What is Social Organizing?
(Forthcoming in *Philosophy and Phenomenological Research*)

Megan Hyska
megan.hyska@u.northwestern.edu

Abstract

While scholars of, and participants in, social movements, electoral politics, and organized labor are deeply engaged in contrasting different theories of how political actors *should* organize, little recent philosophical work has asked what social organizing *is*. This paper aims to answer this question in a way that can make sense of typical organizing-related claims and debates. It is intuitive that what social organizing does is bring about some kind of collectivity. However, I argue that the varieties of collectivity most amply theorized by analytic philosophers in recent years, including grouphood and collective intentionality, are not the right kinds to embed in a theory of social organizing. I ultimately argue that the sort of collectivity that organizing characteristically brings about is a special kind of causal complementarity among agents' actions— and that while this can exist alongside grouphood and collective intentionality, it is not the same thing as either. The notion of social organizing that emerges is one that can clarify, without trivializing, a number of pressing contemporary debates about how normal people should conduct themselves as interconnected political actors.

1 Introduction

Scholars of, and participants in, social movements, electoral politics, and organized labor are deeply engaged in contrasting different theories of how political actors *should* organize— vertically or horizontally? Around electoral goals or mass movements? Within or across identity groups? By leveraging the power of an existing base of support or first converting many more people to the cause? However, little philosophical attention has been paid to the question of what social organizing *is*[1]. This paper makes a contribution to answering that question.

It's not that practicing social organizers have never taken up this question themselves—they certainly have. But those organizers who've written about the nature of organizing are of course fundamentally concerned with being practically useful— they are, if sometimes theorists, first and foremost practitioners[2]. The different, more theoretical register in which I am asking the question is such that practitioners' answers, while instructive, are understandably incomplete, and indeed at times confusingly divergent. We see this divergence in for instance views on where precisely to place the distinction between organizing and mobilizing[3] or on the degree to which organizers can or must bring political commitments to their organizing work[4]. These differences don't just concern what makes for *good* organizing, but what counts as organizing at all. Indeed, a divergence among notions of social organizing are of long standing; at the 1907 International Anarchist Congress in Amsterdam, where delegates held forth at length about "the question of organization," the Italian delegate, Errico Malatesta, announced, "I have listened attentively to everything that has been said before me on the problem of organisation and I have the distinct impression that what separates us is the different meaning we give words" [Antonioli, 2007, 32].

How can the project of analysing organizing proceed in the face of this apparent plurality of concepts? My view is that the concepts of organizing we commonly see invoked, while perhaps diverging with respect to the particular tactics they enumerate for the cultivation of collectivity, share a common core, which concerns the fundamental nature of this collectivity itself. It is an articulation of this core notion of collectivity which practitioners' analyses don't provide, and which I set out to offer here. This paper should then be thought of as offering a substantive necessary condition on organizing that is robust across otherwise divergent concepts traveling under that name. The feature that I'll posit is consistent with organizing concepts' inclusion of paradigm cases while also being non-trivial and appropriately directional (more on these conditions in §2 below), and so able to play the substantive roles that folk concepts of organizing are commonly called upon to play in political discourse.

[1] A recent exception is Nunes [2014, 2021].

[2] See Inouye [manuscript], Inouye [2022] for discussion of this dual role of the organizer.

[3] This distinction has a long history in the US organizing tradition, going back to the Congress of Industrial Organizations. The distinction is typically taken to concern the question of leadership cultivation, with some saying that a tactic that that does not prioritize the cultivation of pre-existing "organic leaders" within a workplace or community amounts to mobilizing rather than organizing [e.g. McAlevey, 2016]. Others find the cultivation of a broader set of leaders compatible with organizing and indeed advocate for "group-centered leadership" [see e.g. Ransby, 2003, on Ella Baker]— an ideal that has developed into the notion of "leaderfull" organizing often invoked in recent years in the context of the Movement for Black Lives [Miranda, 2021, Woodly, 2021].

[4] Whereas the approach to organizing popularized by Saul Alinsky [see e.g. Alinsky, 1971] emphasizes that the organizer must be non-ideological— indeed, Phulwani [2016] reads Alinsky as saying that this is what separates the figure of the organizer from that of the statesman— Inouye [2022] has argued that, for figures like Ella Baker, organizing is *essentially* ideological.

If our task is then to articulate the flavor of collectivity at the heart of a notion of social organizing, it is natural to wonder whether the sort of collectivity involved might not be one of those that is already most amply theorized in analytic philosophy: so in §3 we explore the possibility that organizing brings about some flavor of collective practical intentionality or merely grouphood. What I contend though is these sorts of collectivity do not help us articulate a substantive necessary condition of the kind we want. In §4 I introduce a notion of inter-agential causal complementarity, which I argue is the right sort of collectivity to embed in a theory of organizing. We begin in §2 though with some remarks to further motivate the project, and to establish some guiding principles for our inquiry.

2 Motivations and Starting Points

In late 2020, a debate kicked off within a region of the US Left concerning how to advance its healthcare agenda in the incoming Congress. Underlyingly, this debate was about the role of organizing in politics. One point of view, championed by podcaster and YouTuber Jimmy Dore [Dore, 2020] and former comms staffer for the Sanders campaign, Briahna Joy Gray, was that the Congressional Progressive Caucus should immediately force a vote on Medicare For All by withholding votes supporting Nancy Pelosi's speakership unless she committed to holding the vote on the House floor. While there was no indication that the legislation currently had the votes to pass, advocates of this tactic argued that, whether or not it passed, it would force dissenting lawmakers to brave the optics of voting against a healthcare policy that enjoyed significant support among Americans in the middle of the COVID-19 public health crisis [Gray, 2020].

But high profile progressive electeds like Alexandria Ocasio-Cortez, former Sanders staffers, and others on the Left, opposed the tactic. David Sirota, alluding to Dore and Joy Gray's use of social media to advocate for the tactic, noted in his critique that "it is easier to do YouTube rants and get retweets about a high-profile floor vote than it is to get everyone excited about the much harder *organizing* work of getting esoteric HHS waivers or removing a relatively obscure committee chairman who is blocking progress" [2020]. Ben Burgis, agreeing in the pages of *Jacobin*, noted that "Political theater can be a useful educational tool, but it can't be a substitute for the long, slow, and often dismally unsexy work of *organizing* and mobilizing citizens at the grass roots," and that "The Left's goals can't be won with procedural tricks or exhorting individual leaders to fight harder. They have to be won by *organizing* the working class at the base of society" [Burgis, 2020]. And the national leadership of the Democratic Socialists of America released a public statement likewise gently but definitely rejecting the tactic, which stated that "we are socialists, and we are *organizers*, so we know that there are no shortcuts to liberation" [2021].

In short, critics agreed, forcing a floor vote on Medicare for All at that moment would represent an attempt to do politics without *organizing*. Nor was this first time that US intra-Left tactical disputes would be framed in terms of an orientation to the task of organizing. As Jasper Bernes and Joshua Clover noted in 2012, "The exhortation to

organize has often been heard in the dissolution of the various Occupy encampments in the US, from left thinkers as various as Noam Chomsky, Doug Henwood, and Jodi Dean" [Bernes and Clover, 2012]. And Ella Baker, critiquing an emphasis on charismatic leadership in the civil rights movement of the mid-1960s, once complained that:

> [What was needed was] a greater degree of real concentration on organizing people. I keep bringing this up. I'm sorry, but it's part of me. I just don't see anything to be substituted for having people understand their position and understand their potential power and how to use it. This can only be done, as I see it, through the long route, almost, of actually organizing people in small groups and parlaying those into larger groups. [Payne, 1995, 380]

While these proponents of organizing may, for all we've shown here, have slightly different notions of organizing in mind, it is clear that for all of them the concept of social organizing is intended to cover a *particular sort* of political strategy; it's not that just *any* politically motivated action counts as organizing, or else critics couldn't deny that #forcethevote, the expressive horizontalism of Occupy, or the speeches of Stokely Carmichael (later Kwame Ture), all counted. I don't take it to be a desideratum on a view of organizing that all the above critics' claims come out *true*; but I do take it as a desideratum that their claims come out as *coherent*, and not *trivially* false. The concept of organizing that will meet this desideratum will then be more specific than that of political activity in general.

Beyond their role in strategy-discrimination, concepts of social organizing also appear as weight-bearing components of explanation. When Raphael Warnock and Jon Ossoff won both Georgia Senate seats in January 2021, *The American Prospect* ran a story on "the organizing that brought Georgia to the brink" of this Democratic victory [Gibson, 2021]; implicit in this verbiage is the idea that organizing *explained* the victory.

In light of the role in strategic deliberation and explanation that concepts of organizing are asked to play, it makes sense to ask— what *is* organizing? Is there indeed any concept available that could bear the weight of the above uses? Some are skeptical. Bernes and Clover note that, for political actors, "'organize' must in some regard be the right thing to do in so far as it is a term both common-sensical and capacious in its lack of specificity. It risks being what Fredric Jameson calls a 'pseudoconcept': the imperative to 'organize' comes down to, *do that thing that causes you to be more rather than less effective*" [Bernes and Clover, 2012].

If this were truly all that organizing came down to, then any attempt to use the concept to discriminate one strategy from another or to explain a turn of political events, would be spurious. If "organizing" just means "whatever is effective", then documentation of organizing's absence could not offer a *reason* that forcing the vote, or the horizontal formations of the Occupy moment, were to be rejected— rather, this would merely be a restatement of the speaker's view that they ought to be. And indeed, insofar as advocates of these derogated tactics favored them because they thought they *were* more politically effective, they might well reply that they *were* organizing.

Likewise, in explanatory contexts, if organizing is simply any effective political action, then to say that a policy change, an uprising, or an electoral result was won by organizing is to say that x happened because events that disposed x to take place preceded it— which is evidently explanatorily vacuous.

Note also that the deflationary notion of organizing proposed by Bernes and Clover will cast antipathy toward organizing as incoherent— why, after all, would anyone resist doing what was most effective? However, even comprehensive hostility toward organizing is not unheard of in the history of political thought. Some individualist anarchists have held positions approaching this, with perhaps the most radical articulation coming from William Godwin, who avowed that "every thing that is usually understood by the term cooperation, is in some degree an evil" [Godwin, 1842][5].

What we are seeing then is that for a notion of organizing to play the roles that it is apparently called on to play, any connection between organizing and the effective accomplishment of a goal can't be definitional. Organizing must refer to a type of action such that it is a substantive claim that this type of action is likely to be the most effective in all, or in any given, political project(s).

To summarize, in order for an account of organizing not to confirm Bernes and Clover's deflationary suspicions, it is first necessary that it avoid several varieties of triviality, which I combine in the following desideratum:

> TRIVIALITY AVOIDANCE:[6] it isn't the case that:
>
> a) by definition, organizing is more politically effective than not organizing, OR
>
> b) every, or virtually every action, counts as organizing.

I have framed this desideratum, and indeed this entire paper, in terms of the event-type *organizing*, rather than the gradeable state-type, *organization*, because I take it that a full account of the latter will not exhaust the questions we might ask about the former. One question that arises for the event-type is whether it is always an action– i.e. even if *organization* can sometimes arise by fluke, must *organizing* always be an intentional attempt to increase organization? And if so, intentional under what description? Our modest goal of offering a substantive necessary condition on organizing doesn't answer all these questions. But I do stake out this set of assumptions about the relation between organizing and organization:

> DIRECTIONALITY: Organization is a gradeable state-type. Organizing is an action that increases the grade of this state-type that obtains.

As promised above, our project here will be to offer a necessary condition on organizing, but moreover one that is substantive: a vacuous condition like "exists in the universe" would, after all, probably be a necessary

[5] Of course, many anarchist thinkers do not agree. See the discussion of "the organizing question" in Antonioli [2007].

[6] Throughout this paper, I use SMALL CAPS formatting to mark desiderata for, or challenges to, a satisfactory account of organizing. I then summarize all these items in §4.3

condition on an event's being one of organizing because it would admit of no counterexample. However, it would do nothing to illuminate the phenomenon at hand. I've introduced the desiderata of TRIVIALITY AVOIDANCE and DIRECTIONALITY because I take them to be principled measures of substantiveness.

Before we get started crafting proposals that meet these conditions, it will be useful to consider some paradigm cases of organizing, as well as the form that organizing claims often take. Below then are two vignettes accompanied by some typical claims about the organizing that's going on within them:

Case 1: Maya and her coworkers want to see higher wages and improvements to workplace safety. Maya starts a card check campaign to help form a union.

a. Maya organizes her coworkers.

b. Maya organizes a union.

c. Maya organizes (for) improved workplace safety.

