

Matteo Bianchin

Collective deliberation as epistemic cooperation

Abstract

I suggest that collective deliberation processes should be seen as cooperative epistemic activities. After considering a well-known argument for why they cannot result from a majoritarian aggregating procedure, I focus on the limitations of a functionalist approach to collective agency and suggest that Tomasello's approach to cooperation can shed light on how collective deliberation works. I then argue that understanding collective deliberation in terms of epistemic cooperation sheds light on both the structure and the normative implications of collective deliberative processes.

Keywords

Collective Deliberation, Cooperation, Group Agency, Group Rationality

Received: 01/03/2025

Approved: 12/05/2025

Editing by: Massimo Caon

© 2025 The Author. Open Access published under the terms of the CC-BY-4.0.
matteo.bianchin@uniroma2.it (Università di Roma Tor Vergata)

1. *Group agency and collective deliberation*

We routinely do things together. We participate in groups, play a variety of institutional roles, cooperate on a regular basis. When we do, the relevant intentions and actions are commonly expressed in the first-person plural and reported in the third-person singular, in particular when attributed to institutions. We moved next class to Wednesday, the BCE raised the interest rates, the Catholic Church became more tolerant about homosexuality, yet didn't change its minds about same sex marriage, and so on. Both common sense and the social sciences treat groups and institutions as corporate agents. As long as agency requires deliberation, however, the literature on collective action has been faced with the problem of whether and how the latter can be credited to collectives (Searle 1995, Tomasello 2008, List-Pettit 2011, Gilbert 2014). If we are to take joint and collective actions as something people perform on a regular basis, it looks like the psychology of deliberative processes should be located at the corporate level, crediting groups and institutions with a mind of their own.

It might be tempting to reduce corporate agents to collections of individuals that make up their mind by a majoritarian voting process. Yet a well-known dilemma blocks this move. Suppose an advisory board of three must judge whether global warming will be significant, given present carbon dioxide emissions. To draw a reasoned conclusion, they should make up their mind about whether gas emissions are above a critical threshold (first premise) and whether a significant global warming will follow from their being above that threshold (second premise). Now suppose the first expert takes both the first and the second premise to be true, concluding that there will be a significant global warming. The second expert takes instead the first premise to be true, but not the second, and concludes that there will be no significant global warming. The third expert finally takes the second, but not the first premise to be true, and concludes again that there will be no significant global warming. If we aggregate their attitudes by a majoritarian procedure, it turns out that there is a majority for each of the premises that favor their being true, yet the majority about the conclusion is that there will be no significant global warming. Thus, although individual attitudes are fully consistent, the majoritarian set of attitudes is inconsistent. The upshot is that majority voting may lead to inconsistent group attitudes even if individual attitudes are consistent, failing to express a rational deliberation process (Pettit

2001, List 2006, List-Pettit 2011: 45-6). This result can be pictured as follows:

	Individual 1	Individual 2	Individual 3	Majority
Premise 1	True	True	False	True
Premise 2	True	False	True	True
Conclusion	True	False	False	False

The discursive dilemma is taken to show that majoritarian aggregation conflicts with group agency because, given that group attitudes supervene on individual attitudes, majoritarian aggregation makes groups to hold inconsistent attitudes and thus fails to account for how groups can be rational (List-Pettit 2011: 67 ff.). Rationality can be preserved, however, if groups are differently organized, for instance by a premise-based procedure according to which the attitudes of the group on both premises are set by majority voting, but the group attitude on the conclusion is set by what follows from the accepted premises rather than by majority voting. In this case, the group ends up entertaining a consistent set of attitudes. Of course, a premise-based procedure entails dropping the view that the group attitude on each proposition is determined by individual attitudes on that proposition, because the group attitude on the conclusion is determined by the group attitudes on the premises rather than by the attitudes individuals entertain on the conclusion. Taken as a whole, however, group attitudes still supervene on individual attitudes, because the group attitudes towards the premises supervene on the individual attitudes towards the premises, so that individual attitudes fully determine group attitudes towards all propositions, conclusion included. In other terms, while group attitudes do not supervene on individual attitudes one by one, they nevertheless supervene “holistically” on individual attitudes, because fixing individual attitudes fixes group attitudes anyway (List-Pettit 2011: 69). Which entails that individual and group attitudes can come apart without violating supervenience, because the group attitude towards the conclusion is set whatever the individual attitudes towards the conclusion are. In short, individual attitudes towards premises are sufficient to fix group attitudes towards both the premises and the conclusion, but individual attitudes towards the conclusion are neither sufficient, nor necessary to determine the group attitude towards the conclusion.

The upshot is that preserving group agency requires crediting groups with some rational autonomy, granting collectives a mind “of their own” (Pettit 2011). Both deliberative processes and their outputs must be attributed to the whole group, however individuals contribute to determine them, because they largely depend on how the group is organized, very much as individual cognitive systems depend both on their parts and on how they are organized into a whole (Bechtel 2008: 149 ff.).

