



University of Southern Denmark

Group knowledge

A real-world approach

Klausen, Søren Harnow

Published in:
Synthese

DOI:
[10.1007/s11229-014-0589-9](https://doi.org/10.1007/s11229-014-0589-9)

Publication date:
2015

Document version:
Submitted manuscript

Citation for published version (APA):
Klausen, S. H. (2015). Group knowledge: A real-world approach. *Synthese*, 192(3), 813-839.
<https://doi.org/10.1007/s11229-014-0589-9>

Go to publication entry in University of Southern Denmark's Research Portal

Terms of use

This work is brought to you by the University of Southern Denmark.
Unless otherwise specified it has been shared according to the terms for self-archiving.
If no other license is stated, these terms apply:

- You may download this work for personal use only.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying this open access version

If you believe that this document breaches copyright please contact us providing details and we will investigate your claim.
Please direct all enquiries to puresupport@bib.sdu.dk

Group knowledge: a real-world approach

Søren Harnow Klausen

Received: 20 April 2014 / Accepted: 20 October 2014
© Springer Science+Business Media Dordrecht 2014

Abstract In spite of the booming interest in social epistemology, explicit analyses of group knowledge remain rare. Most existing accounts are based on theories of joint intentionality. I argue that this approach, though not without merit or useful applications, is inadequate both when it comes to accounting for actual group knowledge attributions and for purposes of meliorative social epistemology. As an alternative, I outline a liberal, de-intellectualized account, which allows for the complex distribution of epistemic states typical of most real-world collectives, and makes minimal requirements as to the psychological underpinnings of collective states of knowing and the formal features of groups. The account is inspired by theories of distributed and extended cognition. It is guided by the principle that we should use the same standard when dealing with social and individual epistemology. Careful attention to what is normally required—and, in particular, not required—for attributing knowledge to individuals lends support the more liberal view.

Keywords Social epistemology · Group knowledge · Social knowledge · Distributed cognition · Joint intentionality

1 Introduction

There has been a steadily growing interest in the social dimension of knowledge. Following the pioneering work of Gilbert (1989), a range of analyses of joint intentionality, especially joint belief, has been developed, and philosophers have, in various ways,

S. H. Klausen (✉)
Department for the Study of Culture (Philosophy), University of Southern Denmark,
Campusvej 55, 5230 Odense M, Denmark
e-mail: harnow@ifpr.sdu.dk; harnow@sdu.dk

extended these analyses to the realm of the epistemic.¹ Somewhat surprisingly, there have been few attempts to develop an explicit account of group *knowledge* as such. The somewhat idiosyncratic analyses provided by Corlett (1996) and Tuomela (2004) have remained rare exceptions, and though the topic has attracted more attention recently, it is still most often left implicit or treated en passant.²

In this paper I shall explore the prospects for a theory of group knowledge. I will argue that theories of joint intentionality, while admittedly apt at covering some important cases, are ill suited to function as the basis for a comprehensive and realistic account. A more liberal, less formal and less intellectualist approach is needed, both to account for many ordinary ascriptions of knowledge to groups and to develop a suitable theoretical framework for the evaluative and meliorative project associated with social epistemology. Drawing on inspiration from cognitive ethnography and extended cognition, I provide a sketch of such an account. The highly permissive approach I am advocating comes at a price, however. It might be said to have counterintuitive implications, at least in its more radical forms, and I am not quite sure myself how far we should, eventually, push in this direction. But since the limitations of the existing accounts seem obvious, possible alternatives deserve serious exploration.

2 The reality of group knowledge

It remains a matter of controversy how seriously and literally we should take talk about collective knowledge. The notion of a group mind appears dubious to many. The most widely cited outspoken skeptic is Quinton (1975). There have been few explicit attacks on the idea of a group mind in recent decades, but a latent skepticism is still widespread among mainstream analytic epistemologists and philosophers of mind.³ The view that “group mind” is really nothing but a metaphor can also be found in organization and communication science (e.g. Pavitt 2003a, b), and it has been dismissed as “mystical” in reviews of the sociology of knowledge (Kuklick 1983, p. 295). When conjoined with the very widespread and arguably intuitive *psychology requirement* for knowledge—only things with the appropriate mental states (e.g. beliefs) can be subjects of knowledge—skepticism about the reality of group minds casts doubts on

¹ See e.g. Schmitt (1994) on joint *justification*, Mathiesen (2006) on the *epistemic features* of group belief, or List and Pettit (2011) on the ability of group agents to *track truth*.