Case 2: Amir is upset about recent police sweeps targeting unhoused people who live under his neighborhood's viaducts. He reaches out to sympathetic groups, who offer to provide resources like safety marshals, medics, speakers, and the mobilization of their respective memberships, and he advertises a meeting point at a city park.

a. Amir organizes the sympathetic groups.

b. Amir organizes a coalition.

c. Amir organizes a protest.

d. Amir organizes (for) an end to the sweeps

Statements about token organizing events, like the above, provide information about the specifics of the token in part via the grammatical object taken by the verb "organize" (i.e. "her coworkers," "a union" etc.). It is important to note however that the grammatical objects of these organizing statements don't all fill the same logical role[7]. In 1a, for instance, Maya's coworkers are what organizing is performed upon; they are the *operands* of the action. In contrast, 1b is most readily read to mean that this complex social entity, the union, is the *output* of Maya's organizing. And "improved workplace safety" (1c) is the *goal* that motivates the construction of that complex social entity— the thing that organizing takes place for the sake of. While "the sympathetic groups" of 2a differ from the individuals of 1a in that groups are already complex whereas individuals are not, in the context given, "the sympathetic groups" is nonetheless most naturally read as specifying the operands of Amir's organizing. Likewise,

[7] A piece of evidence that this is so is that one can compose a sentence that takes all these objects at once: e.g. "Maya organizes her coworkers into a union in order for them to achieve improved workplace safety".

a coalition (2b), while perhaps a strictly different sort of entity than a union (1b), seems here to name the output of Amir's organizing. The grammatical object of 2c and 2d are *both* goals, as was "improved workplace safety" in 1c; in almost all cases of course, organizing in some sense aims at a proximal goal as a means to give rise to others, and I include both 2c and 2d to make evident that instances of organizing can be described in terms of goals at various points along this chain.

We might get started on specifying the nature of organizing in terms of the argument places we've just identified: organizer, operands-system, output-system, and goal(-set). We might even say that the task of giving an account of organizing is that of articulating what sort of relations systematically hold among these elements. Here at the outset, I will make it a framing assumption that organizing is a kind of goal-directed activity where the means adopted for the achievement of that goal involves whatever is going on when the above outputs are produced out of the above operands, i.e. when a union is produced out of individual coworkers, or a coalition is produced out of sympathetic groups. The key question that I will address in the body of this paper below is what the difference between operands and outputs systematically is. That is, if there is a type or degree of collectivity present in the output (e.g. the union) that wasn't there in the operands (e.g. the non-unionized coworkers), what type of collectivity is this? In the next section we will consider approaches on which organizing brings various forms of group intentionality or even grouphood simpliciter. But our claim will be that none of these possibilities will do.

3 Accounts Based on Group Intentionality and Grouphood

The first set of suggestions we consider are based on the idea that the operands and outputs of organizing differ with respect to their capacity for collective practical intentionality. Such suggestions have an intuitive appeal; when Maya's coworkers or the sympathetic groups attending Amir's protest become organized, surely the key thing is that they can *do* things *together* in a way that was previously impossible. The idea that this section explores is that the characteristic way that organizing augments the ability to do things together is not through sheer numbers of complementary actions being taken at once, but by bringing it about that these complementary actions are underlied by a kind of collective intention. We will find, however, that it is difficult to spell out a satisfactory account of organizing in terms of such a suggestion.

Practical collective intentionality, the putative phenomenon of collectives as such forming intentions and taking actions, is now the subject of a robust philosophical literature, and is theorized as coming in multiple forms. Since the late 80s and early 90s, two of the most influential approaches to the idea of collective intentionality have been those of Margaret Gilbert and Michael Bratman. Both thinkers at least sometimes characterize themselves as offering accounts of "shared intention," and they offer as paradigmatic examples of this phenomenon cases where two individuals decide to go for a walk, or move a table together. We might then consider a proposal on which the

output of organizing is a system that bears a shared intention:

> **Organizing**$_1$: Organizing is an activity which aims to achieve some goal by cultivating a shared intention to achieve that goal.

Consistent with the mission for this paper set out in the introduction, Organizing$_1$ ("O1") (and the subsequent similarly labeled proposals discussed below) should be understood as offering a substantive necessary condition on organizing. It does not purport to offer sufficient conditions.

Of course, Gilbert and Bratman's accounts of shared intention differ, and so O1 can be interpreted in either Gilbertian or Bratmanian terms. For Gilbert, a set of people share an intention to ϕ when they've established up-front that they are "jointly committed" to ϕ-ing, where a joint commitment occurs when all parties come to commonly know of one another that they have individually and openly expressed a willingness to be so committed [Gilbert, 1989, 180-199 ; Gilbert, 2008, 502; Gilbert, 2014, 23-36].

For Gilbert, where such an up-front commitment is established, the shared intention can persist even when the individuals who share the intention cease to have individual contributory intentions— that is, even when some or even all of them no longer in fact intend to do their part in bringing it about that the collective do the thing they have a shared intention to do. In contrast, for Bratman, the presence of ongoing individual intentions of roughly this sort is precisely what makes it the case that a collection of people share an intention. On his presentation, *we* share an intention 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.
>
> [Bratman, 1993, 106]

Bearing in mind that O1 can therefore be interpreted in accordance with either a Bratman- or Gilbert-style notion of shared intention in mind, how does O1 fare with respect to our desiderata? Because the conditions on shared intention, whether Bratman's or Gilbert's, are somewhat demanding, it certainly won't be the case that just any old action satisfies O1, and moreover it's not trivial that a system of individuals who share an intention will be more effective than a system of individuals who don't— all of which is to say, O1 does fine with respect to TRIVIALITY AVOIDANCE. The situation with respect to DIRECTIONALITY is a little less clear. If a system simply has a shared intention or it doesn't then O1 makes organization a 2-valued discrete variable. Strictly speaking, this can satisfy DIRECTIONALITY so long as we understand one of these values as in some sense an increase over the other one, and organizing as the action that induces a system's proceeding from the lower to the higher of these values.

However, contrary to the verdicts of O1, systems that satisfy neither Bratman's nor Gilbert's conditions on shared intention can indeed be the output of organizing. There are a range of interesting examples that satisfy only a subset of, say, Bratman's conditions[8], but for brevity's sake I skip to one that meets none of them:

Case 3: The boss of a criminal syndicate wants as few people as possible to have a full picture of her plan to monopolize the drug trade in a certain city. She does confide the plan in her two most trusted capos, but on the proviso that neither can disclose the plan to anyone who doesn't already know. And because the boss is fearful of her underlings teaming up to depose her if they ever develop close working relationships, she does not tell either capo who the other is. The capos assign distinct subplans to each of their soldiers to go execute, but, per the boss's wishes, do not explain to any of the soldiers what the overarching plan is. Finally, in line with the boss's penchant for secrecy, no soldier knows who all the other soldiers are.

This total system, including bosses, capos, and soldiers, is intuitively socially *organized* to some degree (indeed, we call it "organized crime") and I think it's equally intuitive that the boss, in conveying instructions to her two capos, is organiz*ing* the syndicate—specifically, organizing to bring about the syndicate's drug monopoly. But note that the intention to develop this monopoly is not shared by all members of the organization— while everyone is in fact acting in service of the plan to develop a monopoly, and while the boss and trusted capos will have the intention to do so, none of the soldiers can have any such intention, because they don't know that developing a monopoly is the plan. The case is even clearer that none of the soldiers, nor indeed even the capos, can have the intention to establish a drug monopoly *together*, as is required by Bratman's first condition. On Bratman's account, where a set of people share an intention to J, each of the individuals in that set has an intention that the members of the set J together (J-ing together does not itself require the members of the set collectively intending to J together, thus dodging any obvious circularity). As Bratman notes, I can not have such an intention, where I do not "see your playing your role in J-ing as in some way affected by me" [1993, 102]. And I cannot do that where I do not conceive of J-ing as something that you and I both know ourselves to have roles in— that is, where I do not conceive of myself as doing anything *with* you and think that you conceive of things likewise. Because the soldiers in case 3 do not know who all the others are, even if they infer that the action assigned to them is part of some larger scheme, they neither know what the scheme is, nor have any sense of which agents roles in the scheme might be distributed across. While the two trusted capos *do* know what the scheme is, they likewise don't know which agents will likewise see themselves as having a role in the scheme, because their identities haven't been disclosed to one another. This system therefore, for many reasons, does not meet the first of Bratman's conditions, and his latter conditions can only be met where the first is— so the system meets no conditions on Bratmanian shared intention.

[8]For instance, Andrea Sangiovanni [2015] has characterized solidarity as a state satisfying Bratman's first two conditions but not the third. That Case 3 is an instance of organizing while the output system meets none of Bratman's conditions therefore makes the case, in passing, that the sort of collectivity brought about by organizing is also not solidarity, or at least not solidarity as Sangiovanni conceives of it.

It's likewise evident that the members of the syndicate have not openly expressed a willingness to be committed to the project of developing this monopoly (this too would require them knowing both who the others were and what the plan was), and so they do not have a shared intention on the Gilbertian view either. The general issue that case 3 raises is that a system can apparently be the output of organizing even when the agents in the system do not have common knowledge of one another's identities, let alone one another's intentional states vis-a-vis the goal of the organizing intervention. We will sum this up by saying that O1 fails to account for cases of organizing which yield NO SELF-CONCEPTION AS ACTING TOGETHER.

But there are other approaches to collective practical intentionality that are perhaps more applicable to the organizing of crime syndicates, corporations, and other large formal organizations—this is the idea of full-fledged collective agency. We'll thus consider the following:

> **Organizing$_2$**: Organizing is an activity which aims to achieve some goal by way of creating a collective agent which will pursue that goal.

We will of course be particularly interested in interpreting this proposal in accordance with accounts of collective agency which don't themselves require shared intention. Prime among these are proposals that to have a belief or desire is really to have a certain functional or interpretive profile [e.g. List and Pettit, 2011, Huebner, 2014, Tollefson, 2015], where group-level states satisfying these profiles are realizable by many different sets of underlying group-member attitudes, including those which include no other-regarding attitudes[9] like our crime syndicate. On a simple functionalist account, the intentional state-type belief is characterized not in terms of a particular underlying physical state (e.g. a brain state), but in terms of a relationship to certain environmental inputs (e.g. evidence), other internal states (e.g. desires), and outputs (e.g. behavior). If a rival organized crime group challenges the crime syndicate we discussed above and the net effect of the syndicate's members' behavior is to attack the rival group, then on a functionalist view, the internal states that mediate this distributed behavior may constitute the belief that the rival group is threatening the syndicate and the desire that this stop, regardless of the fact that syndicate members lack robust other-regarding attitudes. If full blown agency is a matter of having a suite of intentional states (e.g. environment-responsive beliefs, desires, intentions to act) that fit together in the pattern we associate with rationality, the collection's ability to realize these various state types prima facie renders it a candidate for agency. What many of those who are broadly characterizeable as functionalists in the group agency literature will also emphasise is that the crime syndicate's behavior is effectively explained and predicted by adopting what Dennett [1987] famously called the "intentional stance"— i.e. by trying "to attribute beliefs, intentions, and desires

[9]Note that functionalism per se doesn't guarantee an attribution of agency to entities like our crime syndicate. Some, like List and Pettit [2011, 33] affirm these sorts of cases where you have "organizational designers co-opting others into a a structure underpinning group agency, without making them aware of their agency at the group level" as instances of group agency without joint intention. But on the other hand, Huebner [2014, 19-21] takes it as a principle that we should withhold any attributions of "mentality" or agency from such systems.

to groups in the same way we do to individuals" [Tollefson, 2015, 104]. On views like Tollefson's, the very fact that taking the intentional stance toward groups yields good predictions and explanations is what makes it true that those groups are agents.

Because O2 can be satisfied by organized systems whose constituent agents have no self-conception as acting together, we can see why it might be regarded as an improvement over the overly demanding O1. Nonetheless however, we will see that O2 is inadequate. This is because there are organized systems that fail to be group agents. To see this, consider that even as shared intention is not necessary for group agency, it also isn't generally regarded as sufficient. As Christian List and Phillip Pettit point out, a collection of agents acting with a shared intention,

> may perform only one action together...And even if they perform several actions, they may do so without forming and enacting a single system of belief and desire. Thus there may be no basis for predicting what they will do together in future or for speculating about what they would do under various counterfactual possibilities. In short, there may be no basis for thinking of them as a unified agent. [2011, 34]

Collins [2019, 56] similarly argues that "if we were to call Gilberts' groups 'agents' we would have to have a conception of agency on which agents can be constrained to one goal, without an ability to form new goals, beliefs, decisions, and intentions," and that this is unacceptable. As I understand this line of thought, some systems with shared intentions fall short of agency not simply because they're short-lived—there would be very few agents indeed if immortality were a requirement— or even because they are in fact highly focused on one goal—single-mindedness doesn't exclude agency. Rather, what I take List & Pettit and Collins all to be indicating is that agency requires that a system *could itself* form new intentional states under various circumstances, and that it is the lack of this potential that disqualifies even some systems with shared intentions. While of course the two individual agents who decide to go for a walk together can as individuals decide to commit themselves in such a way that a further shared intention arises— to go for ice cream after the walk, say—the presence of a joint intention per se doesn't set up the infrastructure for the collective as such to form further attitudes. This means that an act which brought about a mere shared intention would not as such count as organizing, according to O2. And this is problematic—while I protested above that the creation of joint intention isn't *necessary* for organizing, I suspect there are many cases where organized systems do have shared intentions that fall short of group agency. Consider Case 2 above, in which Amir got a bunch of sympathetic groups to agree to make various contributions to the holding of a protest. Plausibly what Amir is doing here is simply bringing about a shared intention; we may assume that the sympathetic groups don't consistently work with one another, that there are NO SUSTAINED PATTERNS OF COOPERATION among them, perhaps even that they disagree about many issues beyond the one at hand, so that it's intuitive that they don't constitute a collective agent. Nonetheless, what Amir is doing seems, paradigmatically, like organizing. But O2 will not capture this.