2. Collective mindedness and plural subjectivity

The discursive dilemma shows that group agency requires collectives to exhibit at least some rationality. Yet there are different ways to conceive how they can be rational. List and Pettit favor a functionalist approach on which crediting groups with mental states and deliberation processes best explains their behavior. This has been a standard argument in the philosophy of mind. Anything that count as an agent must be credited with representational states that depict how things are, motivational states that specify how the agent intends them to be, and a capacity to process both to bring about actions that satisfy the latter according to the former. Consequently, a system must be taken to be an agent if it is impossible to dispense with the “intentional stance” in explaining its behavior, that is if the best explanation for its observed behaviors is an intentional explanation that makes sense of them by attributing mental states and deliberation processes to the system as a whole – be it an individual, a group, or an institution (List-Pettit 2011: 23).

While this argument is generally taken to support realism about mental states and processes (Fodor 1985, Pettit 2014), it has two limitations. For one thing, the fact that a functionalist approach licenses the analogy between group minds and individual minds can cast doubt about intentional explanations tout court, because it also licenses the reverse claim that folk psychology is an illusion for individual agents as well as for corporate agents like governments and companies. Consider the policy statements issued by corporate entities like governments and companies. They rationalize past behaviors and express commitments that allow predicting future behaviors, but one can be dubious about the idea that they describe actual inner causes. The prediction and retrodiction they allow about the behavior of corporate entities are likely vindicated because the same forces within a corporation that produce the policy statements also produce the policy, not because the latter describe any actual

psychological process (Clark 1994: 419). Given the analogy between individual and corporate agents, reductionists may thus argue that the attributions of mental states to people are very much like official policy statements: they rationalize past behaviors and express putative commitments to future behaviors, but do not describe actual causes (cf. also Mercier-Sperber 2017):

We have official spokespersons so that governments can better predict one another's actions [...] Ordinary psychological terms evolved for similar roots and serve similar functions. Folk psychology is an interpretation manual for such signal system. It too provides a running commentary, a story line that has some predictive value. But when we seek explanations of why the system acted as it did, we inevitably endorse reduction, and seek the inner constituent that jointly conspired to generate the result. (Clark 1994: 222)

A second trouble concerns subjectivity. First personal self-awareness is routinely taken to be a condition for agency, yet functionalism is silent about the first person. The intentional stance works in the third person. Group actions, however, are routinely deliberated and carried out as joint actions in the first-person plural. From an agential point of view, “we” are those who jointly act for reasons we share. If we look at collective deliberation processes from an agential point of view, then, there seems to be some intuitive constraints on what can count as collective actions and deliberations. First, any one of a random collection of individuals may well undergo identical psychological processes as the others, yet this is not enough for the relevant processes to be jointly performed. Tuomela therefore advanced a “collectivity condition” that requires people to reason and act in the “we-mode” (Tuomela 2013: 24). Which means that sharing a goal is not per se enough for a collection of individuals to count as a collective. In order for them to work in the we-mode, their thoughts and actions must be individuated by the group perspective, so that they have a common goal rather than just a goal in common. Sport teams or symphonic orchestras are good examples: the reason why players perform their individual actions is the role these actions play in a collective action plan that can only be grasped from the point of view of the group – winning a game, playing a symphony, and so on (Tuomela 2013: 40 ff.). A further constraint seems to follow. Collective deliberation and action processes entail more than common knowledge, as common knowledge just amounts to the iterated second personal knowledge that “I know that you know that I know that I know that ...”, whereas group agency entails that groups are credited with mental states and processes expressed by

the first-person plural “we” (Schwengerer 2022). Which apparently entails a final constraint. Collective agency requires crediting groups with a sense of “we-ness” that mirrors individual agential self-awareness (Tommasello 2008: 57, Bianchin 2015a, Crone 2021).

Collective actions and deliberation processes can thus be hardly accounted by a minimalistic planning theory of action that makes joint intentions to result from interdependent individual action plans. On this view the fact that “we” intend to perform joint action J reduces to the fact that 1) each intend that we J – e.g. go to the movie, 2) each intend that we J in virtue of (1) – i.e. because each intend that we go to the movie, and 3) both (1) and (2) are common knowledge among participants – i.e. each knows that each other knows that each knows that (1) and (2) (Bratman 2014: 84). And there are two major problems with this. First, the account looks circular, as a first-person plural still occurs in the propositional content of individual intentions “that we J” (Tuomela 2005, Schweikard-Schmid 2021: 18). Second, common knowledge expresses the infinitary second personal knowledge that “I know that you know that I know that you know ... and so on ...”, rather than the first plural personal knowledge that “we J”, and there seems to be no way in which iterating the second personal attribution of propositional attitudes to other people can generate the plural first personal perspective expressed by first-person-plural intentions (Searle 1995: 25). It looks like collective action and deliberation processes must be modelled to preserve plural subjectivity.

