

Consensus as an epistemic norm for group acceptance.

To appear in O. Palermos & D. Pritchard (eds.), *Socially Extended Knowledge*

Joëlle Proust

Institut Jean-Nicod, ENS, Paris

Abstract

What are the propositional attitude(s) involved in collective epistemic agency? In the general case of collective agency, an epistemologically instructive debate features two opposing camps on this question: the ascribers have defended an extended notion of belief, while the rejectionists have claimed that groups form goal-sensitive acceptances rather than beliefs. Addressing this question, however, requires providing responses to four preliminary queries. 1) Are group attitudes reducible to the participants' attitudes? 2) Is epistemic evaluation sensitive to instrumental considerations? 3) Does accepting that p entail believing that p ? 4) Is there a unity of epistemic rationality across levels: should group attitudes be in principle consistent with the attitudes of the individual participants? Both "believing" and "accepting as true", as applied to plural subjects, fail to provide satisfactory answers to these four queries. An alternative analysis for epistemic group attitude called "accepting under consensus" is proposed. This attitude is shown to reflect actual group agency, and to offer consistent and independently justified answers to the queries. On this analysis, an individualist epistemology cannot simply be transferred to collective agents.

Introduction

Our aim in the present chapter is to identify the propositional attitude(s) involved in collective epistemic agency, both at the level of the group and at the level of its participants. Collective epistemic agency is exemplified in scientific research groups, in team reasoning, or in discussion forums. Are there group-level propositional attitudes similar, or indeed identical, to their participants' attitudes? Or do groups have propositional attitudes that they alone can entertain? This question is made all the more pressing by current debates about the nature of knowledge. According to Virtue reliabilism, knowledge possession is explained by a causal relation between agents' cognitive competences and their ability to form true beliefs (Greco, 1999, Sosa, 2007). Such an epistemic relation is seen as the crucial condition that prevents an agent's beliefs from being accidentally true, or from being reliably true as a consequence of some external manipulation (Pritchard, 2005). In contrast with process reliabilism, which claims that an ability to form true beliefs may consist in stable and reliable automatic dispositions that an agent unknowingly possesses, virtue reliabilism requires that the ability be *integrated* within the character of the agent, i.e., in a "web of stable and reliable belief-forming processes" (Pritchard, 2010). A new question, however, that this condition has raised

is whether it is compatible with the social extended cognition thesis. Granting that most of our knowledge depends on collective epistemic agents, what kind of cognitive integration, if any, might apply to extended cognitive abilities? (Pritchard, 2010, Pritchard and Palermos, 2013, Palermos & Pritchard, 2016). The question of the nature of the propositional attitudes entertained by a group constitutes a precondition for addressing this question.

In the first section, we will examine the forms of joint agency that involve group intentionality. The second section will delineate our problem space. What are the constraints that an adequate analysis of group attitude should meet? How are these met by the classical proposals for group attitudes – group beliefs and group acceptances? In the third section, an alternative proposal will be made: groups can only form consensual acceptances. This proposal offers a coherent response to our four queries. It suggests that an individualist epistemology cannot simply be transferred to characterize socially extended knowledge.

1. Group Intentionality: preliminary definitions

1.1 Group intentionality (GI): the general case

As a starting point of our discussion, we might propose a definition of collective intentionality inspired by a definition of individual intentionality. Not every individual system has intentionality: there exist living organisms, such as bacteria, for example, that fail to use stable representations of the world to guide their behavior. Similarly, groups may fail to be intentional if they are structurally unable, as groups, to form stable and coherent attitudes across time, or if they do not need collective forms of representation to act jointly. Tango dancing with a stranger, for example, does not require a unique, agent-level epistemic or conative decision center. Each dancer relies on individually computed "coordination smoothers" to compensate for prediction errors between expected and observed joint behavior.¹ Group intentionality (GI), then, is not an automatic by-product of group agency. Reciprocally, group intentionality seems to have a constitutive link with group agency. Let us therefore propose this first definition:

¹ For a presentation of a minimal architecture for joint action, see Vesper et al. (2010).

(1) Group Intentionality (GI) consists in the ability to represent goals, form goal-conducive attitudes, and to preserve, integrate and revise them over time in the light of new attitude contents, and act on them.²

An immediate problem with the notion of goal-conducive attitudes, in this definition, is that although we can say in common parlance, for example, that "Brussels believes (desires, fears, etc.) that *p*", it is *prima facie* difficult to maintain that Brussels *literally* has a mind, a belief, a desire or an emotion about anything. One might be tempted to rephrase group intentionality as the ability *of the group's members* to form attitudes, and to preserve, integrate and revise them over time in the light of new attitude contents, and act on them. The worry, this time, is that the intentional attitudes (that is, attitudes with representational content) that are verbally reported as belonging to the group may result merely from an aggregation of individual attitudes. If the aggregation is performed by an external observer rather than by a collective activity, the resulting "group" does not need to have group intentionality, nor even group agency. For instance, in Francis Galton's example of the visitors to the 1906 Plymouth farmer's fair who entered the competition for predicting the weight of a certain ox after it was butchered and dressed, each visitor bought a ticket on which they wrote their own prediction. Thus there existed a group, whose members were the ticket holders, who are supposed to have accepted the rule of the game regarding prize attribution and aimed to give the best estimate. No further coordination was involved: no participant knew the others' estimates, or was even aware that the average assessment would be computed.³ "The group" so defined turned out, however, to Galton's surprise, to be more accurate in estimating the final weight of the dressed meat than any of its individual members: averaging the 787 legible responses offered an estimate closer to the truth than any single estimate. Why, then, should such a "group" not be recognized as endowed with intentionality (as a group)? An obvious reason is that its members never attempted to form a collective assessment, and hence did not form, integrate, evaluate, and revise attitudes *from a collective standpoint*. In contrast to our tango dancers, they did not have any minimal collective goal either. Let us call this kind of group "observer-relative", in contrast with the "observer-independent" kind, in which the attitude integration, or the behavioral coordination of the group results from a purposeful coordination. In the latter case, group members have a specific shared goal that requires cooperation to be attained.

² This condition is meant to be understood as a conjunction of these ability clauses. The corresponding dispositions may fail to be manifested in given contexts.

³ Galton, (1907).

We should, then, distinguish observer-relative, non-agentive groups such as a set of consumers, from observer-independent groups, which themselves include non-intentional agentive groups such as tango dancers, and intentional agentive groups such as research teams. Although these three kinds of group deserve philosophical examination, observer-independent intentional groups should form a privileged target of study in epistemology, political, and moral philosophy, because they alone can qualify as responsible collective epistemic agents. Corporations, scientific teams, national and supranational agencies need to perform team-reasoning.⁴ They aim to form consistent sets of attitudes, which may differ in content, and possibly in type, from those of their individual members. To capture this feature, we need our definition for GI to include a condition to the effect that a common goal is being pursued. Having a common goal supposes that each group member is aware that the other members share target attitudes or attitude contents about the particular collective goals being pursued, and in addition, knows that the others have a similar awareness.⁵ Let us define "common attitudes" as the attitudes that are not merely shared by every group member, but also represented by them as shared by the other group members. We might, then, rephrase (1) as follows:

(2) GI consists in the ability to represent goals, to form common goal-conducive attitudes, to preserve, integrate and revise them over time in the light of new attitude contents, and to use them to guide action.⁶

The definition of (2), as it stands, however, needs further clarification. Although a clause about common knowledge makes good sense, we need to understand what it means, for a group, to form common attitudes. Are groups themselves genuine believers, as proposed by Gilbert (1989) and Tollefsen (2015)? Or is a group's attitude rather distributed among its members, with each having the ability to take a "group-oriented" attitude, as proposed by "we-mode" theorists such as Raimo Tuomela⁷ (2005, 2007) or Gallotti & Frith (2013)? Or do groups have their own group-level acceptances and associated contents, independently of those entertained by their members (Meijers 1999, Wray 2002)?