Perhaps it will be claimed that, even if neither collective agency nor shared intention looks to be a necessary condition of organization, and so neither O1 nor O2 plausible necessary conditions on organizing, bringing about collective intentionality *of some kind* is necessary. This suggests a disjunctive proposal like the following:

> **Organizing**$_3$: Organizing is an activity which aims to achieve some goal by way of either creating a group agent which will, or a shared intention to, pursue that goal.

But this is unsuccessful too. Consider:

Case 4: Frieda doesn't know her neighbors well, and has never coordinated with any of them on any project. However, she knows that many of them likely oppose a local condo developer's request to be issued an exemption from the neighborhood's affordable housing quotas. She knocks on her neighbors' doors to let them know where and when the public comment session on the developer's request will be held, in case they would like to attend, and a significant number of them end up attending.

I think Frieda might still be counted as doing some organizing here. But notice a few things: we needn't imagine that she is even actively persuading neighbors to go, let alone securing commitments from them— she's just providing them with information that she thinks they'll use in a certain way. The system composed of Frieda's neighbors doesn't look to have a shared intention—they don't know one another, they aren't acting with each other's actions in mind, and they aren't making any commitments to one another. Still less does it seem that Frieda has made them into a group agent; whatever might happen down the road, there is nothing in place right now, either a formal decision procedure or a common susceptibility to the manipulations of a puppeteer, to suggest that their actions will converge on common effects in this way going forward. Frieda then is organizing without bringing about any sort of collective intentionality.

We've seen then that accounts that require of organizing that it create a shared intention (O1) a group agent (O2), or a disjunction of the two (O3), are all susceptible to counterexample. While the collective intentionality literature is now vast, and there are other related notions that might be plugged into a proposal of the above forms, whether by itself or disjunctively, I contend that the following conclusion will persist: the notion of collective intentionality, in any of its forms, simply will not help us articulate a substantive necessary condition on organizing. My proposal then has something in common with other recent work in social ontology [e.g. Epstein, 2015, Burman, 2023] which advances a skepticism that intentionality can do all the work it is often asked to do in social theorizing.

When we reflect back upon the cases of organizing we've discussed so far, we can see that what they have in common is that their outputs are bigger, more complex social entities apparently built out of the smaller, less complex social entities that are their inputs (e.g. individual workers into a union, petty criminals into a crime syndicate, sympathetic groups into a coalition). And it would be natural to call the complex social entities that

are the outputs here *groups*. Of course, in Case 2, Amir's organizing takes pre-existing groups as the operands as well, with the coalition-group being a perhaps more complex group because it has these already complex parts. All this suggests the following: perhaps organizing characteristically has something to do with the creation or modification simply of groups, where grouphood simpliciter is not nearly as demanding as shared intention or collective agency. I will ultimately find this approach unsatisfying, but note that whereas my claim above was that the notion of collective intentionality couldn't furnish a plausible necessary condition on organizing, and indeed *some* group-based proposals will likewise fail to be plausibly necessary, my ultimate conclusion here will instead be that the notion of grouphood can't furnish a *substantive* necessary condition. It is not possible to articulate even a disjunctive action-type in terms of grouphood that is a) necessary and b) satisfies both TRIVIALITY AVOIDANCE and DIRECTIONALITY, our two principled measures of substantiveness: not the creation of groups, and not any particular sort of modification that makes reference to their properties qua groups. The notion of grouphood simply doesn't illuminate the action of organizing per se.

To work our way to that conclusion, let's start by considering a proposal:

> **Organizing**$_4$: Organizing is an activity which builds a new more complex group as a means by which to achieve some further goal.

To evaluate this characterization, we will of course want to think harder about what groups are, and about what it takes to create them. I think, as a start, it is evident that if we're talking about instances where organizing creates groups that didn't exist before, we're probably talking about groups like clubs, teams, and organizations rather than groups like left-handed people, red-heads, or Asian-Americans. Katherine Ritchie [2015] calls the former Type 1 groups and the latter type 2 groups[10]. Where Type-2 groups are based merely on a common characteristic, what could it mean to "create" a new such group? It could mean bringing it about that some people instantiated a characteristic that had never been instantiated before. There is more to be said about the different ways this might work for constitutively social properties as opposed to non-social ones. But we bypass that conversation here because it is just not clear that our paradigm cases of organizing involve bringing it about that the involved individuals all have some never-before-instantiated property, of any kind[11]. What is a plausible candidate for such a property in the case of Maya's unionization campaign, for instance? I doubt there is one.

Type 2 grouphood excluded, we might explore whether Type 1 grouphood does the trick— and we might be encouraged in this exploration by the fact that Ritchie explicitly distinguishes Type 1 groups as those that are "organized" [2015, 210]. However even Type 1 groups have been theorized in a number of different ways by social

[10]She also refers to them as "social kinds" and "organized groups" respectively—for our purposes though, it would be confusing and question-begging to use this vocabulary. Tollefson [2015, 3] tracks a similar distinction with the terms "corporate groups" and "aggregative groups".

[11]The property "being a member of this particular group" or any variant thereon may well be previously uninstantiated, but pointing to this characteristic here would amount to grounding membership in the group with membership in the group.

ontologists. Our interest here is of course not in determining which of these views is right as an account of Type 1 groups per se, but rather is in considering whether any of the notions of grouphood on offer can help us articulate a satisfying account of organizing. It will be just as well to start by experimenting with one particular view, to see whether plugging it into O4 yields something promising. We'll start then with Katherine Ritchie's [2013, 2015, 2020] structuralist view: on this account, groups are realized structures, where a structure is composed of nodes connected by functional relations. A structure is realized when its nodes are filled by individuals who stand in the functional relations that define that node [Ritchie, 2015, 316]. On this account, a group may persist through changes to both its membership and its structure, both over time and between worlds [Ritchie, 2020, 412].

How does O4 fare with this view of grouphood plugged in? It seems capable of respecting TRIVIALITY AVOIDANCE: it won't be every action that brings about the realization of a more complex functional structure, and moreover it will be a substantive claim that more complex structured groups are better-positioned to effect political goals, so the connection between organizing and effectiveness wouldn't be trivial. Moreover, where organization is equated with complexity, O4 looks compatible with DIRECTIONALITY.

But an apparent counterexample to O4 presents itself: there are instances of organizing that *do* yield a new Ritchie-style group, but not one that is more complex than what existed before. This appears to be the case with what the applied organizing literature dubs STRUCTURE-BASED ORGANIZING [e.g. McAlevey, 2016, 12]— that is, instances of organizing that take existing structures, rather than collections of unconnected individuals, as their operands. When the labor movement targets workers on the shop floor, or when the US civil rights movements targeted individuals in Black churches under Jim Crow, it's not as though the targeted individuals bore no antecedent relations to one another. In an established work place for instance, coworkers already stand in functional relations to one another having to do with the labor they perform for their employers; they often also train one another, and arrange to cover one another's duties or shifts when needed. This set of functional relations constitutes an existing structure on these individuals; a structured group exists here prior to the project of, say, unionizing. Now, the functional structure of a union is distinct from the functional structure of the workplace as such; after the act of organizing then, its seems that there are two different functional structures on the workers and so, on Ritchie's view, two extensionally equivalent but numerically distinct groups. So far so good: the act of organizing has created a new group[12]. However, while the workplace and union structures are distinct, it's not at all clear that the latter is more *complex* than the former, whatever complexity amounts to; it doesn't involve more people, and needn't involve a more intricate variety of functional relations. Cases of STRUCTURE-BASED ORGANIZING— indeed cases like Maya's (Case 1) above— do not satisfy O4. But since STRUCTURE-BASED cases are paradigmatic instances of

[12]It's much more natural to think that the union is the new group here, since the workplace group may be at least qualitatively identical to the one that pre-exists the act of organizing. But if someone wanted to quibble that Ritchie doesn't give (and doesn't purport to give— see Ritchie [2020, 412]) the fully precise persistence conditions on groups that would allow us to be totally clear on which of the post-organizing entities, if either, is the continuation of the pre-organizing group, that wouldn't trouble the point we want to make here: at least one new group has been created, but none which is more complex than the one that existed pre-organizing.

organizing, this means O4 can't be right.

The natural revision is to toss the requirement that the new group be more complex, while leaving the rest:

> **Organizing**$_5$: Organizing is an activity which builds a new group as a means by which to achieve its goal.

But O5 is no better equipped to handle cases of what is commonly called "internal organizing", than O4 was:

Case 5: Jeeyeon discovers that, although her organization has several hundred dues-current members, a mere dozen or so are currently doing the lion's share of the organization's work and becoming burnt out. She sets about recruiting less active members to do a rotation in these leadership roles, which she hopes will allow current leadership to rotate into less demanding roles where they can recover, and will distribute a sense of organizational ownership more widely across members.

In Case 5, Jeeyeon is shuffling which members fill which functional roles, but on a Ritchie-style account, the group survives such shuffling. Here, organizing doesn't yield a new group at all, but modifies an existing one. So it isn't an instance of organizing, per either O4 or O5.

At this point, it is natural to wonder whether the problem is with O4 and O5 themselves, or the particular account of Type 1 grouphood that we've plugged into these conditions: maybe an account of groups other than Ritchie's describes a kind of object which, plugged into O4 and O5, would make them viable. I however am skeptical. An account which conceived of groups as sets, mereological fusions, or aggregates of members would be a nonstarter; one among several reasons is that these entities are individuated by their elements/ parts/ member-components, so that without new members there is no new group. As a result, plugging any of these entity types into O4 or O5 yields the incorrect judgment that no organizing takes place in cases of both structural and internal organizing.

Effingham [2010] offers a "sophisticated set" account on which a group is a set of ordered pairs, in each of which the first element of is a world, and the latter is a further ordered pair containing a time and a set of individuals. Effingham's account, like Ritchie's, articulates a kind of entity that can survive changes in its membership over time and between possible worlds. It also allows that two different groups can be constituted by the same members at a given time, but Effingham's account, unlike Ritchie's, gets this result only so long as these two different groups have different cross-temporal and cross-modal memberships. That is, on Effingham's notion of grouphood, two groups can be distinct only so long as they don't have the same membership at every possible world. But as Uzquiano [2004, 147] points out, this doesn't allow one to distinguish between two *necessarily* coextensive groups. And this is relevant to us because, at least under certain regulatory circumstances, the set of employees in a work place and the set of people in a union, might be necessarily coextensive. This suggests that STRUCTURE-BASED ORGANIZING does not create a new Effingham-style group. Plugging an Effingham-style account of groups into O4 and O5 won't help

these conditions meet the challenge of internal organizing, and will do worse than Ritchie-style groups by failing to account for STRUCTURE-BASED ORGANIZING as well.

If some instances of organizing don't create a new group at all, on any of the going accounts of grouphood, then one determined to see the notion of grouphood appear in a necessary condition on organizing might set aside group *creation* in favor of any number of varieties of group *modification*, or better yet a disjunction of these: maybe organizing sometimes creates a new group but other times instead changes the members in a group, or changes which members fill which nodes, or changes the responsibilities associated with the nodes.

> **Organizing**$_6$: Organizing is an activity which builds or modifies a group as a means by which to achieve its goal.

As foreshadowed above, my objection to this proposal does not take the form of a counterexample to its necessity. It may well be that every instance of organizing does either create or modify a group. Indeed, it might be very hard to come up with a counterexample to this as a necessary condition because virtually all operations on the social world could be counted as at least *modifying* a group, where both the nature of modification and the nature of grouphood are interpreted broadly, as the impulse behind O6 favors. But for this very reason, O6 is not a *substantive* necessary condition. If indeed any operation on the social world alters some group, then O6 can't satisfy TRIVIALITY AVOIDANCE. And certainly O6 cannot satisfy DIRECTIONALITY. Where any modification to a group counts as organizing, then diametrically opposed alterations will both count as instances of organizing: at any given moment, both adding and subtracting a member from the system, and both increasing and decreasing the responsibilities of any given node in the structure, would count as organizing. Organizing wouldn't be an action that induced change in a *particular* direction but in *any* direction. I don't argue that there are never cases where seemingly opposite actions on a system can both in fact serve to move a system in the same direction. But the fact that diametrically opposed alterations to the system will systematically be regarded as equally entitled to be called organizing means that O6 struggle to respect our desideratum of DIRECTIONALITY. Indeed this result will hold regardless of the conception of grouphood that once plugs in. An undesirable upshot of this is that, where all alterations to a system count as organizing it, there is no room to distinguish some actions as *disorganizing*, as surely we should want to.

What we lost when we gave up O4 and then O5 was a specification of a particular direction in which the operands were shifted to become the outputs. While we saw that neither greater complexity nor new grouphood were the right ways to specify this direction, it turns out that we can't just dispense with such specification altogether. And once we find the right specification of direction, I think we will find that it does not make reference to the properties that make an entity a group, in which case grouphood, as a specific notion of collectivity, falls out of the picture.