Notice, however, that acknowledging that collective mindedness cannot be reduced still does not meet the demand for plural first personal self-awareness that is placed on group agency. Searle’s social ontology, for instance, is adamantly individualistic in this respect, because it takes collective intentionality to be irreducible, yet biologically primitive (Searle 1995: 24). As biological properties are intrinsic properties of individual organisms, collective intentionality is also an individual property. Moreover, Searle’s internalism entails that collective intentionality does not depend on any actual relations among people, because no factor external to the individual contributes to determine mental states. Thus, although “no sets of I-Consciousness” can possibly add up to a “We-Consciousness” (Searle 1995: 24), the latter cannot be a property of plural subjects either. The claim that collective intentionality cannot be reduced just amounts to the claim that mental states expressed by the first-person plural cannot be analyzed into mental states expressed by the first-person singular:

It is indeed the case that all my mental life is inside my brain, and all your mental life is inside your brain, and so on for everybody else. But it does not follow from that that all my mental life must be expressed in the form of a singular noun phrase referring to me. The form that my collective intentionality can take is simply “we intend”, “we are doing so-and-so,” and the like. In such cases, I intend only as part of our intending. The intentionality that exists in each individual head has the form “we intend.”? (Searle 1995: 25-6)

On Searle’s reading, the we-mode is therefore nothing but a feature of individual attitudes like beliefs, desires, intentions, and so on, for which the claim that entertaining psychological attitudes in the we-mode expresses a primitive biological capacity provides no further explanation. Searle’s social ontology is consequently silent about the mechanisms that underly we-ness. We-ness is posited as a primitive feature of collective intentionality on conceptual grounds, because it cannot be analyzed in terms of individual intentionality, but no further thought is devoted to considering whether and how it can possibly emerge from biological or psychological mechanisms that do not involve reducing collective attitudes to individual attitudes.

Mutatis mutandis, similar considerations hold for Gilbert’s normative approach. On Gilbert’s view, groups are constituted as plural subject by a joint commitment to act and think “as one”, where joint commitments are primitive in that they cannot be unilaterally rescinded and therefore cannot be reduced to any set of individual commitments on pain of reducing their normativity to the normativity of individual commitments, which may well be unilaterally rescinded. Joint commitments therefore need “a joint author” to create a collective obligation the parties cannot unilaterally rescind (Gilbert 2023: 85 ff.). Gilbert normative approach unquestionably departs from Searle’s internalism because joint commitments are relational – two or more people need be involved. When it comes to locate the author of joint commitments, however, a paradox looms large. On the one hand, joint commitments cannot be authored by the collection of individuals that participate in the group, because a mere collection of individuals cannot count as a joint author. On the other hand, they cannot be authored by groups themselves, because groups are constituted by joint commitments and therefore cannot be brought into existence by a joint commitment they should themselves set up – they are just not around prior to the joint commitment that brings them about. In other words, making joint commitments to result from mutual individual commitments would dispense with irreducibility, yet preserving the latter requires groups to author the joint commitments that brings them

about – which of course they can't. Which blurs the very idea that groups are plural subjects of irreducibly collective thought and action. Positing we-ness as a primitive biological or normative feature of collective mindedness just points to a blind spot in current theories of group agency.

3. *Phenomenology and beyond*

The reason why subjectivity is a blind spot in group agency is arguably that corporate agents lack the unity of consciousness individuals enjoy (Bayne 2010, Schechter 2013, List 2018). Groups do not possess a phenomenal self-awareness of their own. Therefore, we-ness must be explained on different grounds. Phenomenological approaches indeed routinely convey the view that a sense of plural first personal we-ness or “for-us-ness” emerges from intersubjective we-relations that crucially involve a peculiar kind mutual pre-reflective second personal self-awareness as a basis for reciprocity and interdependence, which is prior to any predicative attribution of propositional attitudes to others and rather involves an “interlocking” of subjective experiences such that each individual agent is pre-reflectively aware of her own experience as being experienced by others, where such awareness involves that others are similarly self-aware of their own experience as experienced by each (Leon 2020, Zahavi 2023). While we-relations are not already joint, cooperative activities, as they do not per se entail sharing a common goal, they provide the cognitive underpinnings of cooperation by equipping agents with a mutual experiential access to each other that constitutes “the pure spere of the We”, in which “We participate without an act of reflection” (Schutz 1962: 175). From a phenomenological point of view, we-relations thus account for how we-ness emerges and enables people to function as a group in deliberating what to do. While plural self-awareness cannot rely on collective phenomenal consciousness, cognitive processes that involve a plurality of individuals can, on this basis, display enough integration to express a shared mental architecture (Szanto 2014).

