



Habgood-Coote, J. T. (2020). Group Inquiry. *Erkenntnis*.
<https://doi.org/10.1007/s10670-020-00232-5>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1007/s10670-020-00232-5](https://doi.org/10.1007/s10670-020-00232-5)

[Link to publication record on the Bristol Research Portal](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via Springer Nature at <https://link.springer.com/article/10.1007/s10670-020-00232-5>. Please refer to any applicable terms of use of the publisher.

University of Bristol – Bristol Research Portal

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/brp-terms/>



Group Inquiry

Joshua Habgood-Coote¹

Received: 23 August 2019 / Accepted: 10 February 2020
© The Author(s) 2020

Abstract

Group agents can act, and they can have knowledge. How should we understand the species of collective action which aims at knowledge? In this paper, I present an account of group inquiry. This account faces two challenges: to make sense of how large-scale distributed activities might be a kind of group action, and to make sense of the kind of division of labour involved in collective inquiry. In the first part of the paper, I argue that existing accounts of group action face problems dealing with large-scale group actions, and propose a minimal alternative account. In the second part of the paper, I draw on an analogy between inquiry and conversation, arguing that work by Robert Stalnaker and Craige Roberts helps us to think about the division of labour. In the final part of the paper I put the accounts of group action and inquiry together, and consider how to think about group knowledge, deep ignorance, and the different kinds of division of labour.

1 Introduction

Group agents do stuff: corporations avoid taxation, research teams publish their results, and running clubs compete in races. Group agents also know stuff: a corporation might know that its profits are decreasing, the ATLAS collaboration might know that the mass of the Higgs boson is such-and-such, and a running club might know that it stands a good chance of winning the six-stage relays. In this paper I want to connect these two topics, considering the nature of *group inquiry*: collective actions that aim at producing knowledge. Advances in communication as well as increased centralised funding for research has allowed an explosion in large-scale

✉ Joshua Habgood-Coote
josh.habgood-coote@bristol.ac.uk

¹ Department of Philosophy, University of Bristol, Cotham House, Bristol BS6 6JL, UK

collaborative work in science. Some recent notable examples include work in high energy physics,¹ work on Climate change modelling,² the proof of the classification of simple finite groups,³ the mathematical tables project,⁴ the Polymath project,⁵ and a host of citizen science projects.^{6,7}

Group inquiry has a distinctive pattern of epistemic benefits and costs. When people pool their epistemic resources, they can increase their epistemic capacities and broaden their knowledge base, allowing them to take on intellectual tasks would be too time-consuming or complex for one person to take on.⁸ There is also evidence for a range of group-level epistemic effects that boost the reliability of group inquiry (intellectual diversity, transient diversity, and virtuous adversariality).⁹ However, group inquiry also comes with distinctive vices (polarisation, group think, and irreconcilable disagreement). Insofar as a normative epistemology should provide advice to inquirers,¹⁰ social epistemology should be able to offer advice for the design and practice of group inquiries, engaging with and critiquing the epistemic cultures of scientific research groups. My aim is to provide some theoretical groundwork for thinking about group inquiry, as a precursor to this normative inquiry.

There are two central challenges to providing an account of group inquiry:

1. The paradigm for understanding group actions has been small-scale, centralised, and co-operative activities, but group inquiry is in many cases large-scale, distributed, and hierarchical. In order to understand these cases, we need an account of group action that can handle large-scale cases.
2. Group actions typically involve a division of labour. Although we have a fairly good intuitive grip on the division of labour for actions like making a car or

¹ Cetina (1999). A collaboration between the CMS and ATLAS projects with around 5000 authors currently holds the record for the paper with the most authors (Cho 2011; Castelvecchi 2015; Aad et al. (ATLAS collaboration, CMS Collaboration) 2015).

² Edwards (2013), Winsberg (2018).

³ Steingart (2012).

⁴ Historically the majority of the mathematical tables used in navigation, mathematics, and astronomy were the product of collaborative work. See Campbell-Kelly et al. (2003), Grier (1998, 2003, 2005). The Handbook of Mathematical functions—which was initially based on tables produced by a team of 200 human computers working at the Mathematical Tables project has a good claim to be the most widely circulated scientific book in history (Grier 2003: 288).

⁵ Gowers and Nielsen (2009), Martin and Pease (2013).

⁶ Watson and Floridi (2018). Citizen science includes projects to classify galaxies (<https://www.galaxyzoo.org/>), to measure the number and location of birds (<https://www.bto.org/volunteer-surveys/birdatlas>), and to predict protein folding (<https://fold.it/portal/info/about>).

⁷ For further discussion of collaboration in science, see: Bird (2010, 2014), Wray (2002, 2006, 2007, 2014, 2018), Kukla (2012), Winsberg et al. (2014), Huebner et al. (2018), Andersen and Wagenknecht (2013), Wagenknecht (2016), and the papers in Boyer-Kassem et al. 2018.

⁸ Hardwig (1991): 694.

⁹ Hong and Page (2001, 2004), Strevens (2003), Bishop (2005), Zollman (2007, 2010), Mayo-Wilson, Zollman, Dank (2011, 2013), Solomon (2005).

¹⁰ Bishop and Trout (2005).

playing a piece of music, it is unclear how to understand the division of labour involved in epistemic tasks.¹¹

I will address these challenges in turn. First, I will offer an account of group action that is able to deal with the spectrum of cases from small-scale co-operative actions like two people going for a run together, to large-scale distributed actions like thousands of people running an election (Sect. 1). Although minimal, this account will fix the central elements of group action. Secondly, I will present what I call the Stalnaker-Roberts model of conversation, which understands conversation as a kind of collaborative inquiry (Sect. 2). This account gives us a set of fruitful theoretical tools for thinking about inquiry in general. In particular, I will argue that Roberts' idea that we can make progress on a question by splitting it into subquestions gives us a natural model for thinking about the division of epistemic labour. Putting the account of group action together with the Stalnaker-Roberts account of inquiry gives us an account of group inquiry that is answerable to both its epistemic and practical aspects (Sect. 3). I will then draw out some consequences of the account for how we think about the relation between group inquiry and group knowledge (Sect. 4), deep ignorance (Sect. 5) and the different kinds of division of epistemic labour (Sect. 6).¹²

2 Group Action

When we think about group actions, it is natural to start with simple cases. In the literature we find many two-person cases: going for a walk (Gilbert 1990), painting a house (Bratman 1992), carrying a piano up the stairs (Tuomela and Miller 1988), and cooking a hollandaise sauce (Searle 1990: 410–12). Methodologically this is sensible: in theory building it is nearly always a good idea to start with simple cases. (In the next section, I will pursue exactly this strategy). However, solely focusing on these examples runs the risk of skewing our philosophical imagination, leaving us without the tools for thinking about more complex cases.

¹¹ The notion of the division of labour has a double life, occurring both in discussions of co-operation and market-based competition. This double life traces back to Adam Smith's argument that markets can deliver the benefits of planned specialised work (Smith 1776/1982). Many discussions of the division of epistemic labour focuses on benefits of competition (Kitcher 1990; Strevens 2003; Weisberg and Muldoon 2009). Our focus will be on the notion of the division of labour applied to co-operation (Muldoon 2018), but depending on one's view about deal with invisible hand cases (see below), one might think of market mechanisms as a kind of group inquiry.

¹² One might think that group inquiry will just reduce to testimony, removing the need for a separate account of group inquiry. Something like this thought might motivate the centrality of testimony to social epistemology (Goldman 1987, 1999, 2000, 2010; Goldman and Blanchard 2016; Goldberg 2016). While it is true that testimony is an important component of group inquiry, (Rossini and Porter 1979; Wagenknecht and Andersen 2013), there are cases in which a group inquiry can successfully reach a conclusion, without any individual knowing that proposition (Bird 2010: 34; De Ridder 2014; Huebner et al. 2018), which would not be possible if group inquiry were reducible to testimony.