⁴ For a review: see Pacherie (2011).

⁵ Tuomela (2007).

⁶ This condition is again meant to be understood as a conjunction of the ability clauses. For a defense of group intentionality along these lines, see Pettit (2003), pp. 181-2.

⁷ According to Raimo Tuomela (2007), group members are able to adopt a shared "we-perspective", from which they think and act as a group member. In the "I-mode" perspective, individuals think and act as private persons. A we-attitude, however, can be entertained either in the I-mode or in the we-mode. For Tuomela, a group is not an extra-agent over and above the group members (2007, p. 4).

Collective attitudes, as suggested above, cannot be analyzed independently of the task that groups are constituted to perform, and hence committed to carrying out. Groups have a "realm of concern" that inspires their goals.⁸ To save space, we will limit our analysis of group attitudes to those that directly concern epistemic agency. This restriction turns out to be theoretically justified. As will transpire, just as strategic world-directed individual actions need to involve one or more epistemic decisions, strategic world-directed collective actions require from the group that a number of epistemic decisions be made.⁹

A second adequacy condition for (2) is that groups with the described integrative properties exist, i.e., that some groups are indeed able to form a consistent set of common attitudes. As argued by Pettit (2003) and List & Pettit (2011), having common attitudes is not sufficient to ensure consistency, because even if a collective decides what to do on the basis of a commonly accepted majority rule, it will end up forming inconsistent attitudes over time unless specific ("premise-based") decision rules are devised. If inconsistent attitudes prevail across time at the group level, however, the group in question will not qualify as intentional.¹⁰ Keeping this constraint in mind, let us turn to the task of defining group epistemic intentionality. or

1.2 Group intentionality (GI): the epistemic case

Our definition of group intentionality above does not mention whether the goal is to act *on the world* or rather to *acquire knowledge* about the world. Making war, or promoting a new public policy, exemplify actions with world-to-group direction of fit. Like most actions requiring instrumental reasoning, however, they will also need to include epistemic actions in the planning process. Given that a group must adjust its common attitudes to the distributed information that is made available to it, the direction of fit is rather "group-to-group".¹¹ How, then, should collective epistemic intentionality be defined?

To explore this question, it may be useful to take individual epistemic agency as our functional template. In individual reasoning, one part of the activity is devoted to performing the task, while another is supposed to monitor performance. In collective agency, does a similar organization of epistemic labor exist? For example, let us assume that investigating

⁸ This term is borrowed from Tuomela & Tuomela (2003).

⁹ On the contrast between strategic, impulsive and habitual actions, see Proust (2015).

¹⁰ An individual with prevalent dispositions to inconsistent preferences and epistemic decisions should equally be denied intentional agency, because the latter is defined by action guidance through representations.

¹¹ There is also a sense in which the epistemic collective work aims to reflect evidence. Hence distributed information also has a group-to-world direction of fit.

whether the Ebola virus is likely to spread to a given area, would, ideally, require i) evaluating whether this issue can be reliably addressed by the collective entity in question, ii) identifying the various variables that might be predictive for the task at hand, iii) collecting data about each of them, iv) combining them as required by the best model available, and, finally, v) retrospectively evaluating the correctness of the output. These steps, however, are not as readily completed by groups as they are by individuals. Individual epistemic agents have a capacity that is difficult for groups to replicate: epistemic feelings are a built-in evaluative system able to track the reliability of the outcomes of epistemic actions. Individuals can feel that their perception is blurred; they can feel that they know, do not know, or vaguely remember something, or that they can solve a problem.¹² Groups, however, cannot sense epistemic affordances, because they do not perceive, remember or reason in a first-person way. Hence, they are structurally unable to entertain non-conceptual evaluative attitudes about their own current cognitive processes, as individuals do.¹³ Research has shown, however, that feedback, whether feeling-based (through individual reports) or provided by (collectively) available factual observation,¹⁴ is necessary for improving collective accuracy.¹⁵ Granting the role of feedback in epistemic agency, two routes are open to groups. First, they can capitalize on the epistemic feelings of individual members about feasibility and reliability, by communicating the associated confidence levels to one another. A second route is to use feedback from the world, i.e. use prior success and failure to assess epistemic competence (both of the group and of its individual members). In summary: epistemic groups are confronted with the self-evaluative requirements for truth-responsiveness. Although groups are unable to form epistemic feelings, they can collect internal feedback through specific organizational policies and decision protocols. The preceding discussion allows us to propose a preliminary definition of group epistemic intentionality (GEI):

(3) GEI consists of the ability to form common truth-responsive, and consistent attitudes based on distributed information about target epistemic issues, and to plan future actions on their basis.¹⁶

¹² On noetic feelings, see Koriat, (2000) and Proust (2013, 2015).

¹³ This observation is based on the case of noetic feelings, i.e., epistemic emotions, which only individual agents can experience. This case suggests that groups cannot experience pleasure, fear or shame. For group consciousness, see Schwitzgebel (2015) and List (forthcoming).

¹⁴ Such as citation indexes, CVs, and bibliometric measures used.

¹⁵ Bang et al. (2014), Henry et al. (1996), Lehrer & Wagner (1981).

¹⁶ For a defense of group intentionality along these lines, see Pettit (2003), pp. 181-2.

Definitions (2) and (3) characterize group intentionality in terms of "an ability to form common attitudes". Definition (3) proposes that plural epistemic agents are groups able to form and pursue common goals related to matters of truth, knowledge and justification, which requires them to form common truth-responsive and consistent attitudes. Definitions (2) and (3), however, leave unspecified the type of the common attitudes that collective intentional agents are able to entertain as groups. The structural differences in reliability monitoring observed above suggest that it is quite likely that the group attitudes involved cannot be identical to those which their members entertain in isolation. Before examining how to address this issue, however, we need to list the conditions of adequacy for a correct answer to our question. This method should illuminate our problem space, by spelling out what the underlying constitutive issues are.

2. Are individual and collective attitudes of the same kind?

2.1. Four preliminary queries.

We will first list these queries and indicate why they are relevant to our question. Addressing them in a way that is both individually appropriate and globally coherent will be a background constraint for the subsequent discussion of the alternative proposals.

A first query concerns the issue of reducibility or irreducibility of attitudes belonging respectively to the individual and to the collective levels.

Q1: Are the propositional attitudes of an agentive group reducible to the propositional attitudes of its participants?

The "*Compositional*" or "*Summative thesis*" takes it that group attitudes are reducible to the attitudes of the participants. The "*Irreducibility thesis*", in contrast, claims that collective attitudes cannot be reduced to those of the participants. It does not claim, however, that one and the same *type* of attitude cannot be present at both levels. What it claims is rather that the *content* of a collective attitude is not generally inferable from the contents of the individual members' attitudes.

Our second and third questions bear on an issue that is crucial both in the philosophy of mind and in epistemology:

Q2: Is group epistemic agency autonomous with respect to group instrumental agency?

Some philosophers have claimed that accepting a proposition p as true may in various contexts depend on the role that p has in practical reasoning, thus denying the autonomy of epistemic norms with respect to instrumental norms.¹⁷ There are weaker and stronger forms of such denial; accepting the thesis of autonomy, in contrast, does not admit degrees.

Q3: Is Clarke's *entailment thesis* true, according to which the act of accepting that p conceptually entails having the belief that p ?¹⁸ Q3 crucially relates to Q2 because the thesis of autonomy must be denied if we allow accepting as true to vary across contexts without influencing global belief.¹⁹ Hence individuals or groups may in some cases accept what they don't, or don't fully believe. Autonomists, in contrast, claim that accepting p as true entails believing it. Although it is rarely noticed in this context, Q2 directly affects the issue of the nature of collective knowledge.

Q4: Should there be a *unity of epistemic rationality* between levels, such that the epistemic states of each individual participant would in principle have to be consistent with the epistemic states of the group to which they actively belong? This issue is also extremely important for the stance already endorsed about Q1. Any negative response to Q1 will have to deal with the existence of clashes between what is accepted by the members of a group and by the group to which they belong. Is there a rational way of resolving such likely conflicts in terms of the attitudes respectively involved at each level?