Tomasello's approach to collective action and thought is likely the most developed attempt to follow a similar path in current debates. On the one hand, it takes we-mode thought and agency to entail a sense of “we-ness” that conveys the perspective of a plural subject of thought and action. On the other hand, it takes the latter to arise from a general capacity for sharing intentions that builds on joint attention and social cognition (Tomasello-Rakoczy 2003, Tomasello 2008: 67 ff., Tomasello 2014:

80-1). Joint attention enables sharing actions and thoughts because it involves a “dual-level attentional structure” with a higher-level focus on shared goals and a lower-level focus on the complementary roles agents must perform to achieve the goal, so that both are represented in a single format (Tomasello 2008: 68-70). Engaging in joint attention, however, requires agents to possess two underlying abilities that draw on social cognition. First, they must be able to engage in mutual perspective taking, or “role reversal”, which entails understanding the equivalence between self and other. Second, they must be able to integrate different perspectives into a we-centric representation of actions and goals (Tomasello 2008: 65). The mechanisms that explain how we-ness emerges must therefore be located at the psychological level in a capacity for social cognition that allows agents to take up and merge each other’s perspective.

Tomasello does not explicitly endorse a specific theory of social cognition, but it seems clear that the latter is best understood in this context along the lines of a simulation theory of mindreading (Bianchin 2015a: 446-7). According to Tomasello, social cognition relies indeed on a capacity to “simulate the perspective of others” that underlies cultural learning because it enables individuals to imagine themselves in the mental shoes of others and reenact their reasoning processes (Tomasello-Kruger-Ratner 2003: 503, Tomasello 1999, 5-6, 30, 71 ff.). Mentalizing thus involves engaging in a mental simulation to devise a scenario “as one would experience or undergo it if it were currently happening” (Goldman 2013, 102). As we engage in social cognition, we are consequently induced to “re-center our egocentric maps” to reenact one another’s cognitive processes in a way that “multiplies the first person” (Gordon 1995a, Gordon 1995b, Hurley 2008, Goldman 2006: 27 ff.). Multiple perspectives can thereby be brought to overlap to engender a we-centric representation of actions and goals that displays what is like for “us” to think and act as a group (Bianchin 2015a: 453).

The analogy between interpersonal cooperation and intrapersonal action planning may help understanding how mental simulation enables integrating multiple actions in a single action scheme across times and persons. Carrying out a plan requires keeping track of past actions and figuring out future actions to coordinate what to do at different times, which entails internally merging different perspectives towards the goal across time as mine. Taking a joint action involves performing a similar task interpersonally. To cooperate towards a shared goal, we need merging different perspectives across people as ours. A common mechanism for self-projection underlies both tasks. Thinking about the future, recalling the

past, and mentalizing others all involve imaginatively shifting our perspective from actual to counterfactual scenarios to simulate what it would be like to undergo the latter, which is consistent with existing evidence that they share a common functional anatomy and the underlying brain networks underpin handling simultaneously multiple perspectives (Bruckner-Carroll 2007: 50-3, Gallagher-Frith 2003: 79, Hurley 2008: 19-20). We-centric representations of actions and goals would then result from merging different perspectives into a double layered representation in which the joint action as a whole is self-attributed by group members to themselves as a group in the first-person plural, while related partial actions are attributed to individual agents.

The capacity to engage in joint actions emerges early in childhood to support small-scale cooperative activities based on “dyadic” second personal relationships, but the underlying mechanism for social cognition support cultural learning, which later allows developing a set of broader social conventions and norms to manage large-scale cooperation (Tomasello 2014: 5, Tomasello 2009: 73 ff., 100 ff., Tomasello 1999: 56 ff.). As group size increases, a new kind of group-mindedness consequently emerges that is based on sharing a cultural common ground. “We” become “an enduring culture” which extends across space and time to whoever identifies with the group by sharing the relevant cultural practices (Tomasello 2014: 83-5). Collective deliberation is thus reconfigured into a “cooperative argumentation” that drives group decision making in accordance with shared norms of rationality and allows individuals to reason from the group’s perspective:

In the context of cooperative argumentation in group decision making, linguistic conventions could be used to justify and make explicit one’s reasons for an assertion within the framework of the group’s norms of rationality. This meant that individuals now could reason “objectively” from the group’s agent-neutral point of view (“from nowhere”). (Tomasello 2014: 5, see 138 ff.)

On this reading, people are not given a primitive “we-consciousness”, but come to endorse a plural first personal stance by merging multiple perspectives into a single we-centric representation of actions and goals they are self-aware of as expressing the perspective of a plural subject whose deliberative processes are structured by shared epistemic norms and enacted by individual agents according to the role they play in a cooperative argumentation.