We can think about group actions as lying on a spectrum between *small-scale* and *large-scale* group actions. Small-scale actions involve a small number of people working together on an activity together in a highly co-operative way with a good deal of mutual awareness and common knowledge and without any significant power asymmetries. By contrast, large-scale actions involve a large number of people working together on an activity in a highly distributed way with significant division of practical labour and with the participants knowing comparatively little about what each other are up to (or even who the other participants in the action are). Examples of large-scale group actions include producing pins in a workshop (Smith 1776/1982), writing a newspaper (Babbage 1832: 216, quoted in Shieber 2013: 276), campaigning to abolish slavery (Anderson 2016), and a crew navigating a difficult route into harbour (Hutchins 1995). This distinction is not binary: there are a range of cases with different degrees of each of these characteristics (although cases will tend to have clustered sets of characteristics).

Existing accounts of group action are poorly positioned to make sense of large-scale group actions. Different accounts of group action face slightly different problems.

Searle and Bratman claim that for a group to be intentionally V-ing, its active members must have plural intentions for the group to V (Searle 1990; Bratman 1992). It is easy to imagine participants in a large-scale action lacking suitable commitments to the group's action: a pin-maker might just be indifferent to the group's activity, solely concerned with earning a living, or be actively trying to undermine the group's activity (they might be a spy from another workshop) (Shapiro 2014).¹³

Tuomela and Miller (1988) appeal to a structure of mutual belief between the participants in a group action, requiring that each participant believes that others will do their parts.¹⁴ In large-scale cases, the participants may not know each other, or what the joint project is. In the Manhattan project, a large group of people were involved in enriching Uranium, without even knowing what they were doing, let alone who the other participants were. In this kind of case each participant may perform their part of a joint action while suspending on whether the other participants will do their part.

Gilbert (2009) claims that it is a condition on group action that the members form a so-called 'plural subject', where each member jointly commits to the activity of the group. This model applies to what we might call *democratic groups*, where a group intention must stem from the agreement of all members, but it is difficult to see how to apply it to hierarchical groups where a small set of operative members dictates the collective's plan of action, or groups characterised by a division of labour where a plan splits the task into sub-tasks to be performed by individual agents (Bird 2014).¹⁵

¹³ In a response to Shapiro, Bratman concedes that his framework is not well suited to large-scale cases (Bratman 2014a, b: 333–5).

¹⁴ Even in the individual case, it is questionable whether intending to V requires believing that you will V, or even that you will probably V (see Holton 2008).

¹⁵ Gilbert (2006: C8) defends the possibility of large, disjointed plural subjects, characterised by impersonality, anonymity, and hierarchy. This does not demonstrate that her account can deal with large-scale

The distinction between small-scale and large-scale actions also applies to collective epistemic actions. We can think of *conversation* as our prototypical small-scale group epistemic action. Conversation typically involves just a few participants, a body of shared information, co-operation on a strategy of inquiry, very little division of labour, and a flat power structure. By contrast, the examples of large-scale scientific collaboration we started off with involve more developed power structures, teams being split up into sub-teams within a hierarchical structure, limited shared information, and significant power asymmetries. If we work with existing accounts of group action, then we will not be able to give a general account of group inquiry.¹⁶

I propose that we strip back our understanding of group action to a minimal account that focuses on the central properties of group action, allowing that these properties can be realised in different ways. There is a great deal of disagreement about the analysis of intentional action, but it is common ground in this debate that intentional action involves an agent acting in way that is explained by her intentions.¹⁷ This gives us the following minimal account:

MINIMAL-INDIVIDUAL: An agent *S* is intentionally *V-ing* iff i) *S* is *V-ing* ii) *S* intends to *V*, and iii) condition i) obtains because of ii).

I want to leave things at this fairly abstract level, making no commitments about the nature of action, intention, or the explanatory relation involved.¹⁸ The point is that if there is genuine collective action, then these functional properties will be realised by groups. This suggests the following minimal account on group intentional action:

MINIMAL-GROUP: An group *G* is intentionally *V-ing* iff i) members of *G* are either *V-ing* or performing sub-tasks of *V*, ii) *G* intends to *V*, and iii) condition i) obtains because of ii).

This account claims that a group action is constituted by a bunch of individual actions which are explained (in some suitable sense) by a collective plan or intention.

Condition i) takes the individual action condition in MINIMAL-INDIVIDUAL and replaces it with a requirement that members of the group be performing suitable

Footnote 15 (continued)

group actions. It might well be possible for large groups to be characterised by certain ‘thin’ joint commitments (such as the commitment to obey a set of social norms), but this does not establish that *all* large-scale groups have the rich commitments required by joint action on the plural subject account.

¹⁶ For an example of this problem, see Tossut (2014). Tossut uses a Bratman-inspired account of scientific collaboration, meaning that her account claims that collective knowledge requires agreement and shared epistemic goals amongst the members of the group (Tossut 2014: 361).

¹⁷ The action-intention connection (iii) is often associated with what Bratman calls the simple view, which claims that intentionally *V-ing* requires an intention to *V* (Bratman 1987: C8). I mean the connection to be compatible with intentionally *V-ing* being explained by the intention to do something else.

¹⁸ A sense of the options. Actions might be: events, states, or processes. Intentions might be a species of belief, a species of knowledge, or a *sui generis* mental state. The explanatory connection might be causal, non-causal, teleological, or a kind of reasons explanation. The direction of explanation between intentional action and these conditions might go either way, either offering a reductive account which explains intentional action in terms of intention, action, and action-explanation, or an action-first account which explains intention, action, and action-explanation in terms of intentional action (Levy 2013).

actions: either the action being performed by the group, or suitable sub-activities.¹⁹ A team of removers might move a piano from A to Z either by all moving the piano from A to Z together, or by one agent moving the piano from A to B, the next moving it from B to C and so on. This expresses the idea that a group can only act by means of its members acting, meaning that there is no group-level basic action (Searle 1990: 410; Lackey 2014).

Condition ii) requires that the actions of individuals must be accompanied by a collective intention. Exactly what intentions are at the individual case is contested. Nonetheless, the literature on individual intentions suggests a broad functional characterisation of group intentions. A group intention is a state that is: (i) the characteristic output of collective practical deliberation, (ii) that will (typically) dispose a group to either V, or to take plausible means to V-ing, (iii) that requires the group to form appropriate means-ends intentions, and not to form incompatible plans' (Bratman 1987), and (iv) that requires that the group has knowledge of how to carry out its plans (Habgood-Coote 2017).²⁰

Condition iii) connects individual action with group intention, expressing the idea there needs to be a collective plan that explains the pattern of individual-level actions (see Tuomela and Miller 1988: 369–70; Searle 1990: 402–3). Exactly how to understand the explanatory relation between action and intention is a difficult question but for our purposes we can work with an intuitive understanding of it.

This account is deliberately non-committal about the nature of action, intentions, and action explanation (see footnote 18). This means that it leaves open a number of difficult questions, including how to think about uncooperative participants, and whether invisible hand cases involve collective action. I suggest that the answers to these questions will depend on general issues about the nature of intentional action.

Depending on how one thinks about the explanatory relation between group intention and individual action, MINIMAL-GROUP is compatible with individuals being part of a group action without intending for the collective's action to be successful, or even whilst intending that the collective fail. A worker in a factory making cars may perform her tasks because she has been told to, although she is indifferent to whether the factory makes cars. A participant in a citizen science project may categorise galaxies because she likes playing free computer games, although she couldn't care less about the project's scientific aims. A saboteur in the factory might be performing the normal duties of a factory-worker whilst intending that no cars are made. A participant in a citizen science project may miscategorise galaxies to hamper the project. In each of these examples, there is a case to be made that the individual actions can be explained by the group's intention in the sense of explanation relevant to intentional action. In the factory worker case, the manager's

¹⁹ We will need to understand 'sub-task' in a capacious sense to allow that groups that groups with incorrect means-ends beliefs are still engaged in the relevant activities.