As with our consideration of O1-O3, the argument here has not been that there are never organized systems

that are groups, but that the notion of grouphood doesn't help us articulate a substantive necessary condition on organizing. The notion of grouphood might of course still be helpful in articulating some of the forms that organization can take, and indeed perhaps in developing a taxonomy of kinds of organizing: for instance, Ritchie's structural notion of groups allows us to nicely describe the sort of reorganizing or internal organizing we saw in case 5 as involving agents moving between nodes in a functional structure. It just doesn't offer us a substantive necessary condition on organizing per se.

We have considered whether organizing can be understood as the creation of groups, group agencies, or joint intentions, and found all such proposals wanting. Here is a perspective on what went wrong with these former proposals: they all supposed that organization was fundamentally a property of collections of people, one that *potentiated* sets of actions that fit together in complementary ways. On such views, organization is a property of sets of *actions* only in a derivative sense, i.e. only if these actions were performed by people who were themselves organized. But in fact, I suspect that our attributions of organization generally track the presence of a certain kind of actions—those that complement one another, in a sense to be elaborated on. And because many different kinds of relations among people potentiate complementary actions, our attributions do not track any one potentiating state-type thick enough to satisfy our conditions on substantiveness. So it makes sense that we failed to find a potentiating state-type that gave us a substantive necessary condition on organizing commensurate with the way the term is actually used.

The suggestion I want to advance in the next is that social organization is fundamentally a property of sets of actions, and derivatively of the collections of people who have performed these actions—an inversion of what was supposed by the foregoing accounts. It will still be accurate to say, as we initially assumed, that organizing brings about a new sort of collectivity among agents, but this collectivity will be of a kind that exists among agents because of how their actions connect. And this connection among actions will be one that obtains because of the way the actions fit together as causes of a common effect, i.e. the way in which they are *causally complementary*.

4 Organizing Increases Causal Complementarity

4.1 Causal Complementarity: A non-technical first pass

That organizing is a causal notion is, in one sense, uncontroversial. Our framing remarks in §2 assumed that any act of organizing is committed to causing both the operands-to-output transformation and, by this means, the achievement of the goal-set. This much was implicit even in the above views. On the view I will suggest though, organizing is a causal notion in the further sense that the nature of the transformation of operands into output is also to be characterized in terms of causal relations. Namely, organizing for a goal, G, brings it about that

the actions individual agents respectively undertake as intentional means to accomplish G (call these *G-directed actions*) support one another's status as causes of G to a greater degree than would otherwise have been the case. This sort of support among actions I will call *interagential causal complementarity.*

My proposed substantive necessary condition is then as follows:

> **Organizing$_7$**: Organizing is an activity which, in order to achieve a further goal, makes some agents more causally complementary than they otherwise would have been.

While this diverges from the accounts that we discussed above, it also helps make sense of why they initially appealed to us. When we reflect on what makes entering into configurations of grouphood or collective intentionality significant in the first place, we notice that these configurations allow the individual's actions to be *causal* where otherwise they couldn't be A worker withholding their labor until the boss makes concessions is unlikely to result in their desired changes if they are the only one doing it; but if they form part of a group agent like a union, the members of which all participate in a strike, that same worker's same action may now form part of a causal web terminating in the desired changes.

It's evident then how others' complementary actions can make it possible for an agent's actions to form part of a causal chain when that would otherwise have been impossible. Configurations of shared intention and group agency tend to require more than a mere conjunction of complementary actions of course (e.g. individual commitments to be jointly committed, certain higher order beliefs, role assignments), and we can see the further requirements of these configurations as ways to make the co-occurrence of complementary actions more likely. These configurations in their own ways tend to cultivate interagential causal complementarity with respect to a given goal. Whereas O1-O6 connected the notion of organizing to these various *means by which to get* to interagential causal complementarity, the view I endorse cuts out these intermediaries and connects organizing to interagential causal complementarity directly. In principle, O7 allows for organizing to take place without communication, intentional coordination, or indeed recognition of any kind among those who are organized. And while all the actions that satisfy O7 may also satisfy O6, O7, unlike O6, will turn out to be able to satisfy our principled measures of substantiveness. To appreciate those merits of the account however, we first need to take a deeper dive into the notion of interagential causal complementarity. I will proceed at first with some discussion designed to give us an intuitive grip on the phenomenon before, in §4.2, offering a rigorous formal characterization. In §4.3 I then make explicit how O7 fares better than O1-O6 with respect to the cases and theoretical desiderata we've considered.

Consider a case:

> **Case 6a:** Neighbors Anne, Bob, and Carl do not know each other and have not coordinated in any way. However, each of them independently has the goal, G, of getting more bicycle infrastructure built in their shared neighborhood. Each of them meets with the alderperson to request the infrastructure, and each also presents

their case to the municipal budget committee. The infrastructure does get built as a result of these individuals' G-directed actions— that is, as a result of union of the sets of actions that Anne, Bob, and Carl respectively took *as an intentional means to getting more bike infrastructure in their neighborhood.* However it turns out that it was completely unnecessary for any of them to have lobbied the city's budget committee when the funds to build the infrastructure were already available should the alderperson choose to use them. Moreover, the alderperson stood ready to commission the building of the bike infrastructure with these funds so as long as at least two of her constituents came to see her about it—three was unnecessary.

The agents in this set, Anne, Bob, and Carl, were—we may for short just say "the system" was, so long as we bear in mind that the system is just the sum of its agents and our account does not rely on any prior notion of collectivity—therefore, effective with respect to G[13]. But what we might further ask of this system is whether all three individuals' actions supported one another's status as causes.

In such cases we may judge that although the system was effective, the application of its constituents' agencies wasn't economical—they did a poor job of causally complementing one another.

Now consider a variant on the situation we just entertained:

Case 6b: Dana, a neighbor of Anne, Bob, and Carl, reaches out to them before they start trying to get the biking infrastructure built. Dana communicates the fact that the alderperson has all the relevant power here, and that it will take exactly two letters from constituents to get her to act. Dana moreover connects Anne, Bob, and Carl to one another so that they can decide amongst themselves which two of them will write letters. None of them now wastes time lobbying the budget committee.

In both 6b and 6a, the system constituted by Anne, Bob, and Carl is effective in getting the infrastructure built. But in the latter variant, this system features a higher degree of interagential causal complementarity than in the former. Because it was Dana's action that brought this higher degree of causal complementarity about, her action then satisfies the necessary condition on organizing articulated by O7.

Assuming for now that a goal, G, is achieved, perfect causal complementarity exists among a set of events with respect G when the subtraction of any one of those events would have been enough for G not to have taken place. In order for causal complementarity to play the role in a theory of organizing that we want it to though, we need it to be a gradeable notion. We get this by devising a way of quantifying how close some events come to the status of perfect causal complementarity.

A first pass at rendering interagential causal complementarity a gradeable notion in the way that we want is easily constructed. Non-redundant causes get weighted at 1, non-causes at 0, and various sorts of redundant causes

[13]Where G is satisfied but none of the individual constituents' G-directed actions were causally relevant, we might not say that the system was even effective.

at somewhere in between. Finally, we calculate the interagential causal complementarity of a system of agents relative to G, based on the average weightings of the G-directed actions within it.

It turns out though that there are some tricky issues that this first pass won't quite deal with—for instance, how do we calculate the interagential causal complementarity of systems that aren't effective? Consider:

Case 6c: Everything in 6a occurs with the difference that, just as the alderperson is preparing to commission more cycling infrastructure, an unforeseeable and unavoidable natural disaster hits the city. As a result, a state of emergency is declared and discretionary spending is suspended for the foreseeable future. The cycling infrastructure therefore doesn't get built.

Case 6d: Everything in 6b occurs with the difference that, just as the alderperson is preparing to commission more cycling infrastructure, an unforeseeable and unavoidable natural disaster hits the city. As a result, a state of emergency is declared and discretionary spending is suspended for the foreseeable future. The cycling infrastructure therefore doesn't get built.

Surely, if it was right to say that Anne, Bob and Carl were organized in 6b, we should want to say it in 6d too. But notice: if the cycling infrastructure wasn't built, then none of their G-directed actions was a cause of G, and the average weighting of their actions was therefore 0 in both 6c, where Dana doesn't make her intervention, and 6d, where she does. Dana's action therefore doesn't increase the system's causal complementarity and so, per O7, doesn't count as organizing. This is the wrong result, but here I think the problem is not with O7, but with a flat-footed method of calculating causal complementarity. While I will affirm that a system whose goal is *impossible* truly should be said to be disorganized with respect to that goal (see §4.4.1), systems of agents who take up a goal that is possible but not actually achieved do get a shot at being counted as organized.

In §4.2 below I lay out a more detailed method that will deal with these issues, by using some ideas from the interventionist paradigm in the causation literature. Broadly, the strategy is to consider the causal profiles of agent's G-directed actions in nearby scenarios where their goal *is* achieved. A reader who would prefer to skip a detailed exposition of this strategy might get by with something like this first pass in mind and so may proceed straight to §4.3.

4.2 Ascertaining Degrees of Causal Complementarity

Properly formulated, a question about whether an action counts as organizing is relativized to a set of agents (A) and a goal (G). So when we ask whether what Dana, above, did counted as organizing, a disambiguation of the questions is: did Dana's action organize the collection of agents comprised of Anne, Bob, and Carl with respect to the goal of getting further bike infrastructure commissioned? This section will lay out the apparatus needed for precisely answering this question.

To calculate the interagential causal complementarity of some system of agents with respect to a goal, we want to picture how these agents' G-directed actions fit into the total causal story concerning G. To do this, we will build a specialized kind of causal model. I will tend to interpret this model in terms of a broadly counterfactual paradigm for thinking about causation, in which, "X causes Y" is generally understood to mean that if X hadn't taken place, neither would Y[14]. The importance of the formal machinery I lay out below is then not that it allows for a reductive account of causation, but that it will prove critical in articulating how causal relations are to be weighted and how these weightings should be fed into a calculation of causal complementarity.

As is familiar from the causal modeling literature, the models that we make use of will consist of a set of variables, a set of structural equations that reflect how these variables' values depend on one another[15], and a function that assigns all variables a value consistent with these structural equations. All this information except some of the details of the structural equations is also captured by a directional graph notation, which has the benefit of being highly intuitive. While in subsequent diagrams I will therefore provide just the graphical notation alongside the list of structural equations, in the following I include both notations in full:

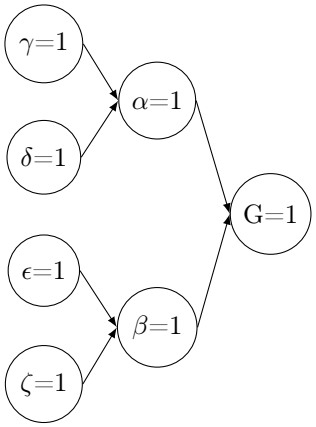

Variables: G, α, β, γ, δ, ϵ, ζ

Equations:
G:= $\alpha \vee \beta$
$\alpha := \gamma \vee \delta$
$\beta := \epsilon \wedge \zeta$

Values:
G=1
α=1
β=1
γ=1
δ=1
ϵ=1
ζ=1

Figure 1: A sample causal model in graphical and equation notations

We can think of a goal event, G, as a 2-valued variable, where a value of 1 indicates that G takes place, and a

[14]My practice here will be to give these informal counterfactualist glosses of the first order causal notions we discuss and then interpret the model in terms of these notions, rather than ever defining the first order causal notions in terms of the technical terms of the model. Such non-reductive uses of manipulationist/ interventionist frameworks to understand particular scientific claims is common in work by non-philosophers like Pearl [2000]. This practice is however deviated from elsewhere, both by those who think of these models as one formal implementation of counterfactualism [e.g. Woodward, 2003, Briggs, 2012], and by those who see this sort of modeling as allied with an alternative to counterfactualism [e.g. von Wright, 1971, Menzies and Price, 1993]. My hope is that the practice I adopt here avoids complexities orthogonal to our project, and preserves an ecumenicism concerning the nature of causation where possible. As will be evident from the form of my structural equations however, I do assume a deterministic picture of causation here. I think however that a probabilistic version of the picture is possible.

[15]Note that the ":=" operator used in structural equations is unlike the more familiar "=" in not being symmetric. G:=A means that the value of A determines the value of G, but doesn't entail the reverse.

value of 0 indicates that it does not. In the model above, the value of G is represented as depending on 6 other atomic variables (α, β, γ, δ, ϵ, ζ)—we will assume throughout that all the variables we are interested in represent events and have two values that correspond to the event's either happening (1) or not (0). Note that while generally when we talk about variables we will mean atomic variables it will sometimes also be helpful to think of compound variables, which for us will be built out of atomic variables and truth functional connectives $\wedge$, $\vee$, and $\neg$. In the directional graph notation, vertices are variables, and edges represent relations of causal relevance, such that arrows move from causally relevant variables to their effects. In such graphs, we will say that α is the child, or more generally the descendent, of β and γ, that β and γ are the parents, or more generally the ancestors, of α.

A note is in order concerning how any variable could have as few as six ancestors— aren't all contemporary events such that their happening is in some way causally dependent on, say, the birth of the universe, and many other major events between then and now? Shouldn't any causal model then contain many, many variables? In the causal modeling literature, in order to keep this set of variables minimal and not have it expand to include events going all the way back to the big bang, models will often include only the *Markovian parents* of the effect whose causal history is being investigated [Pearl, 2000, 14]: the Markovian parents of a variable are the minimal further set of variables that are sufficient for determining the value of that variable[16], so that the inclusion of these variables' ancestors in the graph is strictly speaking redundant. In the above graph, strictly only α and β are Markovian parents of G, since they are the smallest subset of variables such that the value of G can be determined based on them alone.