Thus, while Tomasello's framework accounts for the phenomenology of we-ness in terms of the plural self-awareness of individual agents in enacting a plural subject, it goes beyond phenomenology in articulating group reasoning as a cooperative argumentation. In this respect, his approach rather converges with Habermas' view of communicative action and arguably accounts for how communicative action itself emerges from earlier pre-linguistic capacities for social cognition (Bianchin 2015b). Indeed, both Habermas and Tomasello understand communication as a cooperative interaction that cannot be generated by strategic rationality. Tomasello explicitly argues that agents endowed with strategic rationality alone cannot develop communicative conventions because they are unable to cooperate (Tomasello 2009: 98 ff., Tomasello 2008: 62 ff.). Habermas likewise takes communicative action to merge individual action plans into a social interaction scheme that cannot be reduced to the mere interlocking of individual action plans:

The concept of communicative action is presented in such a way that the acts of reaching understanding, which link the teleologically structured plans of action of different participants and thereby first combine individual acts into an interaction complex, cannot themselves be reduced to teleological actions. (Habermas 1987: 288)

As a matter of fact, Habermas only diverges from Tomasello in taking joint attention, social cognition, and conventional symbols to co-originate (Habermas 2017: 31 ff), whereas Tomasello's account relies on attentional, perceptual and agential capacities that predate language and enable early communicative conventions (Tomasello 2009: 4). This difference fades away, however, as children acquire – around 4 – the linguistic tools to manage the folk psychological vocabulary of propositional attitudes, which allows developing a general theory of mind and consequently expand their capacity for social cognition beyond dyadic, second personal relationships, to enable partaking in cultural institutions and large-scale cooperation schemes based on collective beliefs and practices (Tomasello-Rakoczy 2003: 136-9, De Villiers 2007, Hacquard-Lidz, 2019). As mentioned, joint attention requires people to engage in mutual perspective taking, which entails understanding the equivalence between self and other (Tomasello 2008: 71-2). As the related capacities for social cognition and cultural learning develop, the self-other equivalence accordingly generalizes beyond dyadic relations, enabling individuals to operate in anonymous settings under a general constraint of rationality and

reciprocity, which allows them accessing objective reasons and guiding their actions via the impersonal authority a ‘generalized other’ (Tomassello-Rakoczy 2003: 139, Bianchin 2015b: 599 ff.). It is in this context that deliberative processes become “cooperative argumentations”.

4. *Collective deliberation as epistemic cooperation*

Understanding collective deliberations as cooperative argumentations allows modeling them as cooperative activities aimed at forming beliefs and making decisions that approximate some shared standard of truth, justice, and the like. Cooperative argumentations are linguistically articulated forms of epistemic cooperation in which people engage to achieve shared epistemic goals. Which makes sense of how groups can be rational without crediting them with a metaphorical mind “of their own”, as a cooperative argumentation is nothing but a communicative process in which people jointly make up their minds about what to do by exchanging arguments. The discursive dilemma mentioned in the first section evaporates, because people can be naturally expected to accept conclusions that follow from shared premises.

Understanding collective deliberation as epistemic cooperation thus strengthen the plausibility of a deliberative approach to collective decision making because it accounts for how deliberation can track group rationality. Deliberation is commonly taken to benefit collective decision making because it enables integrating different perspectives, interpretations, and information pools, it allows weeding out good arguments from bad, and it ultimately leads to rational agreements that reliably track the best solution (Landemore 2013: 97 ff.). More specifically, it is common ground that engaging in a deliberative process affects people’s preferences, views, judgements, and social dispositions so as to improve collective decisions because it (a) exposes people to new information, perspectives, and interpretations (b) confronts people with new arguments and allows subjecting them to public scrutiny, (c) encourages people to take a reflective stance on their preferences and perspectives in the knowledge that they have to be justified to others, (d) creates a dialogic situation that elicits people’s cooperative dispositions and induces participants to devise the issue at stake in a “we-frame” (List-Dryzek 2003: 9). If collective deliberations are conceived as cooperative epistemic activities, the structure and the normative features of deliberative processes

can be analyzed within the framework sketched above for understanding cooperation.

First, understanding deliberation processes as cooperative activities allows outlining how they are structured by articulating the shared mental architecture they involve. Here Tomasello's idea that cooperation involves a double-layered representations of actions and goals – with a higher-level focus on shared goals and a lower-level focus on individual roles – explains how agents manage to reason together. For something to count as a collective deliberation, agents must understand their giving and asking for reasons as the complementary roles they play in a cooperative activity designed to reach an epistemic goal they jointly represent as their own. Moreover, complex tasks may require groups to be organized so that different individuals and sub-groups specialize in performing different tasks. Take advisory boards or large corporate decision makers like economic organizations and public institutions. Sometime the deliberative process is the same for all group members, so that they are all equally involved in one and the same process. But suppose a large amount of data must be gathered, analyzing them requires specialized knowledge, and reasoning about different aspects of the issue at stake requires different cognitive skills – economic, legal, sociological, psychological, and so on. In such cases, the overall task is routinely broken down into different components which are fed to different individuals or sub-groups, so that group membership entails performing a specific function within the group, or within a sub-group. Also, some tasks may be performed by technological devices – information and arguments can be stored in written records, disseminated by different media, processed by computer software, and so on. The division of cognitive labor allocates cognitive tasks across the group, so that the group itself works as a distributed cognitive system whose parts cooperate to the overall deliberation by performing different specialized tasks whose outputs are available to the group as a whole (Bird 2014: 80 ff., 2024). The deliberation process here can be suitably regarded as a culturally scaffolded cooperation scheme that still display enough cognitive integration to devise a shared mental architecture.