The queries above will be shown to play a major role in the controversies about the nature of extended knowledge between so-called [belief] "ascribers" and "rejectionists". On one view, represented among others by Margaret Gilbert (Gilbert, 1989, 2002) and by Deborah Tollefsen (Tollefsen 2003, 2015), individuals and groups can both form beliefs. Among those rejecting this view, Brad Wray (2003) and Anthonie Meijers (2002) propose that groups are not equipped for *believing*, but rather just for *accepting* contents. This controversy, however, deals with the general properties of collective world-oriented action. Given our present focus on collective *epistemic* agency, we will discuss the arguments used in this controversy from an epistemic viewpoint, in order to explore the functional and

¹⁷ See Cohen (1992) and Stalnaker (1984).

¹⁸ For a defense of this thesis, see D.S. Clarke (1994).

¹⁹ As proposed in Stalnaker (1984).

semantic basis of what our definition (3) describes as "the ability to form common truth-responsive, consistent attitudes based on distributed information."

2.2. Belief as a collective attitude

2.2.1. Irreducibility (Q1)

Now let us examine how the ascribers have addressed Q1. According to Margaret Gilbert (2013), groups and members have the same *attitude type* available to them: they can form beliefs. Collective belief *contents*, however, are *not* reducible to members' belief contents: this is because collective beliefs do not need to satisfy a distribution condition. A group may have a belief that no individual participant has. The condition for a group belief to be present is rather that the group members "are jointly committed to believe that *p* as a body". This idea fits our definitions (GI2) and (GI3) above: intentional group members must have interconnected intentions to cooperate in order to attain a collective goal, be committed to achieving it, and have common knowledge of these conditions.²⁰ In scientific research, for example, when a collaborative plan of study is collectively established, team members are committed to following it, and entitled to expect the other members to be committed.²¹ They may well privately disagree about several points. While working together, however, they set their epistemic disagreements aside.

To summarize: Gilbert's expression of "X-ing as a body", where X is a propositional attitude of some kind, offers a non-reductive answer to Q1. Joint commitment in X-ing that *p* by the group members allows a "plural subject" to emerge. What are the additional grounds for claiming that X-ing *is* believing? First, common parlance: the concept of belief is standardly extended to collectives. Second, individual and collective beliefs are sufficiently analogous to recognize them as belonging to the same type of attitude in spite of the fact that they have their own specific features.²² What they have in common is their sensitivity to epistemic considerations, with truth as their goal. This analysis denies Q1 for two metaphysical reasons: (1) a plural subject comes into existence through joint action, (2) belief is instantiated by individuals and by groups. From these two points it is inferred that

²⁰ For a review, see Pacherie (2011).

²¹ Unless they have reasons to stop collaborating.

²² See Gilbert (2009), chapter 6, pp. 141 sq.

commitment to believe as a body may overcome private reasons to (dis)believe. This consequence makes Q4 a sensitive issue.

2.2.2 The disunity of rationality (Q4)

Why should it be rational for individual members to be committed to believe "as a body" what they do not privately believe?²³ Gilbert insists that individuals are not invited to try to believe what the group believes. What they should rather try to do is collectively *emulate a single body that believes that p*.²⁴ What exactly does such an emulation consist in, and does it jeopardize the unity of rationality?

To address these two questions, let us consider a case of group-level epistemic agency. Let us suppose that the World Health Organization (WHO) aims to predict the chances of seeing the Ebola Virus under control in Guinea in less than a year. In order to perform such a complex, interdisciplinary type of epistemic action, a committee will be assembled with the required joint competences. This committee will identify and plan various epistemic steps, such as collecting evidence from the local administration, from health professionals, and from anthropologists, devising mathematical models based on prior data evolution, considering various scenarios, discussing the reliability of outcome in each of them, and so on. Failure in a particular step would lead to a decrease in accuracy of the final outcome.

When a collective epistemic task involves various fields of inquiry, as is the case in our WHO group, members may not understand the methods used in some steps of the collective investigation. It is nonetheless essential that all of them be committed to playing their own role as well as they can: for example, screening the evidence they must collect for its accuracy and exhaustiveness. Each of them justifiably believes that the others expect him/her to be a careful and exhaustive data gatherer, and expects that the others will similarly apply the best methods available to perform their own role. The WHO, as an organization, also "counts" on the committee to produce a reliable output at the time agreed. Should the committee not produce the anticipated output, or use dubious methods to achieve it, steps would be taken either individually, against those who indulged in the use of these methods, or collectively, by excluding the whole group from any further expert study.

²³ As claimed by Gilbert (2009) and Tollefsen (2002, 2015) in the general case of action.

²⁴ Gilbert (2009), p. 157.

Our WHO group is jointly committed to using epistemic norms, and hence, it constitutes a plural subject. Is the group of experts, however, "committed to believe as a body" that P , where P refers to the claim that the Ebola virus will be under control in 90% of the target area within a year?²⁵ As objected by the rejectionists,²⁶ this condition seems to conflict with the definition of belief, an attitude that automatically tracks the truth: as emphasized by Bernard Williams (1973), one cannot wish to acquire a given belief. One can intend to deliberate about whether p , but one cannot intend to believe that p , or be committed to saying "I believe that p ". Once the group has deliberated about its conclusions, some of the experts might still believe (and cannot help believing) that key indicators are missing from the study, which, from their viewpoint, should affect the reliability of the common decision.

Confronted by the rejectionists with this objection, Gilbert agrees that the content of beliefs, whether individual or collective, cannot be "directly willed". Emulating a body that believes, she claims, does not require *adopting* its beliefs. Voluntarism applies to the decision to *emulate*, not to a decision to align one's own beliefs on those of the group. Although acting and speaking as a body are "in a sense, representational",²⁷ the individual agent representing the group "is not required to have any particular attitude toward the proposition that p ". Relying on emulation, however, presents three serious problems.

The first has to do with *knowledge acquisition*. Unless participants have individual beliefs that are relevant to the collective epistemic decisions, it is unclear how the latter can be justified. For a group to attain true beliefs, if such is their goal – as ascribers maintain –, it seems mandatory that only truth-conducive methods should be used in deliberating about what to collectively believe. All the evidence individually collected should be made available to the group. No expert should be allowed to hide or distort relevant facts, or merely shy away from having an individual opinion. These various requirements are *prima facie* difficult to reconcile with the emulation mode.

The second problem has to do with the unity of rationality. How can the gap be bridged between individual and collective agents? Our experts are not violating their normative obligations to the group by *individually* believing "not p ". A self-intimation such as "be committed to believe!", Gilbert recognizes, is inconsistent. The normative difference between

²⁵ For a proposal of transferring commitment from "believing as a body" to "using efficient procedures" in forming a collective belief, see Mathiesen (2006).

²⁶ Meijers (1999), Wray (2001).

²⁷ Gilbert (2009), p. 158.

the collective and the individual case boils down to a socially coercive rule concerning the expression of collective belief: "Present group intentions, preferences, decisions and beliefs *as if* they were those of an individual agent."²⁸ In a nutshell: a disunity of rationality results from the social pressure on individuals to think in certain ways, in collective contexts. No epistemic reasons, then, are able or expected to bridge the gap between individuals and group beliefs.