However, the particular purposes for which we are constructing our causal models will require us to adopt a slightly different procedure for choosing which variables to include in our models than this standard Markovian Parentage determination. Consider the variations on case 6, discussed above: it might turn out that the goal variable, getting more bike infrastructure commissioned (G), has as its Markovian parents the 3 variables of the alderman submitting the request (R), the funds being budgeted for the project (F), and there being a natural disaster (D):

Actions, like those of Anne, Bob, and Carl, are of course a special kind of event, and so might be thought of as 2-valued variables just like G, F, R and D— we will call variables corresponding to actions taken by the set of agents whose degree of organization we're interested in assessing (A), A-variables. You'll notice however that there are no A-variables present in Figure 2, because they aren't Markovian parents of G. This by itself does not of course mean that no A-variables are relevant to the value of G, just that their value is superfluous for the determination of G when already taking into account the values of R, F, and D. But where our interest is in determining the causal role of the G-directed actions taken by some set of agents (A), we need a procedure that ensures that these variables get into the model. This observation pulls us to include two special subsets of A-variables: first, we want

[16]Within certain constraints, there will always be a unique such set; see Pearl [1988].

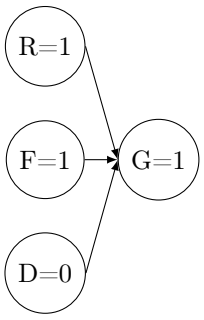

Equations:
G:= (R∧ F)∧¬D

Figure 2: Model of Case 6a/b with Standard Markovian Parentage Variable Determination

to make sure we include variables corresponding to any G-directed actions that members of A *could* have taken that are causally relevant to G; second, we want to include variables corresponding with all the G-directed actions that agents *in fact* took in the system whose organization we want to evaluate. The latter will fail to be a subset of the former where some G-directed actions taken by the agents *could not* have caused G, and so do not correspond with variables that are causally relevant to G; being G-directed and being causally relevant to G are distinct properties. So let's deal with the inclusion of these possibly distinct sets in order. Call the set of variables that are causally relevant to G and which correspond to possible G-directed actions taken by members of A $A_{G-relevant}$. It will not suffice to say that our model should include the union of the set of G's Markovian parents with $A_{G-relevant}$. This method would include the elements of $A_{G-relevant}$, but might omit the variables which connected elements of $A_{G-relevant}$ to G's Markovian parents. Without maintaining this connective tissue, the methods we'll describe below for calculating the weighting of the elements of $A_{G-relevant}$ at a model won't work. Instead, what we should do is apply Markovian Parentage determination recursively until this process yields a set that includes every element of $A_{G-relevant}$. Call the resulting set G's *A-inclusive Markovian Ancestry.* In order to obtain the full set of variables that we want included in the model, we should then take the union of G's A-inclusive Markovian Ancestry with the set of variables corresponding to G-directed actions that members of A in fact took in the system we want to evaluate, or $A_{G-directed}$. Note that we do not have to add a further step to include variables that connect elements of $A_{G-directed}$ to G, because all elements will either also be in $A_{G-relevant}$, in which case that work is already done, or they will be causally irrelevant to G, in which case there is no causal connective tissue to be recovered.

Let's then apply this to case 6a. To keep the case simple, we will assume that the set of G-directed actions that could have been taken by Anne, Bob, and Carl and which are relevant to their goal of getting more bike infrastructure built (i.e. $A_{G-relevant}$) contains only the actions of lobbying the alderperson. We note that their actions lobbying the budget committee are elements of $A_{G-directed}$ but not of $A_{G-relevant}$. In this figure, variables labelled A_x are elements of the union of $A_{G-directed}$ with $A_{G-relevant}$.

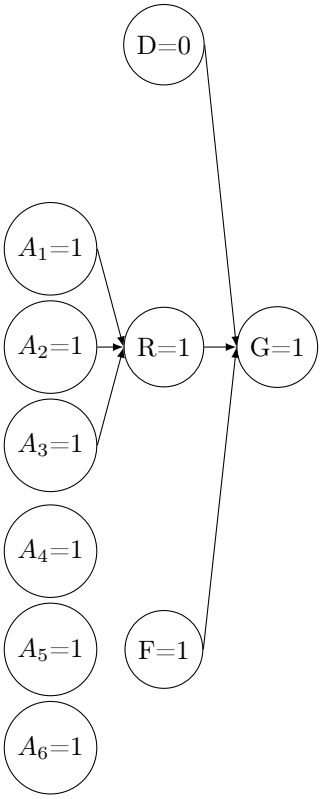

Equations:

G:= (R∧ F)∧¬D

R:= $(A_1 \wedge A_2) \vee (A_2 \wedge A_3) \vee (A_1 \wedge A_3)$

Figure 3: Model of Case 6a with Updated Variable Determination

And apply this same procedure to Case 6b:

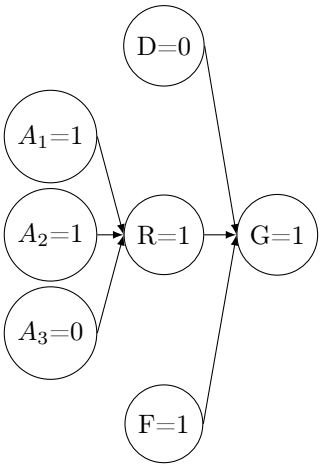

Equations:

G:= (R∧ F)∧¬D

R:= $(A_1 \wedge A_2) \vee (A_2 \wedge A_3) \vee (A_1 \wedge A_3)$

Figure 4: Model of Case 6b

Note that some variables don't appear on the left hand side of any equation (or, equivalently, on the sharp end

of any arrow); some of these will be causally irrelevant variables, but others are causally relevant but such that their own causes aren't included in the model. Either way, we say that these are *exogenous* variables. Those variables that do appear on the left hand side of an equation (or sharp end of an arrow) we say are *endogenous*.

We've now seen how to construct a model that represents the systems we want to evaluate for their interagential causal complementarity. This machinery will turn out to be very helpful for articulating the quantitative weighting of causes with respect to events that we will want to feed into our calculations of causal complementarity, and that is what we turn to next. But it is worth heading off a confusion here, concerning the kind of causal weighting we are interested in. Scientists and social scientists working with causal models are of course interested in a variety of such quantification in the form of determining a "path coefficient" for links between particular variables in a model. To use an example drawn from Sewall Wright's pioneering use of causal models (Wright [1920], Pearl and Mackenzie [2018, 82-84]), where the weight of a new born guinea pig pup is a function of both length of gestation and litter size, and where gestation length is in turn partially dependent on litter size, these path coefficients capture how much of the pup's birth weight should be attributed to the direct effect of an extra day of gestation, as opposed to the size of its litter. While my strategy here draws on the basic idea of quantifying the causal significance of variables with respect to one another, the feature I'm trying to quantify is a little bit different. Where two different causes were required to bring about an effect, we are not interested in asking which was *more* causal and reducing the score of the other— this would, in a sense, amount to putting causes into competition with one another, whereas our goal is to assess how well causes *complement* one another.

The things that we want to detract from the weighting of a variable at a model are that the variable is causally irrelevant, counterproductive or redundant. With all this in mind, the following lays out our method for attributing appropriate weightings to our variables. Note that $Val(V)_{G,m}$ denotes the value of variable V with respect to goal G at model m.

Weighting: Let the weighting of some variable, V, with respect to some goal, G, and at some model m, be written as $\mathrm{w(V)}_{G,m}$.

I *Negatively valued variables:* If $Val(V)_G, m$=0, let $\mathrm{w(V)}_{G,m}$=0.

II *Positively valued variables:* If $Val(V)_G, m$=1, let $\mathrm{w(V)}_{G,m}$ be determined according to the following:

(a) *Reflexivity*: For any V and any m let $w(V)_{V,m} = 1$.

(b) *Irrelevance*: Where, at m, a variable, V, is not an ancestor of G, let $w(V)_{G,m} = 0$.

(c) *Backward Inheritance*: Where X:=Y, let w(X)=w(Y)

(d) *Conjunction*: Where w(X∧Y)=z, let w(X)=w(Y)=z

(e) *Disjunction*: Where w($V_1 \vee V_2 \vee ... V_{|D|}$)=z, call the set of disjuncts D and let w(V_1)=w(V_2)=...w($V_{|D|}$)=$z(|\{x : x \in D \text{ and } Val(x) = 1\}|)^{-1}$

(f) *Negation*: Where w($\neg X$)=z, let w(X)=-z

(g) *Occurrence-summing*: Where an atomic variable, V, appears multiple times in a single equation (i.e. multiple times within a compound variable), let w(V) be the sum of the weightings of each of its occurrences, calculated as per the above.

(h) *Generosity*: Where an atomic variable, V, appears on the right hand side of multiple distinct equations in a single model, let w(V)$_{G,m}$ be the highest of those values resulting from calculating its weighting in each of those equations, reached as per the above.

In the above, clause IIa effectively establishes the idea that where an event takes place, it has a weight of 1 with respect to itself.

Clause IIb suffices to suppress the weightings of causally irrelevant variables in the model. Causal irrelevance, a feature that a variable has independent of the value that it or any other variable is assigned in a model, is when the value of a variable has no bearing on the value of G, given any set of value assignments for the other variables. Some causally irrelevant variables will appear in our models because we are including variables corresponding with those G-directed actions that in fact took place, some of which might have been misguided (i.e. the complement of $A_{G-relevant}$ with respect to $A_{G-directed}$). The model reflects their irrelevance through the fact that these variables do not appear in a structural equation and, graphically, aren't connected to any other nodes of the tree. Clause IIb assigns causally irrelevant variables a weighting of 0.

Our general approach for all other causally relevant variables is to let the weighting of a variable in a model be a function of the weighting of its immediate descendants, so that establishing a weighting for all variables in a model begins with G and works backward from there. Per IIa, G has a weight of 1 with respect to itself. And clause IIc then allows us to start working backward through the generations of G's ancestors to assign them weightings as well. This however will often suffice only for assigning weights to compound variables, when what we will really care about is the weights of the atomic variables within them. Clauses IId-IIg detail how to determine the weighting of an atomic variable from the weighting of the compound variable it appears in, working recursively from the main connective in the variable.

IId reflects the intuition that if an outcome requires two distinct events to occur, then each of these events should fully inherit the weighting of their descendent; the full weighting of a conjunction is passed on to each of its conjuncts.

IIe handles disjunction and therefore is what will allow us to appropriately handle cases of redundancy. Informally, an event, e, is a redundant cause of G where e is an element of some non-singleton set of events such that

if none of its elements had occurred then G wouldn't have occurred, and such that each element is such that if it but no other element of the set had occurred, G would still have occurred[17]. Our models reflect redundancy mainly through the appearance of disjunction in their structural equations. For instance, in Figure 1 above, the structural equation G:=$\alpha \vee \beta$ reflects the fact that α and β are redundant with respect to G. The intuition expressed by clause IIe above is that where either of two events suffices to bring about their descendent, then the weight of that descendent should be split equally among these events; the weighting of a disjunction is split among the disjuncts. This by itself would suggest dividing the weight of the disjunction, z, by the number of disjuncts, |D|. But there is one further complexity here: we only want to penalize realized redundancy—i.e. cases where two redundant variables both had values of 1— not cases where there was merely a potential for redundancy that was never realized. So we want to divide the weighting of the disjunction, z, by the number of disjuncts that have values of 1 at the model, $|\{x : x \in D \text{ and } Val(x) = 1\}|$. To see how this plays out when applied to a model like that in Figure 1, observe that $w(\alpha)_{G,m} = w(\beta)_{G,m} = w(G)_{G,m}/2$=0.5. This reflects our intuition that redundancy should receive a score in between that of full- and non-causes. But this notion of weighting has one further notable nuance: $w(\gamma)_{G,m} = w(\alpha)_{G,m}/2 = 0.5/2 = 0.25$, while $w(\epsilon)_{G,m} = w(\beta)_{G,m}/1 = (0.5)/1 = 0.5$. In this way, redundant causes of immediate effects that are themselves redundant causes of their immediate G-relevant effect are penalized extra, and this penalty may be further compounded moving up through G's ancestors.

