. It might be true that a complex representational account with symmetric attitudes is required to create an organized group with a particular structure composed of roles and relations. A diachronic account may rely on a view of cooperation like one discussed above. Yet, once an organized group with various defined roles exists, minimal cooperation need not involve anything as mentally complex as that posited by Bratman, Gilbert, or Searle. A synchronic account of minimal cooperation among members of an organized group, might be quite different. The organization with defined roles allows for group members to minimally cooperate by playing particular roles.¹³ Once a group is “built” the ladder of collective intentionality that was used in

¹³ Tollefsen suggests something similar about group agents. She states that “the performance of joint actions on the basis of group ends, shared intentions, joint commitments, or *we*-intentions might very well be the way in which corporate agents form and sustain their agency over time... group ends, joint commitments, shared intentional activity, and *we*-intentions might all be part of what is happening internally within corporate groups, and this produces a pattern of group behavior that exhibits unified agency” (2015: 47). Here I am focused on members of a group cooperating, rather than on whether a group is an agent, nevertheless there are similarities in my argument and Tollefsen’s remarks.

constructing it can be “knocked away”.¹⁴ Once an organizational structure is in place, complex mental work involving mutual knowledge, representations of others’ mental states, and formation of joint commitments are not necessary for cooperation.

§III. (Minimal) Cooperation + Roles

In §I I argued that organized groups are structured wholes. They have structures with roles that are defined in ways that depend on relations to one another. The interdependence of roles relates to the powers, obligations, and rights role-players have. It also relates to the way actions undertaken by one role-player might be part of larger actions that involve other role-players. Organizational structures with defined roles allow for a division of labor and also for a division of knowledge. No one member of a team might understand the entire range of activities and roles involved. Organizational structures can allow many individuals together to carry out complex actions that no one member could fully understand or complete.

Let’s consider an example. Suppose a consulting firm is tasked with determining whether a merger between Company A and Company B would benefit Company A. In order to determine what to recommend, thereby meeting their goal, the firm puts together a consulting team with various roles. Suppose further that the team has many members who are located across multiple offices and that many team members do not know of one another. Roles involve responsibilities and obligations that normatively bind role-players in various ways. The consulting team includes roles that require role-players to research similar past mergers and pass findings on to members who will include them in a report. Other members have roles that require analyzing a merger’s impact on stockholder and customer perception. They too report their findings to members tasked with writing a final report. And so on. Through many members playing their assigned roles—that is, carrying out tasks and interacting in ways team roles require—the team concludes that Company A should merge with Company B as it will benefit Company A to do so.

In the example, some individuals might have been working closely with others in ways that fit with accounts like Bratman’s, Gilbert’s, and Searle’s. Yet, other members never directly interact and have no knowledge of one another. The team had the goal of determining whether a merger should be recommended. Every member, let’s suppose, knows that she is part of a team that has

¹⁴ The argument made here bears similarities with Bar-On’s (1995) reconstrual of Gricean intention-based semantics. She argues that problems for the Gricean can be avoided if speaker intentions aren’t needed to fix meaning *now* (i.e., intentions are not needed for a synchronic account of nonnatural meaning that has been conventionalized), even if meanings *were* fixed by intentions in the past (i.e., speaker intentions are part of a diachronic account of nonnatural meaning). She also uses the metaphor of “kicking away the Gricean ladder”.

said goal. Each intends to act in a way defined by her role to meet the goal. The team met its goal in a way that involved collaboration and, it would be natural to say, cooperation. Yet, the case does not involve cooperation on any of three accounts considered in the last section. Let's see why.

Bratman requires that each of the individuals participating in shared cooperative activity know that each other individual with whom they are cooperating has an intention of the form 'I intend that we J'. Suppose that Kai and Simona are two members of the consulting team. Neither knows the other personally and neither knows that the other is playing a particular role in the team. Kai does not know that Simona has an intention of the form 'I intend that we determine whether a merger between Company A and Company B should be recommended'. Kai does not even know that Simona exists, so she knows nothing about her mental states! If Bratman's account had to be met for some individuals to be engaged in cooperative action, the members of the team would not count as cooperating even in a minimal sense. The account fails to capture that this is a case of cooperation.