Second, conceiving collective deliberation as epistemic cooperation provides some tools to tackle the normative demands placed by large-scale deliberations on a deliberative view of collective decision making. Large-scale deliberation is realistically described as performances of deliberative systems that involve multiple sub-groups, organizations, institutions, networks, and technologies. As a consequence, dialogic

approaches look descriptively inadequate and normatively vulnerable to the criticism that they convey an idealized view of deliberation according to which deliberation processes are intrinsically democratic and fair. The apparent shortcoming of systemic approaches, however, is weakening the claim that people reason “together”, because on a systemic approach the deliberative process turns out to be fragmented and mediated by a variety of mechanisms that go beyond dialogic reason-giving, which seems to refute the participatory ideals commonly associated with a deliberative approach to collective decision making. This is not the place to enter the debate about deliberative systems, but it is worth noting that understanding collective deliberation as epistemic cooperation suggests how they can be designed to be democratic and fair.

Democracy has received considerable attention, and a variety of measures have been advanced to “democratize deliberative systems” (Parkinson 2012). Democratic equality can be implemented as an organizational feature of the system, by granting people equal access to deliberations within sub-groups and allowing sub-groups to be equally represented in larger groups, so that each perspective ends up being treated equally on the whole, even if not all individuals participate equally in every segment of the overall deliberation process (Bohman 2012: 75 ff.). Moreover, cultural technologies like education, the mass media, and science can be designed to reduce the systemic distortions induced by self-serving interests, cognitive biases, and ideologies that can pervasively operate beyond individual control in culturally scaffolded deliberation processes. Finally, social institutions can be designed to minimize the influence of power relations within the system and be anchored in a broad public sphere that supports reflective critique and contestation (Parkinson 2012: 161 ff.). The basic idea that underlies such proposals is that, while large-scale deliberations require a division of epistemic labor that breaks the deliberative process into different tasks and bypasses dialogic control, democratic equality can be preserved by cultural design. Here again Tomasello’s framework helps understanding how this might possibly work, as it suggests that the self-other equivalence built into the psychological infrastructure of cooperation can itself be culturally scaffolded – for instance, through institutional design – so that individuals still count as free and equal ultimate justificatory authorities across the system, as required by democratic deliberations (Forst 2017: 29, 49, 157).

How deliberative systems can be fair has received far less attention, although it plainly relates to well-known debates about epistemic injustice. Epistemic injustice occurs when people are wronged either because

they experience a “credibility deficit” due to a prejudice associated with the social group(s) they belong to or because they are deprived of the hermeneutical resources to make sense of their own social experience due to their being excluded from the social practices in which the relevant conceptual, linguistic, and interpretive tools are created and reproduced – e.g. from medical, legal, educational, and political practices (Fricker 2007: 27, 151). Understanding collective deliberations as cooperative practices helps understanding what is wrong with that, as it suggests setting the standards of epistemic justice along the lines of a general theory of justice designed to govern cooperation. The idea of a “well-ordered science” has been advanced by Philip Kitcher, for instance, to convey the idea that scientific practice should be designed to meet the different viewpoints distributed in a society about what problems are significant, which values should be met, and so on (Kitcher 2001: 122-3). Understanding deliberation as a cooperative practice, however, allows outlining a specific conception of epistemic justice that applies to any deliberative domain. If deliberations are cooperative practices, it should indeed be possible to specify some general principles of epistemic justice along the lines of Rawls’s conception of justice (Rawls 1999: 266) to devise how the burdens and benefits of epistemic cooperation should be distributed wherever collective deliberations are structured by the division of epistemic labor. It is unlikely that people enjoy identical epistemic resources in such cases, and not all differences are likely unjust, but deliberations would look intuitively fair if epistemic inequalities are arranged to the greatest benefit of the least epistemically advantaged and attached to roles that are open to all under conditions of a fair equality of opportunity. Epistemic justice thus would arguably allow for differences that can be justified as arising under conditions of a fair access to epistemic resources – allowing experts, for instance, to enjoy a special epistemic authority in their field – but would prevent arbitrary differences and lay out a standard that might shed light on epistemic injustice as well as it might inform institutional design.