²⁰ The idea that group intention is the output of group level deliberations means that we need to eschew a pure judgement aggregation approach to group intention, which would entail that a group's intentions at a time supervenes on its (operative) members' intentions at that time. In addition to looking to the relation between group and individual mentality, we need to pay attention to the way groups' mental lives develop over time.

commands embody the plans of the group, and the worker does follow those commands, meaning that the worker's actions are in a sense explained by the group's intention. If we want to explain why the participant in the citizen science project is classifying galaxies, we would need to appeal to the fact that the online game was disseminated by a scientific project with certain epistemic aims. In the saboteur cases, we could explain the sabotaging actions by citing the group's intention (although we would also need to cite the saboteurs' disruptive intentions). In each of these cases there is also a case to be made that individuals' intentions are *not* explained by the group's intention in the way distinctive of intentional action. Perhaps these explanations are merely causal, and should be distinguished from action-explanations. If we take this line then the indifferent saboteurs are not *part of* the collective agent, and are more like tools for the collective action.

MINIMAL-GROUP is also neutral on whether 'invisible hand' cases in which a group of agents pursuing a variety of aims brings about some unintended consequence count as cases of group action. If we allow that a group intention can supervene in a complex way on a body of individual intentions with different contents, and that the kind of explanation involved in intentional action is causal, then invisible hand-type groups might be engaged in group actions. By contrast, if we think that group intentions either require individual intentions with the same content, a group-level process of deliberation, or that action explanation is distinctive, then this kind of group will turn out *not* to be engaged in group action (Searle 1990). This issue is important in the context of thinking about group inquiry, since if the invisible hand creates group agents, prediction markets, scientific disciplines, and democratic societies will turn out to be engaged in group inquiry.²¹

The aim of MINIMAL-GROUP is to provide an account of group action that is sufficiently general to cover the full range of small-scale and large-scale cases. In small-scale actions, the group's plan may be determined by the participants agreeing on a plan, which they all commit to acting out, leading them to perform suitable sub-activities because of their collective plan. By contrast, in a large-scale action, the plan may be determined unilaterally by the leaders of the group or by some decision-making procedure, and the participants in the group may perform their tasks because their managers have told them to. In both cases, we find a pattern of individual actions which are animated and explained by a group-level plan, but the way in which the collective intention is realised differs dramatically. In small-scale actions, the participants will take the cognitive load associated with an intention onto their own shoulders, with each participant engaging with the deliberative process, thinking about whether the plan is possible, about how it might be performed, and trying to resolve any tensions between the various plans the group might have, as well as actually carrying out the relevant action. In this kind of case, we will find a fairly close alignment between the group's intention, and the participants' intentions. (We may find that in cases toward the small-scale end the conditions proposed by Searle, Bratman, Tuomela and Miller, or Gilbert are sufficient for group intention).

²¹ On 'invisible hand' explanations in science, see Hull (1997), Kitcher (1993), Wray (2000), Strevens (2003), Zollman (2018).

By contrast, in large-scale cases, the bundle of functional properties distinctive of intention will be distributed across the various members of the group. There may be one team that decides what the group will do, another that checks whether the plan is compatible with the group's beliefs, and further groups that determine how the plan will be carried out, resolve any tensions between the group's plans, and actually carry out the group's plan. This kind of division of labour across the functional properties distinctive of intention allows for considerable divergences between the group's intentions and the intentions of participants within the action.

3 Conversation

Let's start our investigation of group inquiry with the simplest joint epistemic action: a conversation between two agents. An important thread of work in pragmatics, tracing back to Robert Stalnaker and Craige Roberts, uses the idea that conversation is a kind of collaborative inquiry to explain a range of linguistic phenomena. The goal of this section is to unpack the core of this framework, and use it to understand the division of epistemic labour involved in group inquiry.²² The Stalnaker-Roberts model involves various idealisations, and we should not expect every feature of this model to be reflected in conversations or bigger group inquiries.²³

The core of the Stalnaker-Roberts model is a Lewisian picture of conversation as a game. This game involves three kinds of moves: *assertion*, *asking*, and *directing*. These moves are associated with different grammatical moods, different kinds of content,²⁴ and different effects on the conversational scoreboard (Lewis 1979).

Start with *assertion*. Following Stalnaker (1999) we assume that the participants in a conversation start with a set of propositions which they take for granted for the purposes of that conversation. The acceptance of these propositions is taken to be a matter of common belief, and they are represented on the conversational scoreboard by a set of propositions which make up the *common ground*.²⁵ The intersection of the common ground propositions is the *context set*. In a possible worlds framework where we think of a proposition as a set of possible worlds, we can think about each proposition in the common ground as ruling out the worlds in which the proposition is false, meaning that the context set is the set of worlds which the participants

²² For presentations of this framework, see Stalnaker (1999 especially chapter 4), Stalnaker (2014), Roberts (2012, 2018).

²³ This is a general problem for philosophy of language, see Cappelen and Dever (2019), Beaver and Stanley (2019).

²⁴ Exactly how to think about the content of declaratives, interrogatives, and imperatives is a vexed question. I'll use a possible worlds framework and associated interrogatives and imperatives with non-propositional contents. Other approaches are available.

²⁵ Gilbert and Priest argue that the combination of acceptance and common belief associated with the propositional attitude of presupposing is close to Gilbert's notion of joint commitment, suggesting that we might model the common ground as a set of collective beliefs (Gilbert 1989: 294–8; Gilbert and Priest 2013).

consider to be possible for the purposes of conversation.²⁶ The basic way to manipulate the common ground is by asserting: uttering a sentence in the declarative mood expressing a proposition, thereby proposing to add that proposition to the common ground. If I utter the sentence ‘Tahlia came to the party’, I express the proposition *Tahlia came to the party*, and propose to add this proposition to the common ground. If this assertion is unchallenged, we add the proposition to the common ground and shrink the context set, ruling out the worlds in which Tahlia didn’t come to the party.

Next, *asking*. Conversations typically have topics which constrain what we can say. If we are talking about where to go to dinner, I’ll get annoyed if you start offering hot takes on Lithuanian politics. We can represent the topic of a conversation by adding *questions under discussion* to the scoreboard (Roberts 2012), which represent the questions that the participants in a conversation are committed to answering. As I will use the term, *questions* are entities on the same level as propositions that are expressed by *interrogative sentences* and are associated with the speech act of *asking*.²⁷ On a possible worlds approach, we can think of a question as a *set of sets* of possible worlds corresponding to the possible answers to that question. Following Roberts, I will treat this answer set as a set of *exhaustive answer propositions*, meaning that these sets of worlds are mutually incompatible, and that a question is a *partition* over a portion of logical space.²⁸

Asking is a proposal to add a question to the set of questions under discussion. If the proposal is accepted, two things happen: we add a new partition to the scoreboard, which divides up the context set giving us a new range of alternatives to distinguish between, and we commit to trying to answer that question. If I utter the sentence ‘who came to the party?’ I thereby propose to add *who came to the party?* to the questions under discussion, meaning that we add the following partition to set of worlds in the context set (Fig. 1).

²⁶ We might worry that the notion of the common ground introduces a level of co-ordination which is incompatible with highly distributed inquiries. We return to this issue in Sect. 3.

²⁷ I will put interrogatives in quotes, and italicise questions. The sentence ‘who came to the party?’ expresses the question *who came to the party?*

²⁸ Here is a recipe for generating Roberts-style partitions (Roberts 2012: 6:9–6:12). Start with an interrogative phrase: ‘who came to the party?’. Split it into two parts: the *wh*-word—‘who’—which we treat as a variable, and the question abstract—‘came to the party’—which expresses a property. Take the salient domain, including the restriction triggered by the question word, and generate all of the possible complete assignments of the property to the objects in the domain. If our domain is just Ann and Bernard, then the complete assignments of the property *came to the party* will be:

- (i) Ann and Bernard came
- (ii) Ann came, and Bernard didn’t come
- (iii) Ann didn’t come, and Bernard came
- (iv) Ann and Bernard didn’t come

Each of these assignments is a complete answer to the question. In this case the interrogative carries the presupposition that there in fact was a party, so the partition only divides up the portion of logical space where there was a party.