Gilbert requires that two or more people cooperatively carrying out an action together are motivated by the existence of a joint commitment (this is condition (3) in her analysis quoted above). She also holds the joint commitments must be created by all of the parties bound by the commitment.¹⁵ In the case of the consulting team, both conditions could fail. For instance, Simona might not be motivated by the joint commitment to meet a certain goal, but rather to a commitment to play the role that she has been assigned. She might want to play that assigned role because she cares about the firm, others' perceptions of her, keeping her job, or for any number of other reasons. She still seems to be at least minimally cooperating, even if she is not *motivated* by a joint commitment. This alone is sufficient to show that Gilbert's account would not classify the team members as cooperating. Further Simona might not have contributed to forming the commitment to determine whether Company A should pursue a merger with Company B. In fact, *no member* of the team needs to have been involved in setting this as the team's end. The commitment to that end might have been set by those higher up in the larger hierarchical organization of the consulting firm. Given that the team members are cooperating, meeting Gilbert's account cannot be necessary for minimal cooperation.

¹⁵ As noted above, Gilbert appears to want to allow for cases in which to join a joint commitment that they did not create. However, her view as espoused in Gilbert (2006) carries an explicit requirement that joint commitments require co-creation by all those they bind. If her account is going to handle cases like those involving organized groups considered here, it needs to be modified so that not all parties of a joint commitment are required to be creators.

Searle requires *we*-intentions and common knowledge for cooperation. The case under consideration can be augmented to explicitly involve members of the team who lack *we*-intentions but who are, at least in a minimal sense, cooperating with others. Suppose that Layla is another member of the team. Layla has no thoughts in an irreducible *we*-mode, but Layla does intend to play her role as well as possible. She understands that playing her role involves others playing their roles. For instance, she knows that part of playing her role is to pass information on to others on the team and that others cannot successfully play their roles without her playing her role. Yet, she never has a *we*-intention. She thinks about what she is doing and about the goal she's taking part in, but lacks *we*-intentions. By amending the case in this way, it does not seem that *we*-intentions have been smuggled in. And if they have been an argument to that effect is needed.¹⁶ Moreover, it seems that Layla is cooperating with others on the team to provide information for a report and on determining whether to advise Company A to pursue a merger. This provides evidence that Searle's account also fails to establish necessary conditions for cooperation.

The corporate team case involves roles that *require* individuals playing them to cooperate, at least in a minimal sense. Playing a role places normative demands on how role-players are to interact and what they are supposed to do. If roles are played, role-players will interact in ways that aim at a shared end in what appears to be a case of cooperative activity. This is so even if members of the team do not know of others and do not have representations of others' mental states. The accounts of collective intentionality fail to categorize the behavior of the members of the team as cooperative; a different account of minimal cooperation is needed.

An account of the cooperation in the cases just considered needs to take seriously the relevance of roles in an organizational structure. I propose that there is a minimal form of cooperation in which some individuals cooperate in ϕ -ing just in case they play roles in an organized group structure that work towards a common goal. Roles in a group structure function together to allow the larger group to achieve ends or goals. They are defined in ways that are interdependent, that are normatively binding, and that require interaction to achieve goals. Minimal cooperation is achieved when playing roles is combined with a common goal.

The minimal sort of cooperation I propose here is not as robust as cooperative activity as theorized by Bratman, Gilbert, Searle, and others. It need not involve being motivated by a joint commitment, having representations of others' intentions, or having special *we*-intentions. But, I

¹⁶ More generally one might challenge the notion of *we*-intentions as a special primitive irreducible mode of intending. Tollefsen says that *we*-intentions "seem pretty mysterious" and that "[n]othing in our experience and in our everyday practice of ascribing intentions to others...seems to confirm their existence" (2015: 33).

take it that it fits with a common sense understanding of cooperation. In the case studies, individuals are working together by playing roles to achieve a goal. That certainly is, at least, a minimal form of cooperation in the way we ordinarily understand it.