² This holds even for work that appears to be explicitly concerned with group knowledge, like the papers of Goldman (2004) and List (2005). Symptomatically, Lackey’s (2012) account of “paradigmatic instances of group knowledge attributions” is presented more as a means for correcting certain views of knowledge in general than as positive theory of collective knowledge. Of all the extant accounts, Bird (2010) comes closest to mine, both by being fairly directly concerned with collective knowledge and, not least, by being more liberal than the received view, acknowledging the division of epistemic labor typical of most real-world cases. An earlier, sketchier version of my own account was presented in Klausen (2010). Tollefsen (2002) has defended a reliabilist account that bears some resemblance to mine, but remains focused on group intentionality and justification on group beliefs. Hardwig’s (1985) ideas about “epistemic dependence” also pointed in a similar direction.

³ Huebner (2014) advocates a partial skepticism, arguing that group mentality is, in principle, a sound notion, but that it is seldom instantiated in the real world, and that many attributions of mental states to groups found in the literature should be read rather as a shorthand for more complicated claims about the group members. Rupert (2005, 2011, forthcoming) also argues for a tentative skepticism.

the reality of group knowledge as well. This is probably a main reason why work on collective knowledge has centered strongly on its psychological underpinnings and issued in theories of joint belief, group agency and the like.

I think this strong emphasis on collective psychology has been a mistake. But before going on to diagnose the limitations of the received view, let me present some reasons for taking talk about group knowledge seriously. The main such reason is, I submit, the pervasiveness and apparent seriousness of collective knowledge attributions. Such attributions are very common, and they play important roles in our explanatory and evaluative practices. We explain actions and events with crucial reference to the knowledge or lack of knowledge of groups and institutions, and we form and justify moral and political judgments on the basis of such knowledge ascriptions. Here are some samples of collective knowledge attributions:

- (a) “By then, the Russians knew how to build the bomb”
- (b) “The CIA did not know the identity of the drone attack victims”
- (c) “The CIA did not know of the attack in advance”
- (d) “The government did not know what they [sic] wanted from the website”
- (e) “They knew immediately that their lunar landing would have to be aborted”
- (f) “Ensure that the crew knows how to handle the boat should the captain not be on board”
- (g) “What cell biologists already knew is that the EGF receptor–ubiquitin complex binds to a protein called Hrs”
- (h) “He saw his theory as following logically from what biologists already knew about natural selection”

All these utterances appear to have been meant quite seriously. Some—(a), (e), (g) and (h)—figure in explanations of historical events. Some—(b), (c), (d)—have been used as premises in political, ethical or legal reasoning. The imperative (f) has practical and presumably also legal consequences, since it places a certain responsibility on the person in charge of the boat.

It may be objected that the sincerity of an utterance says nothing about whether its constituent expressions should be taken literally or metaphorically (and vice versa). We often use metaphors for serious communicative purposes. Moreover, eliminativists in other fields of philosophy are well aware that the notions they urge us to dispense with—e.g. psychological or moral notions—are used quite seriously, with full ontological commitment, by the folk and even by some alleged experts. Talk about witches also had serious practical and legal consequences.

It is of course true that pervasive and serious talk about an entity does not show it to be real. It shows, at most, that we—those who talk so—are presently committed to its being, in some sense, real and important. But the objection highlights a point of general significance: Collective knowledge should be treated fairly. That is, we should avoid using a *double standard* and not treat group knowledge more restrictively than other phenomena of the same type. In particular, we should not be more demanding in our treatment of group knowledge than in our treatment of individual knowledge. So the question must be, not if our attributions of group knowledge show it to be among the basic furniture of the world (which it almost surely is not), but if they can be taken just—or almost—as seriously as our individual knowledge attributions. It should be

borne in mind that individual knowledge is, arguably, not any natural kind. Though the opinions of philosophers differ, most contemporary views depict knowledge as being in some way dependent on human interests, often including the *attributor's* interests.⁴

Maybe skeptics are right that the requirements for knowledge can never be met. Maybe eliminativists about psychology are right that notions like belief should be abandoned, in which case knowledge will probably have to go as well. As long as these possibilities, serious though they are, have not led philosophers to abandon individual knowledge as a legitimate object of study, they should not be considered particularly damaging to theorizing about group knowledge.

It is surely not a sound objection to taking collective knowledge seriously that it can (perhaps) be *reduced* to states of the individuals and their relations to their environment.⁵ Reduced entities are no less real than their reductive basis. Thus List and Pettit are, for example, quite right in defending the compatibility of methodological individualism and realism about groups (2011, pp. 2ff.).

When this is granted, the metaphoricity issue also dissolves. It is characteristic of metaphors that they can be paraphrased into more literal expressions. I doubt that most speakers will find translations of group knowledge attributions into individual knowledge attributions (assuming that such translations are possible) more “literal” than the original expressions. But at any rate, the availability of alternative, more or less equivalent expressions does not indicate that the original expressions do not signify something real. The same goes for the tendency (demonstrated by Phelan et al. 2012) to replace group names with plural pronouns. So as long as our attributions of group knowledge seem to be on par with our individual knowledge attributions, they can be assumed to be sufficiently (putatively) real to merit philosophical analysis.