Furthermore, we commit to trying to resolve the question by ruling out all of the false answers.²⁹

Finally, consider directive moves.³⁰ Directive moves are associated with sentences in the imperatival mood, which we can think of as expressing an action directed toward the addressee (Portner 2007, 2018).³¹ If I utter the sentence ‘stand up!’ addressing Trey, then I express the action *standing up*, directed towards Trey. Directive acts are proposals for the addressee to change their intentions, meaning that we can represent their conversational effects by adding a *To-Do list* of explicit intentions to the scoreboard (Portner 2007). Since accepting a question into the set of questions under discussion involves a commitment to resolve the question, asking a question involves adding both a question to the question-stack, and the commitment to resolve it to the To-Do list. The combination of intentions to answer questions on the To-Do list will constitute a group’s plan of inquiry, which will represent both the questions they are investigating, and how they plan to resolve them.

Our discussion gives us a threefold categorisation of conversational moves, with associated moods, kinds of content, and changes to the conversational scoreboard:

Move	Mood	Content	Effect on scoreboard
Asserting	Declarative	Proposition	Add proposition into the common ground
Asking	Interrogative	Question	Add question to the <i>questions under discussion</i> , and add answering the question to the <i>to-do list</i>
Directing	Imperative	Property	Add task corresponding to property to the <i>to-do list</i>

The conversational scoreboard corresponds to the conversational dispositions that the speakers take one another to have.³² For a proposition to be in the common ground, the participants need to both be taking that proposition for granted, and believe that they are so doing so. For a question to be included in the questions under discussion, participants need to commit to resolving that question and believe that they are doing so. If the participants in a conversation fail to co-ordinate their attitudes they will end up in a *defective context* (Stalnaker 1999: 85). Representing defective contexts and their resolution is a substantial issue in pragmatics, and

²⁹ What about practical questions? Three options: (i) Treat the cells in a practical question as propositions about the way the world could be, and think of answering a practical question as deciding which proposition to make true. (ii) Treat the cells as propositions about which actions an agent all things considered ought to do, meaning that answering involves forming a belief about one ought to do. (iii) Extend the semantics for interrogatives, allowing questions which are partitions of possible actions and for questions to be answered by forming intentions (See Roberts 2009).

³⁰ I am construing ‘directive’ broadly, to include commands, advice, permissions, suggestions, and warnings.


³¹ This is a simplification of Portner’s view, which involves multiple To-Do lists indexed to different modal flavours (i.e. commands, advice, warnings). For Roberts’ development of Portner’s view, see (Roberts 2018: 338–329).

³² Translated to Gilbert’s idiom we might think that the conversational scoreboard consists of the propositions, inquisitive and non-inquisitive goals which the participants in a conversation are jointly committed to for the purposes of a conversation (see footnote 25).

Question: *Who came to the party?*

Domain: Ann, Bernard.

Key: A: Ann came; a: Ann did not come; B: Bernard came; b: Bernard did not come.



Options: $\{\{AB\}; \{Ab\}; \{aB\}; \{ab\}\}$

Fig. 1 Question

a complete account of group inquiry should also be able to represent inquiries with defective epistemic backgrounds.³³

With this picture of the basic moves in a conversation in place, we can tell a story about how conversational moves contribute to answering the question under discussion. There are four basic ways participants can make progress on a question: (i) by asserting a complete answer, (ii) by asserting a partial answer, (iii) by asserting a casual answer, and (iv) by asking *other* relevant questions. In the remainder of this section, we will run through these ways of making progress in inquiry, introducing relevant concepts along the way.

As we are thinking of things, a question is a set of mutually exclusive alternatives, meaning that we can think about making progress on a question in terms of ruling out possible answers. A *complete answer* rules out all but one of the possible answers, and a *partial answer* rules out at least one of the possible answers (Roberts 2012). Returning to our example above: the proposition *A and B came* is a complete answer to the question *who came to the party?* because it rules out all but one cell in the partition, and the proposition *A came* is a partial answer because it rules out at least one answer to the question. We can define partial and complete answers as follows:

Complete Answer A proposition p is a **complete** answer to a question Q iff p rules out all but one cell in the partition associated with Q .

Partial Answer A proposition p is a **partial** answer to a question Q iff p rules out at least one of the cells in the partition associated with Q .

³³ Explaining defective group inquiries is a question for another day, but see Sect. 5 for a way one might represent disagreement about the question under investigation.

Partial answers can be combined to yield a complete answer to a question. If we put together the partial answers *A came* and *B came*, we rule out all of the cells except *A and B came*, giving us a complete answer to the question of who came.

One can make progress on a question by asserting a proposition that entails a answer to the question when supplemented by propositions in the common ground (Roberts 2012: 6:12). If the common ground includes *A came iff there was beer*, asserting *there was beer at the party* entails a partial answer: *A came*. One can also make progress by asserting a proposition which would entail a partial answer when put together with accessible information.³⁴ Even if *A came iff there was beer* isn't yet in the common ground, asserting *there was beer at the party* might still be helpful because the information about A's alcoholic proclivities might be accessible. This gives us weaker notions of complete and partial answerhood:

Casual Complete Answer A proposition p is a **casual** complete answer to a question Q iff there is some accessible proposition q and $(p \ \& \ q)$ rules out all but one cell in the partition associated with Q .

Casual Partial Answer A proposition p is a **casual** partial answer to a question Q iff there is some accessible proposition q , and $(p \ \& \ q)$ rules out at least one of the alternatives in the partition associated with Q .

We can also make progress on a question by asking more questions. In many cases a question under investigation will be extremely complex and difficult for the participants to deal with. One way that we can deal with this is by overlaying a complex partition corresponding to a question with simpler partitions that divide the same area of logical space into fewer cells. We might overlay the question *who came to the party?* with the simpler questions *whether A came?* and *whether B came?* giving us two simple yes–no questions to consider (Fig. 2).

Let's call the relation between complex and simpler questions the *question-sub-question* relation:

Question-Subquestion Relation: A question $Q1$ is a sub-question of $Q2$ iff every complete answer to $Q2$ entails a partial answer to $Q1$. (see Roberts 2012: 6:6–6:7).

Answering a sub-question of a question will be helpful because the propositions which completely answer the sub-question will be partial answers to the initial question, meaning that completely answering a subquestion will partially answer the original question.

There will also be a casual version of the question-sub-question relation indexed to accessible information. The question *was there beer at the party?* is a *casual* sub-question of the question *who came to the party?* because complete answers to the first question together with accessible propositions about the presence of beer entail partial answers to the second question. We can define this relation as follows:

³⁴ Information is accessible to an agent if she has the capacity to get hold of that information (say by perception, inference, or testimony). That a piece of information is accessible does not entail that it is easy to get hold of.

Casual Question-Subquestion Relation: A question $Q1$ is a *casual* sub-question of $Q2$ iff for every proposition p which is a complete answer to $Q2$ there is some accessible proposition q , such that $(p \& q)$ entails a partial answer to $Q1$.

The question-subquestion relation is important because it gives us a useful tool for thinking about the mereology of questions. The sub-questions of a question are *part* of that question, answering a sub-question gives *part of the answer* to the initial question, and knowing the answer to a subquestion entails knowing *in part* what the answer to the initial question is.³⁵ This part-whole structure is crucial to inquiry. When we face a particularly complex question, we may need to form an plan of inquiry that splits the initial question up into various subquestions (its parts) which we can address in turn. These subquestions may themselves have enough structure to allow further subquestions, leaving us with a hierarchically organised stack of questions to be resolved.

Forming plans of inquiry involving subquestions is an important tool for cognitively limited inquirers. Splitting a question into sub-questions both helps us see how information bears on a complex question, and allows us to divide up our intellectual labour both across time and between people. In general, we can think of the division of labour as the process whereby a complex task is split up into simpler subtasks, which can be performed separately either at different times or by different people. We now have a way to think about the tasks and subtasks involved in inquiry: the subtasks of an inquiry into some question Q will be inquiries into *sub-questions* of Q . These subinquiries may be divided up across time, in a sequential process of resolving subquestions. Or they may be divided up interpersonally, with different members of a group inquiring into different subquestions at the same time. Consider the way in which a team of detectives might organise to investigate a murder. Rather than all inquiring into *who committed the murder?* they might decide to split up, with different detectives considering questions like *who had a motive?*, *what does the forensic evidence show?*, *did A murder the victim?*, and so on. To put the point in our idiom: the detectives collectively inquire into the big question *who committed the murder?* by individually investigating its subquestions because of a collective inquisitive plan.