What about subjectivity? Bird (2014: 102 ff.) suggests that the unity of a subject materializes in distributed cognitive systems due to the interdependence engendered by the division of epistemic labor. On this reading, a plural subject of action and thought emerges from the division of epistemic labor as individual contributions are integrated in a shared cognitive architecture that encompasses both the overall task performed by the system and the sub-tasks performed by its parts. Deliberative systems, however, are not phenomenally conscious and therefore cannot be

credited themselves with the subjective self-awareness agency requires. A phenomenal sense of we-ness can be preserved only along the lines of Tomasello's account of how people identify with cultural groups in enacting large-scale, culturally scaffolded collective cognitive processes. As mentioned earlier, experiencing shared actions as actions "we" perform requires individual agents to manage a double-layered representation of actions and goals, involving a higher-level focus on the joint task and a lower-level focus on related complementary roles. As cooperation scales up, cultural practices replace dyadic interactions in endowing agent with a we-centric perspective that allows reasoning from the point of view of the group. On this view, deliberative systems are just complex collectives whose organizational features are shaped by cultural technologies. On the one hand, the division of epistemic labor is structured by social norms associated with different roles people may come to play in a deliberative system. On the other hand, the production and transmission of knowledge across the whole system is supported by a technological infrastructure that includes both material artefacts and abstract methods. Large-scale argumentative cooperation goes therefore well beyond the minimal turn-taking structure of dialogic exchanges, but there is no reason to think that cultural complexity prevents people from experiencing what they do as their own deliberation. In a democratic society, for instance, electoral processes are routinely experienced as collective deliberations performed by a people, however complex, agonistic, fragmented, and technologically mediated they may turn out to be. In this sense, while deliberative systems cannot be themselves phenomenally self-aware, they can be experienced as plural subjects by those who enact them as long as they are subjectively self-aware of acting as a culturally enhanced group. We the people are culturally scaffolded plural subjects.

Acknowledgments

This research is part of the research project PRIN PNRR 2022 Transforming Habits [TRANSHAB]. Agency, Cognition, and Social Change, funded by the European Union-Next Generation EU, Mission 4, Component 1, CUP E53D23018750001.

Bibliography

- Bayne, T., *The Unity of Consciousness*, Oxford, OUP, 2010.
- Bechtel, W., *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*, London, Routledge, 2008.
- Bianchin, M., *Simulation and the We-Mode: A Cognitive Account of Plural First Persons*, "Philosophy of the Social Sciences", n. 4-5 (2015a), pp. 442-61.
- Bianchin, M., *From Joint Attention to Communicative Action: Some Remarks on Critical Theory, Social Ontology, and Cognitive Science*, "Philosophy and Social Criticism", n. 6 (2015b), pp. 593-608.
- Bird, A., *Knowing Science*, Oxford, OUP, 2014.
- Bird, A., *Naturalized Knowledge-First and the Epistemology of Groups*, "Philosophy and Phenomenological Research", n. 3 (2024), pp. 856-973.
- Bohman, J., *Representation in the Deliberative System*, in eds. J. Parkinson, J. Mansbridge, *Deliberative Systems: Deliberative Democracy at Large Scale*, Cambridge, CUP, 2012, pp. 72-94.
- Bratman, M., *Shared Agency: A Planning Theory of Acting Together*, Oxford, OUP, 2014.
- Clark, A., *Beliefs and Desires Incorporated*, "The Journal of Philosophy", n. 8 (1994), pp. 404-25.
- Bruckner, R., Carroll, *Self-Projection and the Brain*, "Trends in Cognitive Sciences", n. 2 (2007), pp. 49-57.
- Crone, K., *Foundations of a We-Perspective*, "Synthese", n. 2 (2021), pp. 11815-32.
- De Villiers, J., *The Interface of Language and Theory of Mind*, "Lingua", n. 11 (2007), pp. 1858-78.
- Dryzek, J., List, C., *Social Choice Theory and Deliberative Democracy: A Reconciliation*, "British Journal of Political Science", n. 1 (2003), pp. 1-28.
- Fodor, J., *The Fodor's Guide to Mental Representations*, "Mind", n. 373 (1985), pp. 76-100.
- Forst, R., *Normativity and Power: Analyzing Social Orders of Justification*, Oxford, OUP, 2017.
- Fricker, M., *Epistemic Injustice: Power and the Ethics of Knowledge*, Oxford, OUP, 2007.
- Gallagher, H., Frith C., *Functional Imaging of "Theory of Mind"*, "Trends in Cognitive Sciences", n. 2 (2003), pp. 77-83.
- Gilbert, M., *Joint Commitment: How We Make the Social World*, Oxford, OUP, 2014.
- Gilbert, M., *The Life of Groups: How Do We Think, Feel, and Act Together*, Oxford, OUP, 2023.
- Goldman, A., *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*, Oxford, OUP, 2006.