A second reason for taking talk about group knowledge seriously is that it has obvious practical benefits. Knowledge is a positive, evaluative notion, and we are (mostly, if not always) interested in getting more of it. It is possible that we could achieve the same practical goals by using only individual knowledge attributions. But it would be impractical; and we would still need to pay close attention to the social context of knowledge. While there is room for discussion about the priority of knowledge as an individual versus knowledge as a social phenomenon, there can be little doubt that the largest meliorative potential lies in altering the social context (and, more generally, altering the environment in which our—allegedly—individual pursuits of knowledge are carried out). It is very likely that a robust notion of collective knowledge (and related notions of collective error and ignorance, collective epistemic virtues etc.) will help such endeavors. I submit that not only do we have notions of group knowledge and use them quite extensively, because they serve important purposes; we could probably do better by refining them and using them more extensively, since we would thereby heighten our critical awareness of the role of social arrangements and processes in our pursuit of knowledge and, ultimately, in our practical endeavors.

⁴ Stanley (2005) and Foley (2012) are explicit about this, as are (of course) proponents of contextualism or those inspired by Craig’s “reliable informant”-account (1990). But even more orthodox invariantists and absolutists must agree that knowledge is *in some respects* a socially constructed entity.

⁵ As Pavitt seems to assume (2003a, b).

3 Desiderata and expectations

We should now be able to formulate some of the central desiderata for a theory of group knowledge and at the same time point to some of the complications and possible limitations that are to be expected. A good analysis of group knowledge, should as far as possible, enable us to

- (1) Account for common usage of the term “knowledge” as applied to collectives (i.e. actual *ascriptions* of collective knowledge)
- (2) Assess the epistemic performance of collectives and suggest possible improvements of our social arrangements

Obviously, these two desiderata—the first a concern for *descriptive adequacy*, the second a concern for the *practical utility* of the analysis—might conflict with each other and thus have to be balanced. As so often, it seems reasonable to aim at “reforming” or moderately reconstructive conceptual analysis, which allows some tweaking of ordinary concepts and emphasizes certain aspects of common usage over others.

- (3) As with all theoretical enterprises, the theory of group knowledge should be maximally *comprehensive* and *unifying*. It should cover as many putative instances and kinds of collective knowledge as possible, showing them to share the same fundamental structure or having the same general function.

We should, however, keep our expectations modest. Even ignoring our interest in practical utility, there is no hope for a strongly uniform account of group knowledge attributions. A brief glance at the examples given above should make it clear that our actual group knowledge attributions differ considerably, and in several respects. For some, a simple *summative* account, according to which a group knows *p* if and only if each and every of its members knows *p* (Gilbert 1989, p. 241; cf. also Quinton 1975, p. 9), may be right (this seems to hold for (e)). For others, like (a)–(c), it is obviously insufficient. (I’ll say more about the strengths and weaknesses of the summative account, which might have been unfairly neglected, later on). A related way of characterizing the diversity is to say that some examples should be read “collectively” (i.e. as being about a group as such) and some rather “distributively” (i.e. about the individuals that make up the group). Note, however, that I will argue that even when a distributive reading is the most appropriate, it still makes sense to speak of a “group knowledge attribution”. Even in those cases where plural subject term, rather than a group name is used, the attributor must still be considering the individuals in question under *some* common aspect and thus as a “group” in a admittedly wide sense of the word (for more about this, see Sects. 7 and 8).

Group knowledge attributions differ particularly with respect to the (mostly implicit) understanding of the *subject* of group knowledge. They have different implications both when it comes to the *scope* of the attribution (which individuals are taken to matter, i.e. which members of a group are considered relevant for the epistemic assessment in question?) and the assumptions about the group’s *structure* (including the distribution of epistemic work and credit within the group). I suppose that “the CIA” does not refer to the same set of individuals in (b) and (c) (and that in none of the cases it refers to all members of the CIA). “The Russians” in a) probably refers to only

a tiny subset of all Russians, and maybe not exclusively to Russians or Soviet citizens. But it is implied that a certain structure of the collective obtained at the time, for example that individuals having made crucial contributions to, or sharing, the knowledge in question were appropriately linked with Soviet leaders and the relevant production capabilities. A secret society of dissident Soviet physicists, keeping their knowledge of how to build the bomb strictly for themselves, would not suffice for making “the Russians” knowledgeable of it.

This example also stands out by being one of collective knowing *how* rather than knowing *that*. But I do not consider this a main obstacle to unification. Leading contemporary epistemologists have proposed unified treatments of knowing how and knowing that (with [Hetherington 2011](#)) advocating a “practicalist” reduction of knowing that to competences, and [Stanley \(2011\)](#) (cf. also [Stanley and Williamson 2001](#)) conversely arguing for an “intellectualist” account of knowing how). I consider it a sensible desideratum for a theory of group knowledge that it should, *mutatis mutandis*, cover both collective knowing that and collective knowing how. But those who disagree can simply ignore the knowing how-cases. There is plenty of diversity to be found even in the field of propositional group knowledge attributions.