To summarise: the Stalnaker-Roberts account provides us with a number of ideas that are essential to understanding group inquiry. The central idea is that inquiry involves a process of ruling out potential answers to a question. Around this, they build an account of the rational structure of inquiry. The question-subquestion relation helps us to think about the kind of planning involved in inquiry, and helps us to understand the division of epistemic labour involved in group inquiry. The idea that inquiry aims at resolving questions also gives us a picture of the success conditions of inquiry, telling us that an inquiry is successful when it reaches a complete answer to a question.³⁶ In the next section, we will put these ideas to work to develop an account of group inquiry.

³⁵ For a discussion of the relation between knowledge and the parts of questions see Pavese (2017).

³⁶ Often we will be satisfied with partial answers. If I ask how to skin a cat, I will be happy with just one way. To make sense of this fact, we might introduce the notion of a *resolving answer*: an answer that gives sufficient information to satisfy the practical and epistemic interests of the speaker(s).

Subquestion 1: *Whether Ann came to the party?*

Options: $\{\{AB; Ab\};\{aB; ab\}\}$



Subquestion 2: *Whether Bernard came to the party?*

Options: $\{\{AB; aB\};\{Ab; ab\}\}$




Fig. 2 Subquestions

4 Group Inquiry

Applying the minimal account of group action from Sect. 2 to group inquiry gives us this schematic account:

MINIMAL-GROUP INQUIRY: A group G is inquiring into question Q iff i) the members of G are either inquiring into Q , or performing suitable subtasks, ii) G has an intention to inquire into Q , iii) condition i) obtains because of the group’s intention.

We are now in a position to fill in this account, using concepts from the Stalnaker-Roberts account of conversation.

We can use our account of the division of epistemic labour to understand the subtasks involved inquiry, meaning that the subtasks of inquiring into some question Q will be inquiries into subquestions of Q .

Roberts’ picture of plans of inquiry also helps us to see what kinds of intentions will be involved in inquiry. A collective intention to answer a question will involve both an aim—answering a particular question—and will create pressure to form a plan for achieving that aim, in the form of an ordered set of subquestions of that question, and an allocation of subquestions to members of the group.

To help us think about the distinction between the aims and the means-ends structure of inquisitive intentions, let's introduce a distinction between *object questions* and *methodological questions*. An object question is the question that the group aims to resolve, and the methodological question is the practical question of how they will resolve that question. The object question is a set of propositions about the subject-matter under investigation (only one of which will be true), and the methodological question is a set of alternative strategies of inquiry by which the group might investigate the object question, (only some of which will be potentially successful).³⁷ If the object question is *whether p?*, the object question will be *how shall we answer whether p?*. These two questions are intimately related, although neither is a sub-question of the other. Forming an intention to investigate some question immediately *raises* the question of how to resolve that question; just as forming an intention to do something immediately raises the question of how to do that thing. Answering a question will often involve an intertwined process of answering both object and methodological questions, just as action often involves an intertwined process of doing something and working out how to do it see (Habgood-Coote 2019). And, we might think that knowing an answer to the object question may require agreement on a possible answer to the methodological question.³⁸

With this framework in place, we can sketch the structure of group inquiry.

We start with a group of agents who are investigating a question: say *whether p?*. With simple questions, a group can investigate by forming a collective intention to resolve the question by means of one person finding out the answer. However, with more complex questions it might not be possible for an individual to resolve the question. Some questions are *laborious*, requiring a huge amount of epistemic labour (think of complicated but mundane calculations). Other questions are *specialised*, requiring the combination of multiple bodies of expert knowledge or skills which are not possessed by any individual (think of a question in sociology that requires ethnographic skills to collect the data and statistical methods to analyse it).³⁹ When a group faces a question which is laborious or specialised, they will need to switch to the methodological question, and agree on a plan of inquiry that addresses the complexity of the object question. A successful strategy of inquiry will involve splitting the object question up into relevant subquestions, assigning these subquestions

³⁷ In a partition-based framework *How to answer whether p?* gives a set of cells which are complete applications of the predicate *is a way to answer whether p?* to *all* of the possible methods in the domain. This means that a complete answer will be a proposition that tells us which exactly which ways are ways to answer the question *whether p?* and which are not. As with many infinitival questions, in this case we favour the mention-some reading, allowing a partial answer to count as a resolving answer. We only require a partial answer which gives a positive application of the predicate *is a way to answer the question whether p?* to one method in the domain.

³⁸ What happens when some members of a group think that a method is a way to answer the object question, and others disagree? Some options: Lackey's account of group justification claims if adding up the bases of justified beliefs doesn't yield a coherent belief set, a group belief can't be justified, making methodological disagreement something like a higher-order defeater for collective knowledge (Lackey 2016). Dang argues that methodological disagreement can *enhance* group justification (Dang 2019).

³⁹ On this distinction, see Hardwig (1991): 695, De Ridder (2014): 46, Wagenknecht (2014).

to members of the group in line with their epistemic capacities, and determining an order to resolve them. With a suitable division of epistemic labour in place, the members of the group can get on with the subinquiries, by finding out new information or by making inferences from their standing knowledge. In many cases the inquiries into the subquestions of the object question will themselves be undertaken by groups of agents, giving rise to further levels of the division of labour.

When the subinquiries of some question are complete, there will be a further stage in which the answers to the subquestions are collated to give a complete answer to the object question. This final stage may be carried out by an individual who understands the partial answers, in which case the inquiry will wind up with the collator knowing the complete answer to the object question. However, in other cases the collation may be done automatically (as in Tollefsen's UN population study case (Tollefsen 2007: 31; Lackey 2014: 289)), may be done by someone who does not understand the partial answers (as in Bird's interdisciplinary modus ponens case (Bird 2010: 34)), or may merely involve papers expressing the relevant propositions being published in suitable journals or repositories (Bird 2010: 35–36). In these cases the collation of answers may be successful without any individual coming to know the complete answer to the object question, leading to group knowledge of a proposition which no individual knows.

Putting together our understanding of the task-subtask structure of inquiry, our picture of inquisitive intentions, and the idea that group action must involve group members' actions being explained by a collective intention, we get the following account of group inquiry:

Group Inquiry

A group G is inquiring into a question Q iff

- i) Some of the members of G are either inquiring into Q , or inquiring into subquestions $Q_1, Q_2 \dots Q_n$ of Q ;
- ii) G has an intention to inquire into Q in accordance with a strategy of inquiry S ;
- iii) The active members of G are engaged in their sub-inquiries *because of* the group's inquisitive intention.

Since it is built on MINIMAL-GROUP, this account inherits its central features, and leaves open the same choice points.

GROUP INQUIRY is neutral on whether a participant in a group action needs to intend the success of the group's activity. A scientist in a very hierarchical laboratory might well be part of a project to answer some question, but be indifferent to finding out the answer, or even intend to sabotage the project. Whether they are part of the inquiry is a difficult question which turns on how we understand action-explanation. This account also leaves open the possibility of 'invisible hand' cases

counting as group inquiries, with the possibility that scientific disciplines (Gilbert 2000; Weatherall and Gilbert 2016), or even science as a whole (Bird 2010) might count as one giant collective inquiry (Wray 2007). In both cases, these questions turn on larger issues about collective action.

This account covers the full range of small and large-scale group inquiries. The basic structure of the account is built from the Stalnaker-Roberts account of conversation, meaning that the account predicts that any conversation counts as a group inquiry. In such cases, a group's inquisitive intention will be realised by closely related individual intentions. The account also has the resources to explain how many agents working together in a complex institution might be part of one inquiry. We can say that the ATLAS and CMS projects at CERN are group inquiries, because all the participants in these projects are pursuing their inquiries *because of* a collective plan to answer questions about fundamental physics.