The way roles are institutionalized or reified in group organizational structures allows for more minimal mental requirements on agents engaged in cooperative activity. Members of certain sorts of groups—those with organizational structures—can cooperate in a way that requires less internal mental representation and less in terms of motivational requirements. Individuals must play their roles that aim at a common goal. Playing a role requires standing in the right relations to other group members, completing particular tasks, and having certain rights and obligations. The complex mental demands placed on individuals by accounts that start with pairs of individuals coordinating action are not necessary when individuals play roles in groups with organizational structures. While individuals in groups will have intentions and might have representations of others' mental states, these are not required for minimal cooperation given the existence of defined roles in an organized group structure.

One might worry that the account I have offered is *too* minimal. Let's consider two cases that might be used to press the worry. First, consider a spy network. The individuals in the network might not know anyone else in the network. Moreover, they might not understand what their roles are, what end they are helping to work towards, or how they are contributing to that end. Nevertheless, the combined efforts of those in the network might fulfill the end. Are they cooperating? Intuitions about such cases are mixed. The account I have given could be adapted to accord with both stances on the case by expanding on what is required for a common goal. If one takes the spies *not* to be minimally cooperating, the account could be adapted to require some mental requirements on having a common or shared goal. For instance, one might argue that minimal cooperation in ϕ -ing requires playing roles in an organized group structure that they work towards a goal that all know. In the spy network case, not all know the goal, so the case does not involve minimal cooperation. In contrast, if one takes the spies in the network to be minimally cooperating, one could take having a common goal to require less. It might require that roles are functionally integrated to achieve an end. Or, one might require just that someone with authority over the group know the goal for it to be a common goal.

Second, consider the parts of a car engine. Each part plays a role so that together they function to convert chemical energy into mechanical energy in order to move the car. They do not, however, seem to be cooperating even in a minimal sense. Just functioning in concord is seemingly not enough for cooperation. While the account of minimal cooperation I have given emphasizes

roles and function, it does not take the parts of a car engine to be cooperating. To see why, note that I argued that there is a form of minimal cooperation that involves playing roles in *an organized group structure*. Roles in group organizational structures involve norms. Role-players are obligated to one another and have responsibilities to complete tasks, give directives, and report to others. The organizational structure of a car engine is defined in terms of ways parts operate, but not in terms of norms. The members of the corporate team are correctly classified as minimally cooperating while the parts of the car engine are not taken to be minimally cooperating. While more could be said to develop a theory of minimal cooperation among members of organized groups in greater detail, I will leave the account as sketched here. In the next section I briefly consider how the account fits into the larger project of theorizing social phenomena and social ontology.

§IV. Internalism, Externalism, and Continuity in Social Theorizing

Accounts of social phenomena, social entities, and social facts have largely centered on mental representation of some form or other. Moreover, as we saw there has also been a common assumption that theorizing about the social can start small and then be scaled up. Torrenco argues that many hold a **continuity thesis** he spells out as follows: “[s]ocial phenomena are determined by collective intentions and their shared contents in small groups as much as (and in the same ways as) in complex social situations” (2017: 70).¹⁷ The preceding discussion could serve as a challenge to certain views that center on mental representations, mental states, or mental content. Whether it challenges the continuity thesis depends on how small scale cases should be understood. To clarify the discussion, I return to the terminology introduced earlier.

As I use the terminology internalist views of the social are views that take social phenomena or entities to be grounded, constituted, constructed ... by mental representation, mental states, or mental contents. On my usage externalist views of the social are those that take social phenomena or entities to be grounded, constituted, constructed... by features that are external to mental states and mental contents, like documents, laws, or material resources.

As we saw above prominent accounts of collective intentionality and cooperation rely on representations of others’ intentions, special ways of intending, and common knowledge. Such views are, at least to a large extent, internalist. Views that are more internalist than externalist are

¹⁷ Note that Torrenco argues that the continuity thesis is false.

also common in social ontology. Collective attitudes, intentions, and representations of a kind itself are often taken to constitute or ground social entities and facts.¹⁸

Other views in social ontology make mental states, contents, and representations less central.¹⁹ Thomasson argues that some social kinds might not depend on representations of themselves. She states social entities like “racism, economic recessions, class systems, and gender-biased power structures are typically not intentionally created (either directly or indirectly) by accepting constitutive rules about entities of that kind” (2009: 549). They are, instead, “byproducts of more basic social and institutional facts” (ibid.). A general view of social ontology should not be so internalist so as to *necessitate* that representations of social entities are operative in the construction of the social world.