2.2.3 Epistemic autonomy partly denied

Finally, Gilbert's view on group belief jeopardizes her positive response to Q2 (epistemic autonomy). For Gilbert, the *epistemic norm* that regulates individual belief is truth, or accuracy: belief is by definition taken to be unconstrained in its content by any other norm, social, practical, or otherwise. Collective beliefs as Gilbert sees them, however, do not seem to enjoy the epistemic autonomy that individual beliefs have. Take her example of an improvement committee that believes that the campus needs a café merely because the Chancellor "is wedded to the idea": social and pragmatic considerations overcome epistemic reasons for believing that there is truly such a need.²⁹ These considerations, however, being unrelated to campus needs, the autonomy of epistemic norms in generating a belief is violated. This is indeed an unhappy consequence for scientific or political groups: if they are not committed to following exclusively epistemic norms in forming beliefs, and if, furthermore, participants are not committed to either retaining or re-examining their own doubts, then, this conception of belief legalizes the divorce between two standards of rationality by making inter-level inconsistency a basic, and frustrating epistemic condition.

An alternative would be to claim that, had our doubting experts in the WHO committee *not* maintained their own reasons for dissenting after the collective decision had been taken, *they* would be epistemically at fault. For they would update their own individual beliefs merely on the basis of a metarepresentation of what a majority of others believed to be the case (and metarepresenting the group's belief seems to be what Gilbert has in mind when she asks participants to emulate believing as a body). The recognition that a belief one has is not shared with others, however, is not a sufficient reason to revise it, even after the group has deliberated about it. Following a heuristic such as "in a group, believe what the group believes" would create epistemic inconsistency between the heuristically revised belief and

²⁸ Gilbert (2009), p. 160.

²⁹ Gilbert (2009), p. 152.

the unrevised one, because the latter is still activated through its own nexus of reasons.³⁰

Let's sum up. Ascribers need to align their responses to Q2 and to Q4. Concerning Q4, they need to recognize that the kind of emulation required for group level belief inevitably entails a potential disunity in the epistemic rationality of the individual participants. The latter are invited to decouple their own reasoning from the group, but they are not offered any epistemic justification for doing so. The reason offered is group coerciveness: a pragmatic not an epistemic justification.

2.3 Strategic acceptance as a collective attitude

A second line of argument, which Gilbert called "rejectionist", has emerged. It emphasizes that a group needs to take certain propositions to be valid premises in its practical deliberations.³¹ To accept that p is to adopt a certain policy. A policy has the function of maintaining coherence in practical reasoning episodes and stability over time – a function that beliefs are not tailored to deliver.³² In epistemic agentive groups, such a policy would stipulate that, once p is found to be true by the group's own accepted set of methods of epistemic evaluation (consensus, ad-hoc algorithms, cumulated evidence, weighted evidence, and so on), the group should take p as a premise in further collective inquiry and in reasoning about how to act. Note, however, that this policy may *also* include prudential, context-relative reasons (such as availability of research funding) for accepting or rejecting a proposition. Granting that the acceptances in question are policy-based, we will call them "*strategic acceptances*"³³ to be contrasted with strictly *epistemic* acceptances (expressing truth or likelihood). Their contextual, goal-relative character, and the inclusion of utility as a major constraint often conflict with evidential requirements. On this view, individual participants are individually able to believe what seems right to them. There is no entailment from accepting to believing. Furthermore, in contrast with the former view, there is no process of belief formation at the group level.

This proposal satisfies our Definition (3) for group intentionality. A *prima facie* argument in its favor is that it takes into account the difference between individual and group "thinking" processes. Let's briefly review its responses to our four queries.

³⁰ See Lehrer (1975), p. 141.

³¹ See Bratman (1999), Engel (1998), Meijers (2002), Wray (2001).

³² See Bratman (1999), p. 30.

³³ For empirical evidence in favor of this distinction, see Goldsmith & Koriat (2008). The distinction is analyzed in Proust (2012).

2.3.1. Irreducibility (Q1)

Irreducibility is defended by rejectionists along a line similar to that of ascribers. The primary idea is metaphysical: "Plural subjects are constituted by their [collective] goals".³⁴ Pursuing goals determines whether a plural agent comes into existence or disappears. Temporary absence of goals, in contrast, does not threaten survival of individual agents. The second argument is based on the epistemic reciprocity at work in groups. Collective goals constitutively create common commitments for defending what the group accepts "as a whole" and for acting on it. The binding, common nature of epistemic commitment does not fit with an aggregative conception.³⁵

The rejectionists, then, defend *two forms of attitude irreducibility*: both types of agents can form acceptances, but only individuals can form beliefs. The respective epistemic contents of collective and individual attitudes do not need to be the same.

2.3.2. Epistemic Autonomy rejected (Q2)

Rejectionists are quite close to ascribers concerning their response to Q2. Only individual thinking can be epistemically autonomous. Collective acceptances compromise between epistemic and instrumental rationality: *S* accepts *p* if *S* decides, *for pragmatic reasons*, to treat *p* as if it were true, or close to true.³⁶ Strategic acceptances, exemplified in prudential reasoning, or in conventional settings, where for example an attorney accepts that *p*, because she is professionally committed to taking *p* as a premise, although presented by rejectionists as a form of epistemic attitude (accepting as if true), clearly violates autonomy.

2.3.3. Rejecting Clarke's entailment thesis (Q3)

As a consequence, Clarke's *entailment thesis* has to be rejected. Collective acceptance may vary across contexts without influencing belief, because the former aims at truth-cum-utility, while the second exclusively aims at truth. This also applies to individual acceptance.³⁷ This response to Q3, again, directly affects the nature of collective knowledge. A group may accept a proposition (for example [that there will, within a year, be a substantial reduction of the unemployment rate]) because it is an instrumental condition that might promote its goals

³⁴ Meijers (2003), p. 369.

³⁵ Meijers (2002), pp. 73-74.

³⁶ See Meijers (2003), Wray (2003), p. 369. Cohen (1992), Engel (1998) and Stalnaker (1984) have pioneered the introduction of strategic aspects in the definition of acceptance.

³⁷ See Engel (1998), Stalnaker (1984).

(for example, being re-conducted in a public function). Obviously, this is an unpromising constraint for an epistemic group.

2.3.4. Disunity of rationality (Q4)

The ascribers take a cross-level disunity of rationality to be an inevitable consequence of the social coerciveness of group-thinking. The rejectionists explain it through the duality of rational norms. Groups need to make plans respecting instrumental norms when they decide what to accept; individuals are not similarly constrained when they form beliefs or acceptances. Again, no properly epistemic reasons are on offer to explain and justify why an individual participant should form *unrelated* epistemic attitudes when the group is meeting and when it is not. Second, the rejection of the entailment thesis is an effort to restore coherence in the position, but it has no independent justification.

In summary, rejectionists' proposal of acceptance as an alternative to collective belief leads to a poorly integrated conception of rationality. It rejects the autonomy of epistemic normativity. It furthermore admits incommensurability between *how* groups and individuals think, and about the contents of their epistemic decisions.

We propose to take a radically different view on three of our four queries. Autonomy, Clarke's entailment, and the unity of rationality, will be the claims underlying our defense of a third way of capturing collective epistemic attitudes.

3. An alternative rejectionist proposal: accepting under consensus

Our alternative proposal is the following: accepting under consensus – consensual acceptance (CA) – is the only epistemic attitude that collectives can entertain as collectives; furthermore, it is an epistemic attitude that only collectives can form. This attitude is a direct way of respecting our definition of GEI: it is a common attitude based on distributed information about target epistemic issues.

A first argument in favor of the proposal is that, as noted by the other rejectionists, a *human* group is not *equipped* to believe, because belief constitutively includes huge, holistic storage of information and automatic revision processes (including error monitoring and coherence monitoring), which only an individual mind (or possibly a future generation of distributed systems of computers) can instantiate. For a similar reason, a group is not

equipped to directly accept propositions *as true*. A major reason for endorsing CA is that it is in virtue of its architecture that a group is able to accept propositions *as consensual*.³⁸ A major teleological reason for performing collective epistemic actions is that they can generate more reliable decisions than the best individual participant can do. Hence, *if and only if* its constitutive rules are respected (these will be presented in 3.2.), a consensual acceptance that *p* is a highly reliable indicator that *p* satisfies a given epistemic norm (such as truth). Defending this proposal should start with a defense of pluralism about the epistemic norms of acceptance.