One crucial difference between small and large-scale inquiries concerns the role of the conversational scoreboard. When just a few agents are investigating a relatively simple question, it is feasible for all of the participants to keep track of contributions to the inquiry, meaning that what is on the scoreboard can be a matter of common knowledge. However, as questions increase in complexity, just keeping track of contributions becomes a complex task. It would be simply inefficient for a massive collaborative project to maintain a co-ordinated common ground and stack of questions under discussion. Instead, I suspect that we will find that the work keeping the score of an investigation will be distributed across different participants, much like the functional properties of collective intentions. A central planning committee might keep track of the answer to the methodological question, with smaller groups keeping track of stacks of subquestions, and progress on them, and collators pulling together the partial answers that issue from the subinquiries to resolve the object question.

One might worry that the lack of a co-ordinated scoreboard undermines the analogy between conversation and group inquiry. I don't think we should take this worry too seriously. Recall Lewis' idea that the conversational scoreboard is analogous to the score in a baseball game (Lewis 1979). Some games involve all of the players keeping the score together. All of the players in a game of pick-up basketball might take responsibility for keeping the score. However, in professional games keeping the score is a distinct task, which may itself be subdivided into roles like referee and linesperson. Similarly, in a group inquiry keeping track of the score will be a specialised task. This task might be performed by a specialised group or it can be split into subtasks and distributed between different participants.⁴⁰

5 Group Inquiry and Group Knowledge

There is a close connection between inquiry and knowledge. It is plausible that ignorance is the norm of inquiry (Whitcomb 2017), and that knowledge is both the central aim and the success condition of inquiry (Williamson 2000; Kelp

⁴⁰ An example: In Polymath projects one of the roles of the moderator is to periodically write up progress reports, and each project is associated with an open wiki page to keep track of the project. It is natural to think of these reports and the wiki as the scoreboard for the next stage of the project.

2014). When a group agent inquires into some question by employing a division of labour into suitable subquestions which are then collated to resolve the question, the product of this process will be that the group knows the answer to the object question. This connection raises the issue of what consequences GROUP INQUIRY has for our understanding of group knowledge. Following Fagan (2011, 2012) we might distinguish two ways in which knowledge might be social: being the product of a social process, or being the property of a group of agents. GROUP INQUIRY opens the door to both kinds of social knowledge.

When knowledge is the upshot of a group inquiry, it will be social in the process sense. When a group undertakes a group inquiry, its knowledge will be the outcome of a social process. Here the contrast is with a group coming to know something from an individual, as when a group learns something from testimony from an individual who found it out by herself. Individual knowledge can also be social in this sense. If some individual were to come to know a group's results by reading their published findings, this knowledge would still be social in the sense that it originates in a social process (Shieber 2013; Miller 2015). Given the ubiquity of large-scale scientific collaborations, a considerable amount of our everyday knowledge is social in this sense.

Group inquiry can also yield knowledge which is social in the sense that it is the property of a group. Successful inquiry changes the mental state of the group, and group inquiry will (if successful) yield collective knowledge of the answer to the object question. In many cases successful group inquiry will involve at least one member of the group coming to know the complete answer to the object question. If a group collates the results of its subinquiries by having one member collect and put together the partial answers to reach the complete answer, the upshot will be one member coming to know the complete answer to the object question.

There may also be cases where group inquiry succeeds without any individual coming to know the complete answer to the question. If the process of collating the results of the subinquiries is automated, is carried out by multiple compilers, or academic publication suffices for compiling, inquiry can be successful without any member of the group coming to know the complete answer to the question. The UN Population study case provides a nice example: we are to imagine each of the members of the commission inquiring into subquestions of the object question *what are the trends in world population?* with the results of these subinquiries then being collated either by a team of collators (Tollefsen 2007: 301), or by an automated process (Lackey 2014: 289). In these cases, the group knowledge which outputs from group inquiry has a distributed or fragmented character, with the group knowing the answer to a question in virtue of members of the group knowing the answers to a suitable set of subquestions (see Habgood-Coote forthcoming).

6 Inquiry and Deep Ignorance

One worry is that this account assumes that inquirers have the resources to represent the possible answers to a question. It is clearly possible to investigate a question without knowing what its *possible* answers are (Friedman 2013). The physics

community might investigate the question *what is the correct fundamental theory of physics?* before they know what kind of thing a fundamental theory of physics is.⁴¹ An adequate account of inquiry ought to have the resources to represent not just ignorance of which answer is correct, but also the kind of deep ignorance that occurs when we don't know what the answers to a question are (Bromberger 1992; Willholt 2020), and the distinctive intellectual achievement involved in understanding a question.⁴²

We have already distinguished object questions concerning a worldly subject-matter from methodological questions concerning how to resolve an object question. I propose that we add a third kind of question: *option questions*. An option question concerns the alternatives associated with a particular question, and will be of the form *what are the possible answers to $Qn?$* , (where $Qn?$ might be either an object or methodological question). An option question is a question about a question, meaning that each possible answer of an option question will be a set of possible answers to another question. If we think of an answer as a set of possible worlds, and of a question as a set of sets of possible worlds or a partition, an option question will be a set of sets of sets of worlds, or a set of partitions (Fig. 3).

In a condition of deep ignorance, there may be very many ways to chunk up logical space to consider, and resolving an option question will in many cases be no mean feat. Option questions raise a host of difficult questions: What are the reasons for deep ignorance? How can we find out what the possible answers to a question might be? Can we resolve a question, despite not knowing what all of its possible answers are?

7 The Division of Epistemic Labour (Again)

For the most part, we have focused on the division of epistemic labour regarding the object question. However, our discussion of distributed scoreboards and deep ignorance helps us to see that group inquiry can involve different kinds of division of labour.

⁴¹ Here I will focus on complete deep ignorance. There are a number of more complex cases: ignorance of some of the options, ignorance about where the division between options fall, and ignorance about the meaning of the terms used to express the question. I will also set to one side the question of how to think about the relation between an inquirer and the question under investigation in cases of deep ignorance.

⁴² A related problem with the framework arises in the case of mathematical inquiry. Once we start thinking about mathematical questions, the possible worlds framework for thinking about questions gets into trouble, because mathematical propositions are necessarily true or false, meaning that we cannot represent mathematical questions as partitions over logical space. There are a couple of ways to get around this: i) introduce partitions over epistemic possibility space., ii) allow impossible worlds into our metaphysical possibility space, iii) identify mathematical propositions with partitions over possible worlds (Pérez Carballo 2016), and then identify mathematical questions with sets of partitions (sets of sets of worlds).

The most important kinds of division are:

1. *Object question division*: whereby the object question is split into smaller sub-questions that are easier to answer;
2. *Methodological question division*: whereby the methodological question is split up into simpler methodological questions, which might involve either the splitting of one method into small sub-tasks, or different groups pursuing complementary routes to the same answer hoping to triangulate on one answer;
3. *Scoreboard division*: whereby the members of the group distribute the task of representing the scoreboard;
4. *Evidence division*: whereby different members of the group pool evidence that supports an answer without clinching the matter.
5. *Option question division*: whereby different members of a group investigate different divisions of logical space, in the interests of finding out what the right question is.

These kinds of division may occur in isolation, but my suspicion is that most real-world cases will involve a mix of different kinds of division of labour, and that much of the difficulty involved in interdisciplinary research involves managing these different kinds of distribution of labour, as well as negotiating the scoreboard of inquiry to address respects in which it is defective.⁴³

8 Conclusion


An adequate account of group inquiry faces two challenges: the lack of a framework for thinking about large-scale group actions, and a lack of clarity about how to think about the division of epistemic labour. In the first part of the paper, I proposed a minimal account of group action which was sufficiently general to cover both large-scale and small-scale group actions. In the second part of the paper, I presented the Stalnaker-Roberts model of conversation before generalising to give an account of the epistemic dimension of group inquiry. Putting the two parts of the story together gives us an account of group inquiry that is answerable to both its practical and epistemic dimensions, whilst being able to account for large-scale group inquiries, and the division of epistemic labour which they involve.

In closing, I want to consider some directions for future work:


For the most part, we have focused on examples of large-scale group inquiry from science, neglecting some prominent examples of democratic institutions that aim to harness collective intelligence: elections, citizens juries, public inquiries and social media (Anderson 2006; Landemore 2012). How might we think of these groups as group inquiries, and how might this model be used to ameliorate democratic institutions?