It is worth noting that, in the above, I don't distinguish between symmetric and asymmetric causal redundancy. A paradigm instance of symmetric redundancy is this: two rocks are thrown at a window and strike at precisely the same time; the window shatters as both rocks make contact. In contrast, in a case of asymmetric redundancy, also known as preemption, the two rocks are thrown at a window, but one strikes it a few seconds before the other; the window shatters upon contact with the first rock, but had that rock not shattered it, the second rock would have. My characterization of causal redundancy captures both, and I am happy to be read as proposing that redundant causes, in both symmetric and asymmetric cases, be treated alike. It might be argued that, in cases of asymmetric redundancy, the preempting cause should be "rewarded" to a greater degree than the preempted one by weighting. If our notion of weighting was ultimately in the service of comparing the causal importance of individual events to one another this might be so: the first rock really does seem to have a stronger claim to causal importance than the second rock does. But ultimately, weighting will be fed into the below notion of causal complementarity which is to be used to assess and compare *sets* of actions, rather than individual actions. When causal complementarity is evaluated for a system members of which performed asymmetrically redundant actions, this counts against the

[17]Some have suggested that all causes are redundant causes [Mellor, 1995, Mills, 1996, Sturgeon, 1998, Schaffer, 2004]: events have both a macro- and a micro-level causal history (and indeed, perhaps innumerable meso-level ones), any of which would be sufficient for their occurrence, and so, the suggestion goes, any macro-level cause is co-redundant with a micro-one. Evidently, when I distinguish between redundant and non-redundant causes here though, I am distinguishing between those that are redundant with another cause of the same level, and those that aren't. There are of course interesting questions about what sorts of events will be of the same level as G-directed actions; however I take it that these questions are to be addressed to the entire literature discussing overdetermination amongst macro-level causes, and not to me in particular.

system as a whole, and whether the coredundant causes are weighted at 0.8 and 0.2 respectively or else 0.5 and 0.5 will make no difference to this final calculation[18].

Clause IIf effectively deals with counterproductivity. Counterproductivity is when an action is causally relevant to G only insofar as the occurrence of that action contributes to G *not* occurring. Counterproductivity can be roughly discerned in our models where the value of one variable depends on the negation of another. Note though that the fact that a variable, V_1, appears with a negation sign in front of it in some structural equation doesn't necessarily mean that it is counterproductive with respect to G; it might be that switching V_1 on switches off a further variable, V_2 that is genuinely counterproductive to G, in which case we would expect to see $V_2{:=}\neg V_1$ among the structural equations in the model, but it isn't in the end true that V_1 is counterproductive with respect to G. Our method of weighting tracks genuine counterproductivity.

Clause IIg deals with cases where a variable appeals multiple times in the same equation. For instance, case 6b as represented in figure 4 sees variables A_1, A_2, and A_3 each appearing twice in the equation R:= $(A_1 \wedge A_2) \vee (A_2 \wedge A_3) \vee (A_1 \wedge A_3)$.

Finally clause IIh reckons with the possibility that a variable may have multiple distinct children, and that in such cases there will be more than one path to determining the weighting of this variable.

We now have a method for weighting individual variables at models. The next step is to feed these weightings into a calculation of the overall causal complementarity of a system. As suggested above, the basic idea here is that to determine the complementarity of a system of agents, A, with respect to a goal, G, in a given situation, we should take average of the weightings of the G-directed actions members of A take in the model we have constructed to represent that situation. And for some cases this might work perfectly well. For instance, consider the result it gives for dealing with cases 6a and 6b. Let m be the model we constructed (in figure 3) of case 6a, and m' be the model we constructed (in figure 4) of case 6b:

In case 6a, there are 6 actions taken by Anne, Bob and Carl. The above system of weighting yields the result that A_1, A_2, and A_3 (their respective acts of lobbying the alderwoman) each get weighted at 2/3. And A_4, A_5, and A_6, their respective acts of lobbying the budget committee, are each weighted at 0 because they are causally irrelevant. Per our intuitive first pass at calculating, their interagential causal complementarity with respect to their goal is (3(2/3)+3(0))/6= 0.33.

In case 6b, there are 2 G-directed actions taken by Anne, Bob, and Carl. The above system of weighting yields the result that A_1 and A_2 are each weighted at 1. Per our intuitive first pass at calculating, their interagential

[18]A purpose for which it might make sense to weight preempting and preempted causes differently is if we wanted to compare the respective causal complementarities of two different systems in the same situation with respect to a single goal. Supposing that agents in one of these systems were consistently preempting the actions of those in the other system, it might seem like the preempting system ought to come out with a better complementarity score, and only different weightings for asymmetrically redundant causes would assure this. I opt not to implement this idea here though because it would add significant complexity to weighting calculations without serving the purpose that is our present focus: characterizing whether organizing has taken place at all, not whether one set of people organized better than another.

causal complementarity with respect to their goal is (1+1)/2=1.

Where Dana's action brought it about that the world resembled case 6b rather than case 6a, she is responsible for rendering the relevant system of agents more causally complementary with respect to their goal. So she satisfies the necessary condition on organizing articulated by O7.

However, as discussed above, this method will not serve us well in all cases. We've seen that a collection of people can be effective without being optimally causally complementary (e.g. in 6a). But in order to make sense of the reverse, a collection being causally complementary without being effective (e.g. 6c and 6d), we need to consider cases in which G does not in fact come to pass. If G is never in fact accomplished, then the weighting of G in the model is 0 (from clause I), and the weights of all its ancestors will, by the method described above, also be 0. Clearly averaging a bunch of zero values would likewise yield zero. So the simple averaging method, if used to calculate causal complementarity, would yield the verdict that a system that didn't accomplish its goal wasn't organized with respect to that goal at all. Now, because our method of weighting can yield negative values (see clause IIf), there can be negative causal complementarity scores as well. A system with a complementarity score of 0 is then such that it could have a greater score than some other system, and so an act that brought it about might still satisfy the condition on organizing articulated in O7. But it is nonetheless clear that the simple averaging method will prevent many activities that looked paradigmatically like organizing—say cases like 1 and 2 from earlier in the paper— from counting as organizing where they fall short of actually bringing about their ends. I regard this as undesirable; I want it to be the case that certain kinds of "failed" organizing are still organizing.

To address this, we want to consider counterfactual weightings—i.e. weightings of actions in causal models other than m, where m describes the system we're evaluating. While perhaps, in m, G doesn't take place, and so no actions taken by members of A are causes of G, this doesn't mean that these actions wouldn't be among the causes of G, were other complementary causes to have been present alongside them. The general thought here is to assess the weighting of the relevant actions in a slightly different causal model in which G does occur. We can create a submodel of m in which this is the case via a special kind of *intervention*. An intervention can be intuitively characterized as an "ideal experimental manipulation" of a variable in a causal model like ours [Woodward, 2003, 94]. An intervention on an exogenous variable simply reassigns its value. An intervention on an endogenous variable erases the equation in which that variable appears on the left hand side— in the graphical notation, the intervention erases any arrows that point to the endogenous variable—and simply stipulates a value for it, effectively making it exogenous[19]. This direct manipulation of the variable then may indirectly alter the values of its descendants in the model, but the equations that connect it to those descendants are left intact. An intervention thus produces a submodel which allows us to consider a state of affairs which is different from that in the original model but which preserves most of its causal dependencies. We will use the expression "intervention" in a way that allows a single

[19]For more detailed discussion of the notion of an intervention, see Woodward [2003, 94–114]

intervention to involve this sort of direct experimental manipulation of more than one variable.

In calculating the causal complementarity of a system of agents, A, with respect to a goal event, G, in a model, m, we need to look at a submodel produced via an intervention on m. But where many different interventions would produce many different submodels, which one do we want? I argue that it makes sense to compare m to the submodel reached via a *minimal agency-preserving G-intervention*, or *MAG* for short.

To find the MAG, first take the set of interventions on m which, indirectly—i.e. by directly manipulating G's ancestors, not G itself— give G a value of 1. If this set is empty, because G is impossible and there is nothing that could cause it, let the MAG of m be m itself[20]. Note that manipulations on m might involve switching variable values from 0 to 1 (i.e. making events happen that didn't in m) or switching values from 1 to 0 (i.e. making events not happen that did in m). Call the former a positive manipulation and the latter a negative manipulation. Second, find the subset of the remaining interventions that involve manipulating the smallest number of variables, n— the minimal intervention(s)— and throw out the rest. And third, of those interventions remaining, select the one in which the highest proportion of the n manipulated variables are A-variables— this effectively amounts to holding the A system maximally responsible for what its agents could have done given the situation modeled by m[21].

Note that if m is already a G-model, the minimal intervention that "makes it into one" will involve the manipulation of 0 variables, which will be the only one left after step two above. So, as defined, where m is a G model, it is also yielded by the MAG on itself.

We're finally now in a position to present our method for computing interagential causal complementarity:

Interagential Causal Complementarity: Where A is a set of agents and G is a goal event, let $A_{G-directed,m}$ be the set of G-directed actions taken by the agents in A in model m.

i where $A_{G-directed,m}$ includes at least two actions such that they were taken by different members of A, and where m' is the model produced by the MAG on m, let the degree of causal complementarity of A with respect to goal G in model m be the sum of the weights at m' of the elements of $A_{G-directed,m}$, divided by the number of events in $A_{G-directed,m}$ plus the number of positive manipulations of A-variables (|PMAV|) needed to produce m', i.e.:

$$\frac{\sum w(x)_{G,m'} : x \in A_{G-directed,m}}{|A_{G-directed,m}| + |PMAV|}$$

ii where $A_{G-directed,m}$ does not include at least two actions such that they were taken by different members of A let the degree of interagential causal complementarity of A with respect to goal G,

[20] Per the equation laid out below, this will result in a complementarity score of 0, but in this case that will be appropriate. For more discussion of this consequence, see §4.4.1.

[21] In the event that this procedure yields a set of more than one intervention, let the causal complementarity of A in m be given by the averages of the scores calculated by the below method with each of the intervention models in turn

in model M, be 0.

Per clause i, determining a complementarity score starts with determining the weights of variables corresponding to G-directed actions that the agents we're evaluating in fact took. However, it looks at the values of these variables not *at* m, but rather at m', the submodel generated by the MAG on m. Notice a few things about this: if a variable corresponding to a G-directed action by members of A is switched on between m and m', its weighting will not appear in the sum in the numerator, because this variable is not among those that have positive values at m and so not an element of $A_{G-directed,m}$. However, this variable will still add to the denominator because it is a positively manipulated A-Variable, the number of which is captured by the second term of the denominator. Because a positively manipulated variable adds to the denominator without contributing to the numerator, it will drag down the total complementarity score, which is appropriate: after all, the fact that there was something our system of agents could have done to bring about G that they didn't do in the system modeled by m is the sort of thing we want to drive down our assessment of how organized they are. What if a variable corresponding to a G-directed action is switched *off* between m and m'? This will only take place where the action corresponding with this variable was counterproductive to the achievement of G. Per clause I of weighting, because the value of this variable will be 0 at m', its weighting will likewise be 0. So although the counterproductive variable does get summed in the numerator of the equation, it will not drive the value of the numerator up. This variable will still however contribute to the value of the first term of the denominator, and the net effect will once again be to drag down the overall complementarity score score, as makes sense: we should want to penalize systems for being counterproductive. The reason that the second term of the denominator takes into account only the *positively* manipulated A variables is because otherwise in cases of counterproductivity the counterproductive variable would be double-counted in the denominator, driving the score down even further. If one regards counterproductiveness as twice as bad as a mere failure to act productively (i.e. the state of affairs remedied by positive manipulations of AVariables) then perhaps this would be appropriate. I however have opted to penalize them equally.

The way that we have characterized the kind of counterfactual circumstance that we want to consider in cases where organizing isn't effective (i.e. the way we've characterized MAG) is designed to allow us to give non-zero complementarity scores to ineffective systems while not letting them off the hook for their failings. Because the MAG on the model of 6c yields the model of 6a, and the MAG on the model of 6d yields the model of 6b, this method assigns a complementarity score of 0.33 to case 6c and 1 to case 6d. This means that Dana's actions can count as an instance of organizing even in the presence of the natural disaster which prevents the cycling infrastructure being built.

Clause ii of the above notion of complementarity is motivated by two considerations. The first is that it is important that, where $A_{G-directed,m}$ is empty, we not compute complementarity using the equation in clause i,

because this risks giving us a zero in the denominator in cases where PMAV is also empty [22]. This would yield an undefined value for our complementarity score. One might on first glance think that it's appropriate that, where members of a set are not yet taking any actions toward G, they don't get a complementarity score with respect to G at all. To see why this won't do, consider a variant of the paradigm cases of organizing that we discussed above, say union organizing, in which, without the intervention of the organizer(s), the individuals in the workplace would not take any action toward the goal of achieving, say, higher pay, but after the intervention they do. According to O7, this is an instance of organizing just in case the score of the output state is higher than the score of the operands state. But there can be no increase over an undefined value, and so if the operands state here had an undefined complementarity score, the above couldn't count as organizing. We avoid this by defining the causal complementarity of such states as zero.

The second consideration motivating clause ii above is the idea that social organizing should involve bringing about some *interpersonal* synergy— that is, it should require bringing about complementarity between distinct individuals' actions. Without something like clause ii, an intervention could count as increasing the score of a system (and so, present any other necessary conditions, as organizing it) where it increased only the proportion of a single individual's actions that were causally relevant to their goal. Because I think we should withold a verdict of social organizing from such cases, we need to assign low scores to systems in which there is only one individual who is engaging in G-directed action, and clause ii does this. I think it's quite right that there is a broader notion of causal complementarity that includes single agent cases; it's just that this is not the kind of causal complementarity that I think we should plug into an account of social organizing. Although the above presentation has spoken simply of "causal complementarity", we are most accurately seen as offering a notion specifically of *interagential* causal complementarity.