3.1. Range of acceptances

A traditional view in epistemology is that truth being the constitutive and single norm of belief, it is also the only norm worth being investigated.³⁹ This monist view may either concentrate on the synchronic analysis of belief and its norm of truth, or extend its focus to the diachronic, zetetic aspects of inquiry as epistemologically significant, because inquiry ultimately aims at truth.⁴⁰ A diachronic focus, however, reveals the diversity of epistemic actions involved in inquiry. Indeed a scientist in any given domain must be sensitive to evidence in order to generate true beliefs. But she also needs to monitor consistency, plausibility, economy, relevance, and so on. In all these cases, a form of acceptance is produced as a result of the corresponding epistemic activity, such as accepting as consistent, etc. Sensitivity to the variety of epistemic norms and the ability to select the right norm at a given step of inquiry are two major competences needed for any thinker to acquire true beliefs in a non-accidental way, and consequently deserve the status of intellectual virtues. Detecting fallacious syllogisms requires monitoring the coherence between premises and conclusion, rather than the truth of the conclusion, which is irrelevant.⁴¹ In contrast, confronting a memo with independent evidence requires monitoring the truth of what is reported, rather than merely its internal consistency. Truth itself can be monitored through a norm of accuracy (the memo must exclusively contain true facts, with uncertain facts being omitted), or through a norm of exhaustiveness (the memo must report all the true facts, with uncertain facts being

³⁸ On the epistemic value of a distributed cognitive architecture, see Section 1.1

³⁹ A predominant role for truth over other epistemic norms may be grounded on a conception of truth as the constitutive and single norm of belief (see Engel, 1998), or as the teleological norm governing a practice of inquiry, (see Axtell & Carter, 2008). The latter view leads to recognition of the role of intellectual virtues in the process of knowledge acquisition.

⁴⁰ See Alston, (2005), Axtell & Carter (2008), Morton (2006).

⁴¹ Thompson (1996).

included).⁴² To convey any idea, or understand what is said, cooperative communicators must be sensitive to informativeness, relevance, and ease of processing, and not merely to accuracy.⁴³

The role of consensual acceptance in generating group knowledge is ignored because, as emphasized by ascribers, research teams prefer to use truth-related idioms in their reports. It is arguable, however, that, being, as a group, a step away from the evidence and the reasons individually collected and evaluated, their proxy for truth is consensus. *The same holds for all the other types of acceptance for which norm sensitivity is based on individual experience.* A group collectively selecting a plan of action has to consensually accept its *relevance* to a goal and its *consistency*. These steps qualify as epistemic because the requirements of relevance and of consistency are not a matter of utility, but depend on objective inferential relations. These constitutive informational and inferential requirements for a collective plan to be formed, however, are deferred for their assessment to individually sensitive participants or to artificial devices (see Section 1.2). The only *direct* form of normative assessment that a group can perform, in all these cases, is consensual acceptance.

3.2. Consensus is not conformity

Consensual acceptance might raise a principled objection: Does it have epistemic value, or merely social value? If the latter, then is it simply a social motivation for affiliating to a group that drives this attitude, rather than an epistemic goal? Addressing this objection requires differentiating a norm of *conformity*, a social norm, from a norm of *consensus*, an epistemic norm. A social norm of conformity involves rules stating how one ought to behave. Consider how people are willing to adjust to constant changes in dress codes: conformity makes affiliation manifest, an instrumental condition for achieving other social goals. Epistemic norms, in contrast, do not constitute instrumental conditions for being affiliated to a social group, professional or otherwise; they rather constitute objective standards and constraints for processing information and preventing misinformation from infecting cognitive systems. Contrary to the claims of constructivists,⁴⁴ these standards and constraints do not respond to social motivations (they are applicable whether or not they have positive or negative social

⁴² This distinction has been proposed as a way of dissolving the preface and the lottery paradoxes. See Kaplan (1981) and Proust (2012).

⁴³ The conversational postulates proposed in Grice (1989) can be analyzed as epistemic norms also regulating rational thinking and animal signaling. See Proust (accepted).

⁴⁴ See Latour (1989). For a critical analysis, see Wray (2001), pp. 327 sq. and Goldman (2010).

effects), even when they serve them (collective and individual epistemic actions may be embedded in world-directed actions).

A second related objection is that the term of "consensus" has a descriptive rather than a normative meaning: it refers to the objective quantity and type of information (stored, predicted, inferred) that is shared (and recognized as shared, hence, common) in a given group at a given time. Anthropologists and pollsters, for example, aim to find out the majority opinion about a given issue in a given population. Why and when, then, should the term of "consensus" refer to a norm? To address this important objection, consider whether, and when, it makes sense, in a given use of the term, to distinguish "genuine" from "spurious" consensus. "Genuine convergence" on a proposition p emphasizes that the sources of decision about whether p is true are independent from each other for consensus to have normative significance.⁴⁵ This condition maximizes reliability in epistemic decision-making when truth or plausibility cannot presently be assessed. An example of normatively consensus-driven acceptance is offered by the aggregation of evaluations by various experts commissioned to investigate about global warming. Consensus is "spurious", in contrast, when there is only one source of opinion that influences the whole group. Psychological studies have shown that agents are often insensitive to this difference, and tend to assess a proposal by the number of times it was emitted, failing to check the epistemic independence of the proposers.⁴⁶ Consensus, then, should not be used to describe the epistemic situation, because the constitutive conditions for consensus are not instantiated: group decision has no epistemic value.⁴⁷ Just as nothing is actually perceived in a perceptual illusion, no epistemic decision is taken on the basis of spurious consensus (the decision is inspired by a confusion of information-based consensus with redundant conformity of opinion).

Let us now turn to the clauses that together constitute CA as a propositional attitude.

3.3. Consensual acceptance: definition

We now need to provide a definition capturing essential semantic and functional features of CA. We will subsequently examine how this analysis responds to our four initial queries. Five clauses characterize a consensus-based collective acceptance.

⁴⁵ See Douven & Riegler (2010), Koriat (2012), Yaniv & al. (2009).

⁴⁶ Weaver et al. (2007).

⁴⁷ Douven & Riegler (2010). Similarly, as stressed by Lehrer (1975), consensus arrived at "by intent", rather than "unsought in the search of truth", is epistemically irrelevant. (p. 141).

- (a) Consensually accepting is a collective epistemic action, and hence it is irreducible to individual attitudes.
- (b) Its goal is to make an epistemically rational joint decision concerning propositions that are not yet evidentially or demonstratively settled. The outcome of consensual accepting – an action – is a consensual acceptance – a propositional attitude.
- (c) Accepting a proposition under consensus presupposes the integration of multiple informational sources and an appropriate arbitration between conflicting epistemic decisions.
- (d) Once a CA/p is made mutually manifest, it is ipso facto part of the group's common ground for taking further collective actions, such as producing reports and preparing intervention plans. Individual meta-beliefs $B(CA/p)$ will also automatically become common knowledge. However, individual participants do not need to revise their own initial belief about p .
- (e) Once it has become common ground, a CA that p entails a collective commitment to rely on p in relevant further group actions.

In (a), agency is claimed to be a precondition for forming a group attitude, as discussed in section 1. An agentive group is able to control its own epistemic outputs by selecting the norm under which it operates. A plural subject, however, cannot automatically form beliefs, first-order desires, or similar passively acquired attitudes, because it lacks a direct functional connection to the world (such as perceptual evidence). It possesses, however, an indirect connection to the world: a disposition to extract shared information by collecting genuine consensus about propositions worth deliberating about.