⁴³ See Wagenknecht's (2014, 2016) ethnographically-informed discussion of *opaque* and *translucent* epistemic dependence. Wagenknecht's distinction is related to the distinction between methodological and scoreboard division, and (mere) object question division.

p : the set of worlds in which p is true (i.e. {PQR; PQR; Pqr; PqR}).



Whether p ?: the set of sets of worlds corresponding to the possible answers to the question (i.e. {{PQR; PQR; Pqr; PqR}; {pQR; pQr; pqr; pqR}})



What are the possible answers to why S ?: the set of sets of sets of worlds corresponding to the possible sets of answers to why S ?

For the sake of simplicity I've only included two sets of possible answers, corresponding to i) *whether p or not p is why S ?* (blue) and ii) *which of p , r or both is why S ?* (orange) giving us {{{PQR; PQR; Pqr; PqR}; {pQR; pQr; pqr; pqR}}; {{PQR; PqR}; {PQR; Pqr}; {pQR; pqr}}}.




Fig. 3 Propositions, questions, and option questions

Providing advice for the design and conduct of collective inquiries ought to be a central goal for social epistemology. This paper has been descriptive, but it sharpens up some normative questions. When should a group distribute subquestions, and when should they double up on the same sub-questions? How should groups manage different kinds of division of labour? When should a group pursue multiple ways to carve up the same object question or multiple methodological approaches at the same time?

Many philosophers of science use formal models for thinking about the epistemic structure of collective inquiry (See Boyer-Kassem et al. 2018). It would be worthwhile to consider how these models relate to GROUP INQUIRY, and what light this account might shed on the significance of normative results drawn from formal modelling to group inquiry.

Acknowledgements Thanks to Neil Barton, Chris Bertram, Dominic Coote, Catarina Dutilh Novaes, Margaret Gilbert, Katherine Hawley, Megan Hyska, Justin Snedegar, Fenner Tanswell, John Matheson, Richard Pettigrew, Andrew Peet, Alison Pease, audiences in the Mathematical Collaboration I and II workshops in Oxford and St Andrews, the ECT seminar in St Andrews, ECAP 2017, and the 2017 UAM summer school in Social Epistemology.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aad, G., et al. (2015). ATLAS collaboration, CMS collaboration. *Physical Review Letters*, 114, 191803.
- Andersen, H., & Wagenknecht, S. (2013). Epistemic dependence in interdisciplinary groups. *Synthese*, 190(11), 1881–1898.
- Anderson, E. (2006). The epistemology of democracy. *Episteme*, 3(1–2), 8–22.
- Anderson, E. (2016). The social epistemology of morality: Learning from the forgotten history of the abolition of slavery. In M. Fricker & M. Brady (Eds.), *The epistemic life of groups: Essays in the epistemology of collectives* (pp. 75–94). Oxford: Oxford University Press.
- Babbage, C. (1832). *On the economy of machinery and manufactures*. London: Charles Knight.
- Beaver, D., & Stanley, J. (2019). Towards a non-ideal philosophy of language. *Graduate Faculty Philosophy Journal*, 39(20), 501–545.
- Bird, A. (2010). Social knowing: The social sense of 'scientific knowledge'. *Philosophical Perspectives*, 24(1), 23–56.
- Bird, A. (2014). When is there a group that knows? Scientific knowledge as social knowledge. In J. Lackey (Ed.), *Essays in collective epistemology* (pp. 42–63). Oxford: Oxford University Press.
- Bishop, M. (2005). The autonomy of social epistemology. *Episteme*, 2, 65–78.
- Bishop, M. A., & Trout, J. D. (2005). *Epistemology and the psychology of human judgment*. Oxford: Oxford University Press.
- Boyer-Kassem, T., Mayo-Wilson, C., & Weisberg, M. (2018). *Scientific collaboration and collective knowledge: New essays*. New York: Oxford University Press.

- Bratman, M. (1987). *Intention, plans, and practical reason*. Stanford: Center for the Study of Language and Information.
- Bratman, M. E. (1992). Shared cooperative activity. *Philosophical Review*, 101(2), 327–341.
- Bratman, M. (2014a). Rational and social agency: Reflections and replies. In M. Vargas & G. Yaffe (Eds.), *Rational and social agency: The philosophy of Michael Bratman* (pp. 294–344). New York: Oxford University Press.
- Bratman, M. E. (2014b). *Shared agency: A planning theory of acting together*. Oxford: Oxford University Press.
- Bromberger, S. (1992). *On what we know we don't know: Explanation, theory, linguistics, and how questions shape them*. Chicago: University of Chicago Press.
- Campbell-Kelly, M., Croarken, M., Flood, R., & Robson, E. (2003). *The history of mathematical tables: From sumer to spreadsheets*. Oxford: Oxford University Press.
- Cappelen, H., & Dever, J. (2019). *Bad language*. Oxford: Oxford University Press.
- Carballo, A. P. (2016). Structuring logical space. *Philosophy and Phenomenological Research*, 92(2), 460–491.
- Castelvecchi, D. (2015). Physics paper sets record with more than 5,000 authors. *Nature News*. <https://doi.org/10.1038/nature.2015.17567>.
- Cetina, K. K. (1999). *Epistemic cultures: How the sciences make knowledge*. Cambridge, MA: Harvard University Press.
- Cho, A. (2011). Particle physicists' new extreme teams. *Science*, 333(6049), 1564–1567.
- Dang, H. (2019). Do collaborators in science need to agree? *Philosophy of Science*, 86(5), 1029–1040.
- de Ridder, J. (2014). Epistemic dependence and collective scientific knowledge. *Synthese*, 191(1), 1–17.
- Edwards, P. (2013). *A vast machine*. Cambridge MA: MIT Press.
- Fagan, M. B. (2011). Is there collective scientific knowledge? Arguments from explanation. *The Philosophical Quarterly*, 61, 247–269.
- Fagan, M. B. (2012). Collective scientific knowledge. *Philosophy Compass*, 7, 821–831.
- Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, 27(1), 145–174.
- Gilbert, M. (1989). *On social facts*. London: Routledge and Kegan Paul.
- Gilbert, M. (1990). Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy*, 15(1), 1–14.
- Gilbert, M. (2000). Collective belief and scientific change. In *Sociality and responsibility: New essays in plural subject theory* (pp. 37–49). Lanham, MD: Rowman & Littlefield.
- Gilbert, M. (2006). *A theory of political obligation: Membership, commitment, and the bonds of society*. Oxford: Oxford University Press.
- Gilbert, M. (2009). Shared intention and personal intentions. *Philosophical Studies*, 144(1), 167–187.
- Gilbert, M., & Priest, M. (2013). Conversation and collective belief. In A. Capone, F. Lo Poparo, & M. Carapezza (Eds.), *Perspectives on pragmatics and philosophy*. Berlin: Springer.
- Goldberg, S. (2016). A proposed research program for social epistemology. In P. J. Reider (Ed.), *Social epistemology and epistemic agency: Decentralising epistemic agency*. London: Rowman and Littlefield.
- Goldman, A. I. (1987). Foundations of social epistemics. *Synthese*, 73(1), 109–144.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Goldman, A. I. (2000). Veritistic social epistemology. *The Proceedings of the Twentieth World Congress of Philosophy*, 5, 107–114.
- Goldman, A. I. (2010). A guide to social epistemology. In A. I. Goldman & D. Whitcomb (Eds.), *Social epistemology: Essential readings*. Oxford: Oxford University Press.
- Goldman, A., & Blanchard, T. (2016). Social epistemology. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Winter 2016 Edition). <https://plato.stanford.edu/archives/win2016/entries/epistemology-social/>.
- Gowers, T., & Nielsen, M. (2009). Massively collaborative mathematics. *Nature*, 461, 879–881.
- Grier, D. A. (1998). The math tables project of the work projects administration: The reluctant start of the computing era. *IEEE Annals of the History of Computing*, 20(3), 33–49.
- Grier, D. A. (2003). Table making for the relief of labour. In M. Campbell-Kelly, M. Croarken, R. Flood, & E. Robson (Eds.), *The history of mathematical tables: From sumer to spreadsheets* (pp. 266–288). Oxford: Oxford University Press.
- Grier, D. A. (2005). *When computers were human*. Princeton: Princeton University Press.
- Habgood-Coote, J. (2017). Knowledge-how is the norm of intention. *Philosophical Studies*. <https://doi.org/10.1007/s11098-017-0931-4>.