4.3 Putting the positive proposal to the test

Now that causal complementarity has been more rigorously characterized, we're in a position to ask whether O7 succeeds where O1-O6 had failed. Below I consider each of six desiderata on a satisfactory theory that we established over the course of the above discussion and make explicit how O7 meets each one.

4.3.1 Tests of necessity:

In the course of §3, we identified five varieties of cases that a necessary condition on organizing must accommodate. These were cases involving NO SELF-CONCEPTION AS ACTING TOGETHER (excluded by O1), those involving NO SUSTAINED PATTERN OF COOPERATION (excluded by O2), those that involved *neither* a SELF-CONCEPTION AS

[22] You would get a value of 0 for both terms of the denominator in a model where none of the members of A are doing anything relevant to bringing about G but where G takes place anyway, so that the MAG doesn't flip the value of any variables

ACTING TOGETHER nor a SUSTAINED PATTERN OF COOPERATION (excluded by O3), and cases of both STRUCTURE-BASED ORGANIZING and INTERNAL ORGANIZING (both excluded by some interpretations of O4 and O5). O7 is able to include cases of all these kinds.

O7 allows that sometimes organizing will yield a system the agents in which, while individually having intentions to bring about G, may not know about each other's intentional states or see each other as playing a role in bringing G about— consider e.g. case 6a which, while not optimally organized, still received a complementarity score above 0. O7 therefore allows that a system can be organized even though its constituents have NO SELF-CONCEPTION AS ACTING TOGETHER.

Notice also that O7 doesn't actually require that *every* agent in the operands or output systems engage in G-directed action. In Case 3 the boss's intervention (i.e. disseminating instructions to her capos) can still count as organizing the whole syndicate for the goal of monopolizing the drug trade because, as required by clause ii of Causal Complementarity, there are multiple people in the syndicate who do end up acting with the intention to bring this goal about: the boss and each of the trusted capos. It doesn't matter that the soldiers don't act with this intention; they can still be parts of a system that is the output of organizing for this goal. Had the boss not made her modification, neither the capos nor any other syndicate member would have acted with the intention of bringing about the monopoly, and the syndicate's complementarity score with respect to this goal would therefore have been 0. In contrast, where the boss does make the modification, the capos engage in monopoly-directed actions. So long as at least some of these actions are such that they are causes of the monopoly (either in reality or the scenario modeled by its MAG), the system gets a score of greater than 0. And because the boss's action therefore yields a higher score than would otherwise have been the case, it satisfies the necessary condition on organizing articulated in O7.

O7 likewise allows that a system can be organized even with NO SUSTAINED PATTERN OF COOPERATION or disposition to sustain such a pattern. A system can in theory be perfectly causally complementary with respect to a particular ephemeral goal while forming no infrastructure for causal complementation concerning other goals. Nor does O7, like O3, exclude cases featuring *both* NO SELF-CONCEPTION AS ACTING TOGETHER and NO SUSTAINED PATTERN OF COOPERATION. In Case 4 for instance, Frieda is satisfying O7 by bringing it about that her neighbors' actions undertaken as intentional means to hold condo developers accountable complement one another in achieving this goal even though she doesn't necessarily bring it about that her neighbors think of themselves as acting in concert or that they are any more likely to act collectively going forward.

Finally, O7 is easily able to accommodate cases of STRUCTURE-BASED ORGANIZING and INTERNAL ORGANIZING. O4 (and, under certain interpretations, O5) struggled with cases of structure-based organizing because such cases don't necessarily involve changes in membership or functional complexity. And they struggled with cases of INTERNAL ORGANIZING because these cases don't involve the creation of a new object. O7 in contrast doesn't associate

organization with the creation of a new object or with an increase in complexity. So O7 can capture what it is structure-based organizing characteristically changes about a system: it takes a collection of people who already bear relationships to one another and suggests that their existing form of collectivity can be wielded for different purposes than it has been theretofore. This will often involve people taking up actions directed at a goal that they hadn't pursued before, and/ or pursuing that goal via actions that complement one another better. This is precisely the sort of modification of a system of agents that O7 labels as organizing. And if Jeeyeon's internal organizing in Case 5 allows a greater proportion of her organization's members' G-directed actions to be causal with respect to whatever their goal is (perhaps because new leadership is able to bring in new ideas that make for more strategic decisions about how the organization should approach its goal) then O7 allows that she's engaging in organizing too

4.3.2 Tests of substantiveness:

We've seen that O7 passes tests on its necessity that O1-O5 did not. We recall that the failures of these proposals led us to briefly consider O6, which passed tests on necessity but didn't satisfy our principled measures of substantiveness, TRIVIALITY AVOIDANCE and DIRECTIONALITY. O7 however satisfies both.

O7 connects organizing to complementarity and complementarity comes apart from effectiveness in both directions: a system can be effective without being maximally causally complementary, and can be maximally causally complementary without being effective. O7 therefore steers clear of varieties of triviality that involve making organizing just "whatever makes you more, rather than less, effective". O7 also helps us steer clear of varieties of triviality that involve overgeneration, since not every action will boost the complementarity of a system with respect to a particular goal, G— acting on the system in a way that creates redundancy, causal irrelevance, or counterproductiveness, or indeed to bring about actions that, whatever their causal profile aren't consciously undertaken as means to accomplish G, won't count as organizing for G.

And because O7 equates organization with the gradeable notion of complementarity, and organizing with actions that increase this gradient, it satisfies DIRECTIONALITY in a way that O6 was not able to.

4.4 Objections and replies

4.4.1 Objection: This view forces us to say that systems with impossible goals aren't organized

Consider a collective which aims to restore the French monarchy by elevating to the throne a barber who cuts the hair of all and only those men who do not cut their own. They have therefore sworn themselves to the protection of this barber, whenever he emerges, so that he can go about his very important business. Their dedication involves meeting weekly to practice tactical maneuvers, raising money to buy the requisite arms, keeping watch

over barbershops all over France in case their barber should reveal himself, safeguarding one another's identities as members of this secret militia, and so on. There is a sense in which we might be inclined to call this collective quite organized, however deluded. Not only do they never bring about the ascension of this barber to the throne however, but there is no intervention of the kind described above that will produce a submodel in which this goal is accomplished either, since no manipulation of the included variables will change the fact that this barber doesn't exist. The above method for calculating complementarity therefore suggests that the score of this group with respect to their goal is 0. And if organization requires some degree of complementarity, doesn't this seem in conflict with our intuition that they are organized, even if they are doomed, as a matter of logical necessity, never to be effective?

There is one solution which I don't endorse but which I think is instructive to acknowledge before presenting my own preferred reply: we could weaken O7 to require only that the organizer *try* to bring about causal complementarity with respect to some goal, not that they actually do bring it about. Let's grant for the sake of argument that the above collective could satisfy this weaker condition[23]. If this weakened proposal were viable it would then allow us to say that the collective is organized with respect to the goal of elevating the barber to the throne.

Assume that a corollary to organizing's being constituted by intending to increase causal complementarity is that disorganizing involves intending to decrease it. The weakening proposal would therefore prevent us from ever coherently saying that a person intended to organize a system but accidentally disorganized it. This feels like a cost on its face, but it also matters because it attenuates the ability of either organizing or disorganizing to play the roles in social explanation that they are commonly called upon to play. Consider the well-meaning individual who intends to organize a social movement with respect to G but in fact brings about drastically reduced causal complementarity with respect to G, leading to the movement's failure. The unweakened version of O7, or an intuitive corollary of it anyway, says that the action here is one of disorganizing. We can therefore explain the movement's failures as a result, in part, of the fact that this individual disorganized it. And the explanation of this particular failure might fit into a broader theory of social change on which disorganization *generally* has these effects. But per the weakened version of O7, we would not be able to explain this failure as a result of the actor having disorganized the movement—because according to the weakened condition, she actually organized it. A picture emerges on which both organizing and disorganizing sometimes increase causal complementarity, and sometimes decrease it. Even if you believe that causal complementarity is of general explanatory significance, the case that either organizing or disorganizing has such significance would thus be much weaker. If the notion of organizing belongs in a general theory of social change, as many of the thinkers quoted in the beginning evidently believe, this will be because there is a regular relationship between organizing itself and certain further social effects. The weakened version of O7 chafes against this. While it's not our job to vindicate any particular theories of change as correct, we have been

[23]Of course it is highly controversial that one can even try to do something they *know* to be impossible, but we assume the members of the collective just haven't noticed this about their goal— far from being implausible, this of course puts them more or less in the company of Frege.

committed to retaining organizing's formal suitability to be deployed in such theories. Since the weakening of O7 disrupts this, it is evident then that this strategy has costs.

Independent of the above, the weakened version of O7 does not satisfy DIRECTIONALITY. Recall that DIRECTIONALITY required that any condition on organizing be compatible with the fact that organization is a gradable state-type and that *organizing is an action that increases the grade of this state-type that obtains.* If we agree that in the above case the collective does not in fact become more causally complementary with respect to elevating the barber to the throne, and organization just is causal complementarity, we agree that there is no increased degree of organization with respect to that goal. But the weakened version of O7 would allow that there was still an act of organizing with respect to that goal. Permitting organizing to come apart from organization in this way plainly violates DIRECTIONALITY. If we stick to our commitment to DIRECTIONALITY, this is sufficient reason to reject this strategy. The proponent of the weakening strategy might well of course suggest that we carry the strategy further and alter DIRECTIONALITY too. For reasons of space I will simply say that the revisions I imagine this proponent wanting to make would amount to dismantling the barriers to organizing's further explanatory irrelevance.

The response to this objection that I prefer is as follows: I think that the mere intuition that some group is organized underspecifies what it is they're organized with respect to. So the above account can absolutely account for the intuition that members of this collective are organized with respect to the goal of holding regular meetings or buying weapons or surveilling barbershops. It will just be wrong to say that they're organized with respect to the goal of installing the paradoxical barber as King of France. And this isn't such a big bullet to bite.

4.4.2 Objection: This makes it arbitrary what the operands of a given instance of organizing is.

An instance of organizing, per O7, is an action that renders some set of agents more causally complementary with respect to a given goal than they otherwise would have been. Since the complementarity score of the set is a function of the individual G-directed actions' weightings, increasing the weighting of even one agent's G-directed actions, while holding the others' fixed, is sufficient to raise the score of any set in which they're an element. And any given individual is an element of an infinite number of sets. This means that, for any token instance of organizing, there are an infinite number of sets that could be specified as the operands of the instance of organizing—that is, that could be specified as the thing *being organized*. This is an implication that I don't try to dodge: it will be strictly speaking correct to specify any such set as the operands. This certainly has some prima facie odd consequences: if Anne forms a tenants union with her neighbours, the natural thing to specify as the operands of this activity would seem to be the set of residents in her building. But on my account, it's strictly speaking just as true that Anne's organizing activity has had as its operands the set of all her building's residents *plus* Winston Churchill, or the set of half her building's residents plus the members of the Chicago Symphony Orchestra, never mind that neither Churchill nor the CSO had anything to do with the unionization effort.

But consider this: When a militant labor organizer, having helped one nursing home form a union, describes what they're doing as "organizing the working class", they are staking a claim to a set of individuals that extends beyond the individuals they've had any direct engagement with. And this claim is rooted in a particular analysis, one that has something to say about the reason that these workers needed a union in the first place, and what the broader causal and symbolic significance of their getting one is. Rodrigo Nunes [2014], in a discussion of how to individuate social movements (which he calls "network systems"), ultimately resolves that "the individuation of a network-system is itself a political construct." While different ways of carving the world up into movements may then be equally valid once these different ways are indexed to particular political perspectives however, the perspective that underlies any given construct can itself be critiqued based on "whether the relations it illuminates are objectively real" [2014, 25]. In something of the spirit of Nunes' comments then, I allow a huge range of operands specifications to be correct for any token instance of organizing. But I respond to the natural impulse to be able to criticize the gerrymandered operandses by pointing out that the specification of who is being organized by an instance of organizing suggests an analysis of which collections of individuals are politically salient— and these analyses can certainly be critiqued.

4.4.3 Objection: This proposal makes it seem like redundant actions are always a bad thing— but isn't redundancy often good?

Given that we act under uncertainty (what if John gets sick and can't lead the door-to-door canvas after all? What if Linda and Nancy get arrested and we don't have enough safety marshals to block traffic for the march?) it often makes sense for us to double agents up in roles that strictly only require one individual, or double up on strategies to achieve a goal when not all are required.

Yet my view does indeed say that implementing either sort of redundancy doesn't count as organizing a system with respect to that goal. This is not the problem that it might seem, for two reasons. First, deeming the implementation of redundancy not to count as organizing is not the same thing as deeming it bad. Part of avoiding triviality in our concept of organizing is actually making sure the connection between organizing and effectiveness isn't too tight. We should want it to be coherent to sometimes regard organizing as the wrong tool for a job.

Second, and echoing the structure of the point I made in 4.4.1, some forms of redundancy *will* count as organizing for some goal— it just might be a goal that the organizer takes to be merely instrumental in her project as she sees it. Suppose our organization needs one more legislator vote for a piece of legislation we care about. Although we are already lobbying one legislator who we hope will change their vote in our favor, you get us to start lobbying two more, in case the first doesn't come through. This won't count as organizing for the passage of the legislation. But it will still count as an act of organizing— organizing, that is, with respect to the goal of getting legislators 2 and 3 to vote in a certain way. If the intuition we want to save is that these actions count as organizing— never

mind how exactly we specify the goal— the above account is able to do this.