- Goldman, A., *Joint Ventures: Mindreading, Mirroring, and Embodied Cognition*, Oxford, OUP, 2013.
- Gordon, R., *Simulation Without Introspection or inference from me to you*, in eds. M. Davies, D. Stones, *Folk psychology*, Oxford, Blackwell, 1995a, pp. 53-67.
- Gordon, R., *Sympathy, Simulation, and the Impartial Spectator*, "Ethics", n. 4 (1995b), pp. 727-42.
- Hacquard, V., Lidz, J., *Childrens' Attitude Problem: Bootstrapping Verb Meaning From Syntax and Pragmatics*, "Mind and Language", n. 1 (2019), pp. 73-96.
- Hurley, S., *The Shared Circuits Model (SCM): How Control, Mirroring, and Simulation Can Enable Imitation, Deliberation, and Mindreading*, "Behavioral and Brain Sciences", n. 1 (2008), pp. 1-58.
- Habermas, J., *The Theory of Communicative Action. Vol. I: Reason and the Rationalization of Society*, Boston, Beacon Press, 1987.
- Habermas, J., *Postmetaphysical Thinking II*, Cambridge, Polity Press, 2017.
- Kitcher, P., *Science, Truth, and Democracy*, Oxford, OUP, 2001.
- Kriegel, U., *The Three Circles of Consciousness*, in eds. M. Garcia-Carpintero, A. Billon, *Self-Experience*, Oxford, OUP, 2016, pp. 169-91.
- Kriegel, U., Zahavi, D., *For-Me-Ness: What It Is and What It Is Not*, in eds. D. Dahlstrom, A. Elpidorou, W. Hopp, *Philosophy of Mind and Phenomenology*, London, Routledge, 2016, pp. 36-53.
- Landemore, H., *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*, Princeton, Princeton University Press, 2013.
- Leon, F., *For-Me-Ness, For-Us-Ness, and the We Relation*, "Topoi", n. 3 (2020), pp. 547-58.
- List, C., *The Discursive Dilemma and Public Reason*, "Ethics", n. 2 (2006), pp. 362-402.
- List, C., *What is Like to Be a Group Agent*, "Nous", n. 2 (2028), pp. 295-319.
- List, C., Pettit, P., *Group Agency: The Possibility, Design, Status of Corporate Agents*, Oxford, OUP, 2011.
- Mercier, H., Sperber, D., *The Enigma of Reason*, Cambridge (MA), Harvard University Press, 2017.
- Parkinson, J., *Democratizing Deliberative Systems*, in eds. J. Parkinson, J. Mansbridge, *Deliberative Systems: Deliberative Democracy at Large Scale*, Cambridge, CUP, pp. 151-72.
- Pettit, P., *Deliberative Democracy and the Discursive Dilemma*, "Philosophical Issues", n. 1 (2001, supplement), pp. 268-95.
- Pettit, P., *Groups with Minds of Their Own*, in eds. A. Goldman, D. Whitcomb, *Social Epistemology: Essential Readings*, New York, Oxford University Press, 2011, pp. 167-93.
- Rawls, J., *A Theory of Justice* (1971), Cambridge (MA), Harvard University Press, 1999.

- Pettit, P., *Group Agents are Not Expressive, Pragmatic, or Expressive Fictions*, "Erkenntnis", n. 9 (2014), pp. 1641-62.
- Schechter, E., *The Unity of Consciousness: Subjects and Objectivity*, "Philosophical Studies", n. 2 (2013), pp. 671-92.
- Schweikard, D., Schmid, H., *Collective Intentionality*, in eds. E. Zalta, U. Nodelman, C. Allen, H. Kim, P. Oppenheimer, *Stanford Encyclopedia of Philosophy*, Stanford, Stanford University, 2021 (Fall edition) <https://plato.stanford.edu/archives/fall2021/entries/collective-intentionality/> (accessed: 5 August 2025).
- Schwengerer, L., *Towards Collective Self-Knowledge*, "Erkenntnis", n. 3 (2022), pp. 1153-73.
- Searle, J., *The Construction of Social Reality*, New York, Penguin, 1995.
- Schutz, A., *Collected Papers I*, Dordrecht, Nijhoff, 1962.
- Szanto, T., *How to Share a Mind*, "Phenomenology and the Cognitive Sciences", n. 1 (2014), pp. 99-120.
- Tomasello, M., *The cultural Origins of Human Cognition*, Cambridge (MA), Harvard University Press, 1999.
- Tomasello, M., *Why We Cooperate*, Cambridge (MA), MIT Press, 2008.
- Tomasello, M., *Origins of Human Communication*, Cambridge (MA), MIT Press, 2009.
- Tomasello, M., *A Natural History of Human Thinking*, Cambridge (MA), Harvard University Press, 2014.
- Tomasello, M., Kruger, A. C., Ratner, H.H., *Cultural Learning*, "Behavioral and Brain Sciences", n. 3 (1993), pp. 495-511.
- Tomasello, M., Racokzy, H., *What Makes Human Cognition Unique? From Individual to Shared to Collective Intentionality*, "Mind and Language", n. 2 (2003), pp. 121-47.
- Tuomela, R., *We-Intentions Revisited*, "Philosophical Studies", n. 3 (2005), pp. 327-69.
- Tuomela, R., *Social Ontology: Collective Intentionality and Group Agents*, Oxford, OUP, 2013.
- Zahavi, D. (2023)., *The Unity and Plurality of Sharing*, "Philosophical Psychology", (forthcoming).