The argument I gave above goes further. It is not just that minimal cooperation does not require a representational state about *being cooperative*. It is the stronger general claim that collective intentions, recognition, acceptance, and joint commitments are not necessary for minimal cooperation at least when individuals are members of an organized group. The view is more externalist than a view that only goes so far as to reject that a social kind K relies on representations of K.

The argument I gave is closer to views that give accounts of social entities in terms of non-mental features. For instance Epstein (2015, 2017) argues that facts about social entities can depend on facts that do not involve mental states, intentions, or even people. Torrenco (2017) recently argued for an externalist view of social entities. On his view institutional facts and entities are not always grounded in collective belief, acceptance, or intentions, rather institutions and institutional statutes, duties, and rights “are characterized by [and grounded in] the tendency to defer to elements that are external to the content of collective intentions—such as laws, declarations, and contracts” (2017: 67). I agree that not all social entities and not all social facts are characterized by or

¹⁸ For instance, Searle (2010) argues for a social ontology that relies on collective acceptance of constitutive rules. Hacking (1999) argues for a view on which mental representations of “interactive kinds” lead individuals to “modify their behavior” when they become aware of being classified in a particular way (e.g., as a woman, as a refugee, as a Black person) they may “modify their behavior accordingly” (1999: 32). Mallon recently defined social construction projects as accounts that explain “by appeal to our practices of representing” (2016: 1). He argues for an account of social kinds or roles “that are structured by the representations of human categories and, over time, by the causal effects of such representations” (2016: 210).

¹⁹ Khalidi argues along the same lines holding that the existence of some social kinds fail to require that attitudes “be directed towards the kind itself” (2015: 104). See Ruben (1985) for arguments that certain social phenomena can exist without being noticed (e.g., exploitation). See Guala (2010) for arguments that we might lack knowledge of social kinds, even if they depend on collective intentions.

grounded in ways that a largely internalist view of the social requires. Moreover, the argument I have given shows that externalist views of the social can be extended beyond social entities to at least one social phenomenon.

Whether the preceding argument supports a rejection of the continuity thesis depends on what is required for cooperation in pairs or small groups. Here I have argued that members of organized groups can cooperate without shared mental states, *we*-intentions, or joint commitments. I have not, however, argued for a view of what is required for cooperation among individuals who are not members of organized groups. So, the status of the continuity thesis is left open and two possibilities arise. If an account like Bratman's, Gilbert's, or Searle's is correct for cooperation at a small scale, then the preceding argument does require a rejection of the continuity thesis. A more internalist account relying on collective intentions would explain cooperation in pairs or small groups. A more externalist account that relies on playing roles in an organizational structure would explain other cases of (minimal) cooperation. This would be a rejection of the continuity thesis and a methodology that seeks a continuous and single theory for cooperation in all cases.

Alternatively, one might take the preceding discussion to provide the beginning of a framework for a theory of cooperation that upholds the continuity thesis. On this view, cooperation would always involve playing roles in some structure or other—whether it is the structure of a legislative body, that of a family unit, or that involved when two people attempt to move a piano together. Such a view would be very different from the sorts of views that have dominated the literature on cooperation and collective intentionality. It would uphold the continuity thesis at least to the extent that roles would be relied on in all cases of cooperation. It is plausible, however, that any viable version of a role-based view of cooperation would also posit some discontinuity in the mental requirements for cooperation in small-scale one-off interaction and for large-scale long term cooperation in organized groups.