Clause (b) states the relation between CA and the epistemic goal that it is meant to serve. Granting that a condition for epistemic (and instrumental) rationality is the relevance of the goals pursued, a group should ideally not spend its time and energy on forming consensual acceptances about well-established facts or demonstrations. Consensus, in other terms, must be informative to have epistemic value. Collective epistemic agency capitalizes on existing knowledge, as represented in individual minds (beliefs) and in their extended counterparts (reports, scientific publications etc.). This background knowledge should only become the focus of collective consideration if it conflicts with newly formed acceptances. A main

additional constraint is consistency of the various consensual acceptances performed by a given group.⁴⁸

Clause (c) states that the process of decision-making must respect the informational constraints constitutive of CA in order to produce reliable decisions. These constraints are worth closer examination. Forming a consensus-based decision rather than an individual-level truth-based decision about p is epistemically advantageous when uncertainty threatens the reliability of a decision. Uncertainty depends on two main informational parameters. First, the objective distinctiveness of the cues (evidence, reasons) in favor of P and not $\neg P$ directly affects the probability of giving a correct response.⁴⁹ The less distinctive the elements for deciding whether p or whether $\neg p$, the more advantageous consensus will be. Second, noisiness and biases in decision-making are subjective factors that independently influence the reliability of individual and collective decisions.⁵⁰ Agents are known to differ in their sensitivity, i.e. in their ability to monitor their subjective uncertainty. Sensitive agents make reliable decisions even when noise, i.e. ambiguity in evidence, is objectively present.⁵¹

Granting that subjective confidence is determined both by distinctiveness of cues and variance, a decision rule encouraging members' interactive arbitration based on their subjective confidence has been shown to considerably enhance reliability.⁵² CA is likely to be more reliable than individual-level acceptance, then, when i) it allows *more extensive evidence* to be collected, integrated and discussed by different participants (thus reducing the

⁴⁸ Deductive closure – the requirement that any proposition entailed by previous acceptances should be taken to be consensually accepted – may be violated in many cases, however, as is also the case in individual belief systems. Although uninformative consensual acceptances may look rationally permissible, they violate a principle of economy applied to collective effort, as reflected in Grice's first conversational maxim.

⁴⁹ In signal detection theory, this difference can be read in the shape of the distribution of responses (whether by a single individual or by a group): the larger its mean value, μ , the less likely it is that an erroneous decision will be offered (Sorkin et al. 2001). This probabilistic dimension captures distinctiveness, an externalist property related to epistemological safety, which Pritchard (2005) defines as follows: "For all agents, A, if an agent knows a contingent proposition p , then, in nearly all (if not all) nearby possible worlds in which she forms her belief about p in the same way as she forms her belief in the actual world, that agent only believes that p when p is true." (p. 163). Arguably, assessing safety in a given context ultimately relies on assessing distinctiveness.

⁵⁰ This is the standard deviation, σ , i.e. the variance in the distribution of the responses to one and the same stimulus, i.e. the noisiness in the decisions, at the individual or at the group level. Sensitivity of agents in a task is inversely proportional to the standard deviation in that task.

⁵¹ As is the case for safety, sensitivity as discussed in epistemology is a modal concept, defined as follows in Pritchard (2005): "For all agents, if an agent knows a contingent proposition ϕ , then the agent does not believe that ϕ in the nearest possible world or worlds in which $\sim \phi$ ". From the viewpoint of the attitude described here, sensitivity is restricted to actual events, such as neural firings, verbal decisions, etc. In contrast to the modal notion, sensitivity, as it is used here, extends to necessary propositions whose demonstration involves a number of demonstrative steps, where noise (and thus, error) is likely to be present. For example, the consensual acceptance of Andrew Wiles' proof of Fermat's Last Theorem has depended on highly sensitive referees, able to collectively detect errors or incomplete steps.

⁵² For an experimental study of the influence of the various processes of group decision-making, see Bahrami et al. (2012), Koriat (2012), Sorkin et al. (2001). For a conceptual study, see List & Pettit (2011). See section 3.3.1 below for a more extensive discussion.

potential impact of misinformation), and ii) when it is able to use participants' own confidence in their individual opinion in order to weight their impact on collective decision.⁵³ Well-designed group size has also been shown to play an important role in collective epistemic success.⁵⁴

Clause (d) is the central feature that allows consistency between collective and individual attitudes to be restored. As group's consensual acceptance that p does not entail that its members should believe p , nor accept p as true. What individual members do is a form of epistemic scorekeeping:⁵⁵ they automatically update their beliefs concerning the claims that win consensus, and those that don't. They monitor the ongoing consistency and relevance of the debates. This is what it takes for a group to form a consensual attitude about p . The stability of the scorekeeping (its "collective" property) is made possible by the fact that all the participants individually metarepresent the same collective decisions, which are representationally decoupled from their own first-order beliefs.

Clause (e) states the type and the scope of the commitments resulting from CA. The collective commitment involved in CA is not to update common ground, an operation that is automatically performed by individual participants once a CA is made manifest to the group, as stated in (d). A CA/p entails the commitment to act jointly in a way that is consistent with it in order to attain future goals (epistemic or non-epistemic). For example, when the WHO has consensually determined the current relevant facts about Ebola virus in a given area and planned an intervention, the group is committed to taking the steps that were collectively decided on. Collective commitment, then, as observed in section 2.2.1, has to do with acting or planning to act in a certain way. This collective commitment ensures that the decision does not remain idle. Stability and perseverance in epistemic action are two important dynamic virtues of an agentive group.

3.4. Addressing the queries

Let us now examine the responses to the four initial queries that our definition incorporates.⁵⁶

⁵³ See Bahrami et al. (2010, 2012) and Lehrer & Wagner (1981).

⁵⁴ Decision performance does not increase linearly with group size, probably through a decrease of members' efforts ("social loafing") and a higher homogeneity in their cognitive characteristics. See Sorkin et al. (2001).

⁵⁵ Lewis (1979).

⁵⁶ As some readers may have noticed, the responses to the four queries have not been presented in the same order in each of the proposals examined. The order of reasons was deliberately selected to drive the order of exposition, to make the derivation of the arguments clearer.

3.4.1 Irreducibility

Our proposal favors a stronger notion of irreducibility than the notion defended by the other two proposals. In contrast to them, irreducibility does not result from the possibility, for a group, to have beliefs or acceptances that no member has. Irreducibility results rather from the fact that, in virtue of their functional structure, i) only groups can form consensual acceptances and ii) consensual acceptances are the only epistemic attitudes that they can form. Obviously, participants will have to act jointly to reach consensus, and hence, extract the information distributed from the individual beliefs, acceptances, and the judgments of uncertainty that individuals can form on the basis of their own noetic feelings. But a group has no mind of its own, and only has an indirect functional connection to evidence.

3.4.2 Epistemic Autonomy (Q2)

The present proposal firmly endorses the claim of epistemic autonomy. Consider the case of an individual agent. If an agent senses an epistemic decision as goal-dependent, she should accept p as true when, in a given context, it is judged useful to accept it as true. But agents do not fall prey to this. Let us suppose that an attorney has discovered by careful investigation that her client is *not* innocent. Being committed to defending him, she will strategically accept that p ("my client is innocent"), while also personally accepting as true, and indeed believing, that *not* p . There would be a norm conflict if her strategic acceptance were performed under a norm of truth. But granting that the attorney merely aims to *convincingly pretend* that p obtains,⁵⁷ she need only monitor the coherence of p with her other assertions. In the collective case, a similar divide is present between objective (i.e. epistemically grounded) consensus about whether p , and strategic decisions about what to report and how to act.

It is true, as Morton (2006) emphasizes, that the selection of a given epistemic action is driven by desires and by needs.⁵⁸ But this does not need to jeopardize the epistemic purity of the outcome. The judgment concerning the correctness of a particular epistemic decision is not influenced by the reasons that motivated one to act in this or that epistemic way. It might have been better to aim for accuracy than for exhaustiveness: this is a practical mistake. But the correctness conditions for accuracy [only true responses] – or those for exhaustiveness – [all the true responses] do not adjust to circumstances. They may be correctly fulfilled even when the acceptance selected is useless for guiding action. This defense of autonomy applies to any epistemic action, whether performed by an individual or by a collective agent.

⁵⁷ As proposed in the literature on acceptance (Cohen 1992, Wray 2001).