5 Conclusion: Implications and Remaining Questions

I've argued that organizing a group of people requires rendering them more causally complementary with respect to a goal. The project here has been a descriptive one: to say what organizing *is*. But what has hopefully been evident all along is that this descriptive project is motivated by the role that a notion of organizing plays in various normative discussions. One such normative discussion concerns the exact nature of our duties to collectivize.

A basic observation of the literature on collective responsibility is that there are ways of altering the world that no one person can effect; the sort of action needed to rapidly decarbonize the economy in order to avert even more devastating climate change, for instance, can not be executed by a single individual. A common thought has been that, while this incapacity entails that no individual can have a duty to solve the problem directly, the individual can have a duty to bring others together into a collective which can address the problem. Proponents of a duty to collectivize have differed concerning the flavor of collectivity that they think we have a duty to bring about. Collins [2013, 2019] has argued that we have a duty to bring about a group agent. Hindriks [2019, 208] has argued that something weaker than group agency which resembles shared intention and which he calls simply "joining forces" is what is morally required. What is interesting is that these theorists and others [e.g. Held, 1970] paraphrase their proposals as duties to *organize*. If the argument I've made above is right however, then although organizing may often be executed by bringing about group agency or weaker forms of collective intentionality, a duty to organize would be distinct from these other suggested duties. A question that our discussion has then raised is whether the duty to collectivize is truly a duty to organize after all. I don't offer an answer to this question here, but raise it to gesture at the generative potential that a rigorously theorized notion of organization has in revealing unnoticed options in existing debates.

I want to close by noting an opportunity for future work. My proposed condition on organizing remains tethered to the agencies of the individuals in the operands and output insofar as our key notion, causal complementarity, is calculated in terms of the numbers of actions that the organized agents *themselves* regard as means to accomplishing the relevant goal (i.e. G-directed actions). But nor is it impossible for the act of organizing, as my proposed condition has begun to characterize it, to co-opt individuals' agency. And questions remain about how to characterize such co-optation. To see this, notice that the notion of a goal appeared twice in my characterization of organizing. First, it appeared as an argument place in the high-level structure of the organizing event: of any given organizing event, I claimed, it should be possible to state some goal or goals that it was in service of— call this the high-level goal. A mention of goals appeared again in the account of causal complementarity; the claim in O7 was that in an instance of organizing, a set of agents become more causally complementary with respect to some goal. And it was probably

natural to assume that, for any instance of organizing, one goal will fill both rolls; that is, we might assume that it is the high-level goal which the organized system become more causally complementary with respect to. But we should consider whether this will always be the case:

Case 7: Virginia wants to see her alderman reelected. She knows that a proliferation of neighborhood civil society groups during his term is something the alderman will be able to claim as a win come campaign season, even if he didn't have much to do with it. Meanwhile, many neighbors in Virginia's ward all care quite a bit about the park where their children play becoming free of litter, used drug paraphernalia, and graffiti, though they are all either indifferent or hostile to the goal of the alderman getting reelected. Virginia herself doesn't care either way about the cleanliness of the parks, but sees that the neighbors forming an organization to do this is just the sort of thing that would help the alderman. Virginia gets her neighbors together and plays a critical role in coordinating them to clean up local parks.

Call the goal of getting the alderman re-elected g, and the goal of getting the park cleaned up g'. It seems apt to say that Virginia is organizing for g: it is the goal that she undertakes her organizing intervention as a means to. But it just won't be the case that the neighbors are becoming more causally complementary with respect g, because we've stipulated that none of them will have any g-directed intentions (they don't even like the alderman), and where no member of the system has any intentions directed at the accomplishment of a goal, their causal complementarity score with respect to that goal is 0. The neighbors are instead probably getting more causally complementary with respect to g', cleaning up the park. Is it apt to say that Virginia is organizing for g'? Not if this requires that she be intentionally trying to bring about g'. Note that g' is not even instrumental in achieving g—rather, the formation of a neighborhood organization is instrumental in each of them. So while the park getting cleaned up is a possible byproduct of a neighborhood group forming with that as their goal, it's the formation of the neighborhood group that Virginia actually thinks might be conducive to reelecting the alderman, not the park being cleaner.

In summary, there are cases where there's no one goal such that it's both what the operands becomes more causally complementary with respect to and what the organizer intends. Full accounts of organizing which want to render organizing incompatible with cooptation and other forms of coercion or disrespect for individual agency might add a further necessary condition on top of O7 to exclude Virginia's actions as an instance of organizing— and they might do this by availing themselves of the resources provided by my account, stipulating that the goal articulated in the high-level structure and the goal with respect to which individuals become more causally complementary must be the same. Alternatively, if one thinks co-optation is compatible with organizing and is simply one of the risks that is run in allowing oneself to be brought together with other people, I think the resources of my account help articulate this predicament.

References

Saul Alinsky. *Rules for Radicals: A Pragmatic Primer for Realistic Radicals*. Random House, 1971.

Maurizio Antonioli, editor. *The International Anarchist Congress: Amsterdam 1907*. Federazione dei Comunisti Anarchici, 2007.

Jasper Bernes and Joshua Clover. History and the Sphinx: Of Riots and Uprisings. *LA Review of Books*, 2012. URL https://lareviewofbooks.org/article/history-and-the-sphinx-of-riots-and-uprisings/.

Michael E. Bratman. Shared Intention. *Ethics*, 104(1):97–113, 1993. ISSN 0014-1704. URL https://www.jstor.org/stable/2381695. Publisher: University of Chicago Press.

Rachael Briggs. Interventionist counterfactuals. *Philosophical Studies*, 160(1):139–166, August 2012. ISSN 0031-8116, 1573-0883. doi: 10.1007/s11098-012-9908-5. URL http://link.springer.com/10.1007/s11098-012-9908-5.

Ben Burgis. Jimmy Dore Is Right About the Urgency of Medicare for All. But AOC Isn't the Problem. *Jacobin*, 2020. URL https://jacobinmag.com/2020/12/jimmy-dore-aoc-medicare-for-all-pelosi-house-floor-vote-speaker.

Asa Burman. *Nonideal Social Ontology: The Power View*. Oxford University Press, 2023.

Stephanie Collins. Collectives' Duties and Collectivization Duties. *Australasian Journal of Philosophy*, 91(2):231–248, June 2013. ISSN 0004-8402, 1471-6828. doi: 10.1080/00048402.2012.717533. URL http://www.tandfonline.com/doi/abs/10.1080/00048402.2012.717533.

Stephanie Collins. *Group duties: their existence and their implications for individuals*. Oxford University Press, Oxford, first edition edition, 2019. ISBN 978-0-19-884027-5. OCLC: on1083878851.

Democratic Socialists of America. Should House Progressives #ForceTheVote on Medicare for All?, January 2021. URL https://www.dsausa.org/statements/should-house-progressives-forcethevote-on-medicare-for-all/.

Daniel Dennett. *The Intentional Stance*. MIT Press, Cambridge, MA, 1987.

Jimmy Dore. House Progressives Can FORCE a Medicare For All Vote Now – Here's How! *YouTube*, November 2020. URL https://www.youtube.com/watch?v=iIqw-mTX6ro.

Nick Effingham. The Metaphysics of Groups. *Philosophical Studies*, 149(2):251–267, 2010.

Brian Epstein. *The Ant Trap: Rebuilding the Foundations of the Social Sciences*. Oxford University Press, 2015.

Brittany Gibson. How Georgia Got Organized. *The American Prospect*, 2021. URL https://prospect.org/politics/how-georgia-got-organized-stacey-abrams-democrats-elections/.

Margaret Gilbert. *Social Facts*. Princeton University Press, 1989.

Margaret Gilbert. Two Approaches to Shared Intention: An Essay in the Philosophy of Social Phenomena. *Analyse & Kritik*, 30:483–514, 2008.

Margaret Gilbert. *Joint Commitment: How We Make the Social World*. Oxford University Press, 2014.

William Godwin. *Enquiry concerning political justice and its influence on morals and happiness*. J. Watson, 4th edition, 1842.

Briahna Joy Gray. The Case for Forcing a Floor Vote on Medicare for All. *Current Affairs*, 2020. URL https://www.currentaffairs.org/2020/12/the-case-for-forcing-a-floor-vote-on-medicare-for-all.

Virginia Held. Can a Random Collection of Individuals be Morally Responsible? *The Journal of Philosophy*, 67 (14):471–481, 1970. ISSN 0022-362X. doi: 10.2307/2024108. URL https://www.jstor.org/stable/2024108. Publisher: Journal of Philosophy, Inc.

Frank Hindriks. The Duty to Join Forces: When Individuals Lack Control. *The Monist*, 102(2):204–220, 2019. ISSN 0026-9662. URL https://www.jstor.org/stable/26614793. Publisher: Oxford University Press.

Bryce Huebner. *Macrocognition: A theory of distributed minds and collective intentionality*. Oxford University Press, 2014.

Mie Inouye. Antinomies of Organizing. Unpublished manuscript.

Mie Inouye. Starting with People Where They Are: Ella Baker's Theory of Political Organizing. *American Political Science Review*, 116(2):533–546, May 2022. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055421001015. URL https://www.cambridge.org/core/product/identifier/S0003055421001015/type/journal$_a$*rticle*.

Christian List and Philip Pettit. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press, 2011.

Jane F. McAlevey. *No Shortcuts: Organizing for Power in the New Gilded Age*. Oxford University Press, 2016.

D. H. Mellor. *The Facts of Causation*. Routledge, 1995.

Peter Menzies and Huw Price. Causation as a Secondary Quality. *British Journal for the Philosophy of Science*, 44(2):187–203, 1993.

Eugene O. Mills. Interactionism and Overdetermination. *American Philosophical Quarterly*, 33(1):105–115, 1996. Publisher: University of Illinois Press.

Dana Francisco Miranda. The Violence of Leadership in Black Lives Matter. In *The Movement for Black Lives*, pages 243–262. Oxford University Press, November 2021. ISBN 978-0-19-750777-3 978-0-19-750781-0. doi: 10.1093/oso/9780197507773.003.0012. URL https://academic.oup.com/book/39181/chapter/338655590.

Rodrigo Nunes. *Organisation of the Organisationless: Collective Action After Networks*. Mute and Post-Media Lab Books, 2014.

Rodrigo Nunes. *Neither Vertical nor Horizontal A Theory of Political Organization*. Verso, 2021.

Charles M. Payne. *I've Got the Light of Freedom: The Organizing Tradition and the Mississippi Freedom Struggle*. University of California Press, 1995.

Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA, 1988.

Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, 2000.

Judea Pearl and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, New York, 2018.

Vijay Phulwani. The Poor Man's Machiavelli: Saul Alinsky and the Morality of Power. *American Political Science Review*, 110(4):863–875, November 2016. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055416000459. URL https://www.cambridge.org/core/product/identifier/S0003055416000459/type/journal$_a$*rticle*.

Barbara Ransby. *Ella Baker and the Black Freedom Movement : A Radical Democratic Vision*. University of North Carolina Press, 2003.

Katherine Ritchie. What are groups? *Philosophical Studies*, 166(2):257–272, November 2013. ISSN 0031-8116, 1573-0883. doi: 10.1007/s11098-012-0030-5. URL http://link.springer.com/10.1007/s11098-012-0030-5.

Katherine Ritchie. The metaphysics of social groups. *Philosophy Compass*, 10(5):310–321, 2015. Publisher: Wiley Online Library.

Katherine Ritchie. Social structures and the ontology of social groups. *Philosophy and Phenomenological Research*, 100(2):402–424, 2020. Publisher: Wiley Online Library.

Andrea Sangiovanni. Solidarity and Joint Action. *Journal of Applied Philosophy*, 32(4):340–359, 2015.

Jonathan Schaffer. Of Ghostly and Mechanical Events. *Philosophy and Phenomenological Research*, 68(1):230–244, 2004. doi: 10.1111/j.1933-1592.2004.tb00338.x. Publisher: Wiley-Blackwell.

David Sirota. Five Things House Progressives Can Do to Advance Medicare for All. *Jacobin*, 2020. URL https://jacobinmag.com/2020/12/house-progressives-medicare-for-all-m4a-pelosi.

Scott Sturgeon. Physicalism and Overdetermination. *Mind*, 107(426):411–432, 1998. ISSN 00264423, 14602113. URL http://www.jstor.org/stable/2659884. Publisher: [Oxford University Press, Mind Association].

Deborah Perron Tollefson. *Groups as Agents*. Polity Press, 2015.

Gabriel Uzquiano. The Supreme Court and the Supreme Court Justices: A Metaphysical Puzzle. *Noûs*, 38(1): 135–153, 2004.

Georg Henrik von Wright. *Explanation and Understanding*. Cornell University Press, Ithaca, NY, 1971.

Deva R. Woodly. *Reckoning: Black Lives Matter and the Democratic Necessity of Social Movements*. Oxford University Press, 2021.

James Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, New York, 2003.

Sewall Wright. The relative importance of heredity and environment in determining the piebald pattern of guinea-pigs. *Proceedings of the National Academy of Sciences of the United States of America*, 6:320–332, 1920.