§V. Concluding Remarks

I have argued that there is a form of minimal cooperation among individuals who are members of organized social groups that does not depend on shared mental states, special *we*-mode intentions, or joint commitments. Roles in group organizations can be defined so that playing a role and having a common goal involves cooperating with others. The account I have given is not a general account of cooperation. In its present form it applies only to cases in which individuals play roles in an organized group structure (i.e., to members of organized groups). Yet, at least in these cases, externalism about the social holds for phenomena like (minimal) cooperation.

Social phenomena and social ontology have largely been understood in internalist terms. The roles that representations and mental states play in the creation of the social world are interesting and need to be better understood. Yet, they should not be overemphasized. As organizations become complex and institutions are reified, representations and mental states matter less. Even if internalist mechanisms are needed to build some foundational aspects of the social world, the internalist ladder can be “knocked away” once complex roles and institutions have been created. As we extend our focus beyond pairs of individuals or small groups to complex institutionalized groups and contexts, the social—even social phenomena like cooperation—can and should be theorized in ways that are more external.

Works Cited

- Bar-On, Dorit. 1995. “Meaning reconstructed: Grice and the naturalizing of semantics.” *Pacific Philosophical Quarterly* 76 (2):83-116.
- Bratman, Michael E. 2014. *Shared Agency: A Planning Theory of Acting Together*. New York, NY: Oxford University Press.
- . 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press.
- . 1992. “Shared Cooperative Activity.” *Philosophical Review* 101 (2):327-341.
- Epstein, Brian. 2015. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. Oxford University Press.
- . forthcoming. “What are social groups? Their metaphysics and how to classify them.” *Synthese*: 1-34.
- Gilbert, Margaret. 2006. *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. Oxford: Oxford University Press.
- . 1989. *On Social Facts*. Routledge.
- Guala, Francesco. 2010. “Infallibilism and Human Kinds.” *Philosophy of the Social Sciences* 40: 244-64.
- Hacking, Ian. 1999. *The Social Construction of What?* Cambridge, MA: Harvard University Press.
- Khalidi, Muhammad Ali. 2015. “Three Kinds of Social Kinds.” *Philosophy and Phenomenological Research* 90(1): 96-112.
- Koslicki, Kathrin. 2008. *The Structure of Objects*. Oxford: Oxford University Press.
- Mallon, Ron. 2016. *The Construction of Human Kinds*. Oxford: Oxford University Press.

- Miller, Seumas. 2001. *Social Action: A Teleological Account*. Cambridge: Cambridge University Press.
- . 2014. “Social Institutions.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2014. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2014/entries/social-institutions/>.
- Paternotte, Cédric. 2014. “Minimal Cooperation.” *Philosophy of the Social Sciences* 44 (1): 45–73.
- Ritchie, Katherine. forthcoming. Social Structures and the Ontology of Social Groups. *Philosophy and Phenomenological Research*.
- . 2013. “What Are Groups?” *Philosophical Studies* 166 (2): 257–72.
- . 2015. “The Metaphysics of Social Groups.” *Philosophy Compass* 10 (5): 310–21.
- Ruben, D-H. 1985. *The Metaphysics of the Social World*. Routledge & Kegan Paul.
- Searle, John. 2010. *Making the Social World: The Structure of Human Civilization*. Oxford: Oxford University Press.
- . 1995. *The Construction of Social Reality*. New York: Free Press.
- . 1990. “Collective Intentions and Actions.” In Philip R. Cohen Jerry Morgan & Martha Pollack (eds.), *Intentions in Communication*. MIT Press. pp. 401-415.
- Thomasson, Amie. “Social Entities,” in *Routledge Companion to Metaphysics*, edited by Robin Le Poidevin, Peter Simons, Andrew McGonigal, Ross Cameron [London: Routledge, 2009]: 545-554
- Tollefsen, Deborah Perron. 2015. *Groups as Agents*. Malden, MA: Polity.
- Tomasello, Michael. 2009. *Why We Cooperate*. Cambridge, Mass: The MIT Press.
- Torrenço, Giuliano. 2017. “Institutional Externalism.” *Philosophy of the Social Sciences* 47 (1): 67–85.
- Tuomela, Raimo. 2006. “Joint intention, we-mode and I-mode.” *Midwest Studies in Philosophy* 30 (1): 35–58.