⁵⁸ See Morton (2006).

3.4.3. Clarke's entailment thesis

Why are individual and collective epistemic agents not always able to form beliefs on the basis of their acceptances? Our proposal offers a principled answer. Accepting p under a norm of truth entails believing that p is true. When accepting p as true, a reflective acceptor should also believe that [s/he accepts p as true].⁵⁹ Clarke's entailment, however, fails to extend to other types of acceptance, including consensual acceptance. Let us see why. When accepting p under a norm N other than truth, a reflective acceptor should not believe that p ; s/he should only believe that [s/he accepts that p under N], or that [it is true that p conforms to N]. For example, when accepting a syllogism as valid, the acceptor does not have to believe its conclusion. S/he should rather believe [that this syllogism is valid]. The same holds, mutatis mutandis, for: [this answer is informative], [this issue is relevant], or [this proposition is consensual]. *A semantic ascent, then, bridges the gap between collective consensual acceptances and individual belief systems*: each individual participant will automatically form a metarepresentation – [this proposition has been consensually accepted], which offers an important, but defeasible cue about its probably being true. Hence, on the present proposal, Clarke's entailment thesis is preserved in the case of accepting as true. It does not hold, justifiably, for acceptances performed under a different norm.

3.4.4. Cross-level unity of rationality

A major argument in favor of CA is that it provides a representational substrate for the unity of rationality. Even when group participants have divergent opinions about p , they all should agree that what the group accepts as consensual is a final step in a collective deliberation with majority decision-making. Hence, a collective epistemic decision, when meeting the conditions of consensual acceptance, ought to be fully and reflectively shared by participants. Consensual agreement, here, is made possible by a common meta-belief, held by participants, that such and such consensual acceptances were formed, and by a common attention to the normative aspects of the decision process. Common knowledge about the structural epistemic features of this process in turn justifies individually accepting that if p is consensually accepted, it is likely to be true. Again, this justification is defeasible, because the decision process may be opaque in some crucial respects.

⁵⁹ The semantic ascent from p to [I believe that p] is justified to the extent that taking p as a premise constitutes the belief that p . On the inferential view (Byrne, 2011), the self-attribution of a belief is generated through a doxastic schema, i.e. an inference from (perceiving or deriving) p to the self-attribution of the belief that p . Considering the two ways of understanding *inference* (Recanati, 2002), however, the doxastic schema does not need to be deductive.

It is another argument in favor of the proposal that it makes the contribution of participants' meta-beliefs about the decision process explicit. Distributed cognition presupposes that individual members themselves pursue epistemic and practical projects, follow strategies of investigation, with their associated beliefs, desires, emotions and reasons. Individual mental states are discrete sources of information that collective agency needs to extract and weigh.⁶⁰ The parallel accumulation of group-level consensual acceptances and individual meta-beliefs about these acceptances epistemically motivates the participants to reconsider (i.e., revise, maintain or suspend) their own individual convictions in the light of group decisions. In this process, a dynamical feature, decoupling allows meta-believers to shift truth evaluation from the group's perspective to another perspective, most notably, to the viewpoints of the meta-believer him/herself or of an ideally reliable observer.⁶¹ Furthermore, participants' meta-beliefs about other members' respective strategies, the quality of evidence, and ulterior motives explain why they will associate with a given collective acceptance, and pursue their collective activity, or exercise their power to exit, as they should whenever consensual acceptance is not elicited in conformity with its correctness conditions.

Finally, the distinction between individual and group commitment also contributes to promoting the unity of rationality across levels. A collective epistemic group is committed to performing further consensual acceptances in a way that is consistent with those already made: this rule is implicit in the participants' common knowledge about what the group should do. No such commitment is attached to participants' epistemic agency. Individual cognizers may or may not revise their own first-order beliefs; they may stop participating if they have reasons to do so.

4. Epistemological implications

As observed above, consensual acceptance fails to directly track truth for two reasons: (1) information is not directly collected, but rather obtained from individual sources. (2) Sensitivity to error is not experienced, but is rather a product of a deliberate decision to monitor feedback. Still, a gain in reliability is likely when group structure, task decomposition, and decision-making processes are rationally organized. This suggests that virtue-reliabilism might, on this view, make sense for the epistemic group. A specific form of cognitive

⁶⁰ The weighting process should focus on the quality of the evidence rather than on the participants' collegially estimated expertise. For a discussion, see arguments by Schmitt (1985), Baird (1985), and Lehrer (1985).

⁶¹ See Lehrer (1985), Recanati (2000).

integration might establish the cognitive character of a group as a causal basis of its dispositions to form reliable consensual acceptances.

Admittedly, cognitive integration for groups is more limited in scope than that for individuals. The latter are endowed with perceptions, memories, emotions, beliefs, and dispositions to act that are interconnected in all sorts of ways. Groups only deliberate about the questions that are presented to them for adjudication. They can only form consensual acceptances about the propositions that might be relevant to their main epistemic target. However, cognitive character may be identified with a group's integrated abilities for acquiring, combining and evaluating information. A collective strategy for inquiring into a matter may be more or less adequate in targeting relevant issues and in aggregating evidence from various fields of expertise. The deliberative processes selected may be more or less appropriate to allowing a genuine consensus to emerge. The arbitration process may or may not include eliciting participants' level of confidence in the reliability of their own contribution. Because they directly determine the group's ability to pool information and to avoid decision biases, organizational features such as these instantiate a genuine form of collective cognitive integration.

Depending on their cognitive integration, groups may indeed be more or less sensitive to error, more discriminative, and more likely to acquire knowledge. Even if their role is to integrate or evaluate rather than directly acquire knowledge, they still deserve to be given credit, to a significant degree, for their ability to do so, over and above the credit attributable to their individual contributors. A key element in cognitive integration has been identified as cross-level rational integration. Meta-beliefs about consensual acceptances allow each participant to assess their plausibility, their consistency with her own beliefs, their relevance to various further problems, and so on. In the individual case, this evaluation shift is a key to what Sosa has described as *reflective knowledge*.⁶² The specificity of group reflective knowledge is that it depends on *cross-level* evaluation shifts. Only such shifts can allow collectives to extract all the cognitive and metacognitive information distributed among their members. Granting that collective knowledge requires evaluation shifts that only individual members can perform, an important epistemological task consists in determining which sets of collective dispositions or rules for deliberation might encourage them.

Acknowledgements

⁶² Sosa (2009).

I thank Adam Carter, Stéphane Deligeorges, John Michaels, Orestis Palermos and Nick Shea for providing useful comments and relevant references. I am grateful to Dick Carter for linguistic and stylistic advice.