- Habgood-Coote, J. (2019). Knowledge-how, abilities, and questions. *Australasian Journal of Philosophy*, 97(1), 86–104.
- Habgood-Coote, J. (forthcoming). Group knowledge, questions, and the division of epistemic labour. *Ergo: An Open Access Journal of Philosophy*.
- Hardwig, J. (1991). The role of trust in knowledge. *Journal of Philosophy*, 88(12), 693–708.
- Holton, R. (2008). Partial belief, partial intention. *Mind*, 117(465), 27–58.
- Hong, L., & Page, S. (2001). Problem solving by heterogeneous agents. *Journal of Economic Theory*, 97(1), 123–163.
- Hong, L., & Page, S. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46), 16385–16389.
- Huebner, B., Kukla, R., & Winsberg, E. (2018). Making an author in radically collaborative research. In T. Boyer-Kassem, C. Mayo-Wilson, & M. Weisberg (Eds.), *Scientific collaboration and collective knowledge: New essays* (pp. 95–116). New York: Oxford University Press.
- Hull, D. L. (1997). What's wrong with invisible-hand explanations? *Philosophy of Science*, 64(4), 126.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51, 227–232.
- Kitcher, P. (1990). The division of cognitive labor. *Journal of Philosophy*, 87(1), 5–22.
- Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford: Oxford University Press.
- Kukla, R. (2012). “Author TBD”: Radical collaboration in contemporary biomedical research. *Philosophy of Science*, 79(5), 845–858.
- Lackey, J. (2014). Socially extended knowledge. *Philosophical Issues*, 24(1), 282–298.
- Lackey, J. (2016). What is justified group belief? *Philosophical Review Recent Issues*, 125(3), 341–396.
- Landemore, H. (2012). *Democratic reason: Politics, collective intelligence, and the rule of the many*. Princeton: Princeton University Press.
- Levy, Y. (2013). Intentional action first. *Australasian Journal of Philosophy*, 91(4), 705–718.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Martin, U., & Pease, A. (2013). Mathematical practice, crowdsourcing, and social machines. In *CICM'13 proceedings of the 2013 conference on intelligent computer mathematics* (pp. 98–119).
- Mayo-Wilson, C., Zollman, K., & Danks, D. (2013). Wisdom of crowds versus groupthink: Learning in groups and in isolation. *International Journal of Game Theory*, 42(3), 695–723.
- Mayo-Wilson, C., Zollman, K. J. S., & Danks, D. (2011). The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4), 653–677.
- Miller, B. (2015). Why knowledge is the property of a community and possibly none of its members. *Philosophical Quarterly*, 65(260), 417–441.
- Muldoon, R. (2018). Diversity, rationality, and the division of cognitive labour. In T. Boyer-Kassem, C. Mayo-Wilson, & M. Weisberg (Eds.), *Scientific collaboration and collective knowledge: New essays* (pp. 78–94). New York: Oxford University Press.
- Pavese, C. (2017). Know-how and gradability. *Philosophical Review*, 126(3), 345–383.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15(4), 351–383.
- Portner, P. (2018). Commitment to priorities. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 296–316). Oxford: Oxford University Press.
- Roberts, C. (2009). Know-how: A compositional approach. In E. Hinrichs & J. Nerbonne (Eds.), *Theory and evidence in Semantics* (pp. 183–213). Berkeley: CSLI Press.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. In *Semantics and Pragmatics*, 5: 1–69 Originally published in Yoon, Jae-Hak and Kathol, Andreas (Eds.) OSU Working Papers in Linguistics, 1998 (Vol. 49).
- Roberts, C. (2018). Speech acts in discourse context. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 317–359). Oxford: Oxford University Press.
- Rossini, F. A., & Porter, A. L. (1979). Frameworks for integrating interdisciplinary research. *Research Policy*, 8, 70–79.
- Searle, J. (1990). Collective intentions and actions. In P. R. C. J. Morgan & M. Pollack (Eds.), *Intentions in communication* (pp. 401–415). Cambridge: MIT Press.
- Shapiro, S. (2014). Massively shared agency. In M. Vargas & G. Yaffe (Eds.), *Rational and social agency: The philosophy of Michael Bratman* (pp. 257–293). New York: Oxford University Press.
- Shieber, J. (2013). Toward a truly social epistemology: Babbage, the division of mental labor, and the possibility of socially distributed warrant. *Philosophy and Phenomenological Research*, 86(2), 266–294.

- Smith, A. (1776/1982). *The Wealth of Nations* S. Andrew (Ed.) London: Penguin.
- Solomon, M. (2005). Groupthink versus the wisdom of crowds. *Southern Journal of Philosophy*, 44(Supplement), 28–42.
- Stalnaker, R. (1999). *Context and content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.
- Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
- Steingart, A. (2012). A group theory of group theory: Collaborative mathematics and the ‘uninvention’ of a 1000-page proof. *Social Studies of Science*, 42(2), 185–213.
- Strevens, M. (2003). The role of the priority rule in science. *Journal of Philosophy*, 100(2), 55–79.
- Tollefsen, D. (2007). Group testimony. *Social Epistemology*, 21(3), 299–311.
- Tossut, S. (2014). Membership and knowledge. Scientific research as a group activity. *Episteme*, 11(3), 349–367.
- Tuomela, R., & Miller, K. (1988). We-intentions. *Philosophical Studies*, 53(3), 367–389.
- Wagenknecht, S. (2014). Opaque and translucent epistemic dependence. *Episteme*, 11(4), 475–492.
- Wagenknecht, S. (2016). *A social epistemology of research groups: Collaboration in scientific practice*. London: Palgrave Macmillan.
- Watson, D., & Floridi, L. (2018). Crowdsourced science: Sociotechnical epistemology in the e-research paradigm. *Synthese*, 195(2), 741–764.
- Weatherall, J., & Gilbert, M. (2016). Collective belief, Kuhn, and the string theory community. In M. Brady & M. Fricker (Eds.), *The epistemic life of groups: Essays in the epistemology of collectives* (pp. 191–217). Oxford: Oxford University Press.
- Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2), 225–252.
- Whitcomb, D. (2017). One kind of asking. *Philosophical Quarterly*, 67(266), 148–168.
- Wilholt, T. (2020). On knowing what one does not know: Ignorance and the aims of research. In J. Kourany, & M. Carrier (Eds.), *Science and the production of ignorance: When the quest for knowledge is thwarted* (pp. 195–218). MIT Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Winsberg, E. (2018). *Philosophy and climate science*. Cambridge: Cambridge University Press.
- Winsberg, E., Huebner, B., & Kukla, R. (2014). Accountability and values in radically collaborative research. *Studies in History and Philosophy of Science Part A*, 46, 16–23.
- Wray, K. B. (2000). Invisible hands and the success of science. *Philosophy of Science*, 67(1), 163–175.
- Wray, K. B. (2002). The epistemic significance of collaborative research. *Philosophy of Science*, 69(1), 150–168.
- Wray, K. B. (2006). Scientific authorship in the age of collaborative research. *Studies in History and Philosophy of Science Part A*, 37(3), 505–514.
- Wray, K. B. (2007). Who has scientific knowledge? *Social Epistemology*, 21(3), 337–347.
- Wray, K. B. (2014). Collaborative research, deliberation, and innovation. *Episteme*, 11(3), 291–303.
- Wray, K. B. (2018). The impact of collaboration on the epistemic cultures of science. In T. Boyer-Kassem, C. Mayo-Wilson, & M. Weisberg (Eds.), *Scientific collaboration and collective knowledge: New essays* (pp. 117–134). New York: Oxford University Press.
- Zollman, K. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.
- Zollman, K. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 2(1), 17–35.
- Zollman, Kevin J. S. (2018). The credit economy and the economic rationality of science. *Journal of Philosophy*, 115(1), 5–33.