References

- Alston, W. P. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*, Ithaca and London: Cornell University Press.
- Axtell, G. & Carter, J.A. (2008) Just the Right Thickness: A Defense of Second-Wave Virtue Epistemology. *Philosophical Papers* Vol. 37, 3 : 413-434.
- Bahrami B., Olsen K., Latham P. E., Roepstorff A., Rees G., Frith C. D. (2010) Optimally interacting minds. *Science*, 329, 1081–1085. doi:10.1126/science.1185718.
- Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. (2012). What failure in collective decision-making tells us about metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594), 1350-1365. DOI: [10.1098/rstb.2011.0420](https://doi.org/10.1098/rstb.2011.0420)
- Baird, D. (1985). Lehrer/Wagner Consensual probabilities do not adequately summarize the available information. *Synthese*, 62(1), 47-62.
- Bang, D., Fusaroli, R., Tylén, K., Olsen, K., Latham, P. E., Lau, J. Y. & Bahrami, B. (2014). Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and cognition*, 26, 13-23.
- Bratman, M. E. (1999). *Faces of intention. Selected essays on intention and agency*. Cambridge: Cambridge University Press.
- Byrne, A. (2011, June). Transparency, Belief, Intention. In *Aristotelian Society Supplementary Volume* (Vol. 85, 1, pp. 201-221). Oxford: Blackwell.
- Clarke, D. S. (1994). Does acceptance entail belief? *American Philosophical Quarterly*, 145-155.
- Cohen, J. (1992). *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.
- Douven, I., & Riegler, A. (2010). Extending the Hegselmann–Krause Model I. *Logic Journal of IGPL*, 18(2), 323-335.
- Engel, P. (1998). Believing, accepting, and holding true. *Philosophical Explorations*, 1, 140-151.
- Gallotti, M. & Frith, C. (2013). Social Cognition in the we-mode. *Trends in Cognitive Sciences*, 17, 4, 160-165.
- Galton, F. (1907). The Wisdom of Crowds. *Nature*, 1907, vol. 75, 450-451.
- Gilbert, M. (2002). "Belief and Acceptance as Features of Groups," *Protosociology*, Volume 16, 35-69. Reproduced in Gilbert (2014), pp. 131-162.
- Gilbert, M. (2014). *Joint Commitment. How We Make the Social World*. Oxford: Oxford University Press.
- Goldman, A. (2010) Why Social Epistemology Is Real Epistemology, in A. Haddock, A. Millar, and D. Pritchard, (eds.), *Social Epistemology*, Oxford University Press, pp. 1-28.
- Goldman, A. (2011) "Reliabilism", in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.
- Goldsmith, M. & Koriat, A. (2008). The strategic regulation of memory accuracy and informativeness. In A. Benjamin and B. Ross (Eds.), *Psychology of Learning and Motivation*, Vol. 48, pp. 1-60. San Diego, CA: Elsevier.

- Greco, J. (1999). Agent reliabilism. *Philosophical Perspectives*, 13: 273-296.
- Greco, J. (2000). *Putting Skeptics in Their Place*, Cambridge: Cambridge University Press.
- Grice, P. (1989). *Studies in the way of words*. Cambridge, Harvard University Press.
- Haun, D., & Tomasello, M. (2011). Conformity to peer pressure in preschool children. *Child Development*, 82(6), 1759-1767.
- Henry, R. A., Strickland, O. J., Yorges, S. L., & Ladd, D. (1996). Helping Groups Determine Their Most Accurate Member: The Role of Outcome Feedback. *Journal of Applied Social Psychology*, 26(13), 1153-1170.
- Kaplan, M. (1981) Rational acceptance. *Philosophical Studies*, 129-145.
- Koriat, A. (2000). The feeling of knowing: Some metatheoretical implications for consciousness and control. *Consciousness and Cognition*, 9 (2), 149-171.
- Koriat, A. (2012). When are two heads better than one and why? *Science*, 336(6079), 360-362.
- Latour, B. (1989). "Clothing the Naked Truth." In Hilary Lawson and Lisa Appignanesi (eds.), *Dismantling Truth in the Post-Modern World*, New York: St. Martin's Press pp. 101-128.
- Lehrer, K. (1975). Social consensus and rational agnology. *Synthese* 31: 141-160.
- Lehrer, K. (1985). Consensus and the Ideal Observer. *Synthese* 62, 1: 109-120.
- Lehrer, K., & Wagner, C. (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical Study* (Vol. 21). Springer Science & Business Media.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8 (1), 339-359.
- List, C. (forthcoming). What is it like to be a group agent?
- List, C. & Pettit, P. (2011). *Group agency*. Oxford: Oxford University Press.
- Mathiesen, K. (2006). The epistemic features of group belief. *Episteme*, 2(03), 161-175.
- Meijers, A. (1999). Believing and Accepting as a Group, in *Belief, Cognition and the Will*, A. Meijers ed.. Tilburg: Tilburg University Press. 59-71
- Meijers, A. (2002). Collective agents and cognitive attitudes. *Protosociology*, 16, 70-85.
- Meijers, A. (2003). Why accept collective beliefs? Reply to Gilbert. *Protosociology*, 18 (IV), 377-88.
- Michael, J. (2011). Interactionism and mindreading. *Review of Philosophy and Psychology*, 2(3), 559-578.
- Morton, A. (2006). Knowing what to think about: When epistemology meets the theory of choice. Stephen Hetherington (ed.), *Epistemology Futures*, Oxford University Press, 111-30.
- Pacherie, E. (2011). Framing Joint Action. *Review of Philosophy and Psychology*, 2 (2), pp.173-192.
- Palermos, O. & Pritchard, D. (2016). The Distribution of Epistemic Agency. In P. Reider (ed.), *Social Epistemology and Epistemic Agency: De-Centralizing Epistemic Agency*. Lanham, MD: Rowman & Littlefield.
- Pettit, P. (2003). Groups with minds of their own. In F.F. Schmitt (ed.), *Socializing Metaphysics*, pp. 167-193. Lanham: Rowman & Littlefield.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175(1), 133-151.
- Proust, J. (2012). The norms of acceptance. In B. Roeber and E. Villanueva (eds.), *Action theory, Philosophical Issues*, Vol. 22, 316-333. Reprinted in Proust, (2013), ch. 8.
- Proust, J. (2013). *The Philosophy of Metacognition: Mental agency and self-awareness*. Oxford: Oxford University Press.
- Proust, J. (accepted). The evolution of communication and metacommunication in primates. *Mind and Language*.
- Recanati, F. (2000). *Oratio Obliqua, Oratio Recta*. Oxford: Blackwell.

- Recanati, F. (2002). Does linguistic communication rest on inference?. *Mind & Language*, 17 (1-2), 105-126.
- Schmitt, F. F. (1985). Consensus, respect, and weighted averaging. *Synthese*, 62(1), 25-46.
- Sorkin R. D., Hays C. J., West R. 2001 Signal-detection analysis of group decision making. *Psychol. Rev.* 108, 183–203. doi:10.1037/0033-295X.108.1.183
- Sosa, E. (2009). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Vol. II. Oxford: Oxford University Press.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, Mass.: Mit Press.
- Stalnaker, R. (2002) Common Ground. *Linguistics and Philosophy* 25: 701–721, 2002.
- Thompson, V. A. (1996). Reasoning from false premises: The role of soundness in making logical deductions. *Canadian Journal of Experimental Psychology*, Vol 50(3), 315-319. <http://dx.doi.org/10.1037/1196-1961.50.3.315>.
- Schwitzgebel, E. (2015). If materialism is true, the United States is probably conscious. *Philosophical Studies*, 172, 7, 1697-1721.
- Tollefsen, D. (2003). Rejecting rejectionism. *Protosociology*, 18, 19, 389-405.
- Tollefsen, D.P. (2015). *Groups as Agents*. Cambridge, UK: Polity press.
- Tuomela, R. (2005). We-intentions revisited. *Philosophical Studies*, 125: 327-369.
- Tuomela, R. (2007). *The Philosophy of Sociality: The Shared Point of View*. Oxford: Oxford University Press.
- Tuomela, R. & Tuomela, (2003) Acting as a Group Member and Collective Commitment. *Protosociology* 18, 19, pp. 7-65.
- Vesper, C., Butterfill, S., Knoblich, G., & Sebanz, N. (2010). A minimal architecture for joint action. *Neural Networks*, 23(8), 998-1003.
- Weaver, K., Garcia, S. M., Schwarz, N., & Miller, D. T. (2007). Inferring the popularity of an opinion from its familiarity: a repetitive voice can sound like a chorus. *Journal of Personality and Social Psychology*, 92(5), 821.
- Williams, B. (1973). Deciding to believe. In *Problems of the Self*. Cambridge: Cambridge University Press. 136-51.
- Wray, K. B. (2001). Collective belief and Acceptance. *Synthese* 129: 319–333, 2001.
- Wray, K. B. (2003). What really divides Gilbert and the rejectionists. *Protosociology*, 18, 19, 363-376.
- Yaniv, I., Choshen-Hillel, S., & Milyavsky, M. (2009). Spurious consensus and opinion revision: Why might people be more confident in their less accurate judgments?. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(2), 558.