It remains to be seen how much can be achieved in terms of a unification. Some types of collective knowledge attributions may be ignored, if they prove resistant to an otherwise promising line of analysis and seem less central. But we have reason to assume that collective knowledge is a relatively diverse phenomenon, which comes in distinctively different forms. This is part of my rationale for advocating a more liberal approach.

Finally, a further constraint has emerged from my discussion of the reality of group knowledge:

- (4) The requirements for collective knowledge should not be significantly stronger than the requirements for individual knowledge (i.e. there should be no *double standard*).

I will not go so far as to argue that collective knowledge should be treated *more* permissively than individual knowledge, or that we should not look to individual knowledge at all. The latter remains the paradigm notion. The notion of collective knowledge *is* metaphorical in the (innocuous) sense that it is an *extension* of the notion of individual knowledge to a new domain (though in the case covered by the summative account, the transition is completely straightforward). But attending closely to the typical requirements for individual knowledge will show them to be less strong. And reflection on the attribution conditions for collective knowledge can shed light on our understanding of individual knowledge. For example, it might turn out that individual knowledge attributions also rest on implicit assumptions about the nature of the subject of knowledge).

4 Limitations of the received view

Since there are very few explicit, worked-out analyses of collective knowledge, it might seem strange to speak of a standard account. But there is a distinctive and influential, though mostly implicit, view to be found in the literature. A large body of work has

pointed in a common direction, even though the authors have rarely gone all the way. In many cases, they have simply left it to reader to fill in a few blanks.⁶ For example, the account of group justification presented by Schmitt (1994) is readily extendable into an account of group knowledge, by adding the truth of the proposition believed by the group (plus, if necessary, some fourth condition in order to cope with Gettier-cases). Much the same holds for theories of belief-aggregation, group rationality or collective truth-tracking. They are not exactly theories of group knowledge, but they appear to be concerned with the epistemic core of group knowledge (or of certain kinds of group knowledge).

There are many differences in detail between the various existing accounts, some of them far from insignificant (as a matter of fact, they are not all equally idealized or formalist; some are closer to my own alternative than others). Still, most extant contributions to collective epistemology share the following assumptions, and thus fall prey to the same criticism⁷:

- (i) Group knowledge requires *group belief*
- (ii) Group belief should be analyzed in terms of *joint commitment*⁸
- (iii) Joint commitment entails *reflexivity* (i.e. each and every member of an epistemic collective must view herself *as* a member of the collective in question).⁹
- (iv) Joint commitment requires *joint attention* (i.e. each and every member of an epistemic collective must be aware of—having focused her attention on—the same object(s), typically the target proposition (the object of knowledge), perhaps also the evidence for the truth of this proposition.

I will contest all four assumptions, focusing mainly on (iii) and (iv). Let me begin by noting that the standard account scores badly on all four criteria (i.e. desiderata) for an analysis of group knowledge stated above.

First, its descriptive adequacy is obviously limited. It can account for few if any of the real-world examples of collective knowledge cited in Sect. 2. Whoever “the Russians” in example (a) is taken to refer to, it is hardly a set of individuals jointly committed to speaking or acting in a certain way (again, one should not be disturbed by the fact that the example is about knowing how; in any case, a parallel example,

⁶ Though some proponents of the joint intentionality approach may also have acknowledged its limitations as a basis for a full-fledged, explicit theory of knowledge (Klausen 2010).

⁷ Examples include Schmitt (1994), Tuomela (2004, 2007), Mathiesen (2006) (though she requires epistemic collectives to be jointly committed to a *method* rather than to a believing a proposition), List and Pettit (2011) (though they make a number of mitigating statements—e.g. pointing to the potential benefits of distributed cognition (ibid. p. 97) and allowing that a group “may form and enact certain attitudes without all its members jointly intending that these particular attitudes be formed and enacted” (ibid., p. 35)—they do require that group members be “licensed by the group” (loc. cit.), and they do come very close to making a joint intention requirement [ibid., p. 33]).

⁸ The specific term “joint commitment” is associated with the work of Gilbert, but all those to whom I attribute the received view have adopted similar notions. Thus Schmitt speaks about “a willingness to act jointly” (1994, p. 260), and Tuomela uses the term “collective commitment” (2004, p. 113; 2007, p. 5). As I use the notion of joint commitment, it carries no specific Gilbertian connotations, but simply refers to a situation where each member of a group is committed some belief or action in a specifically “collective” sense, i.e. with a (perhaps implicit) understanding that it ought to be held by, or carried out by, the group.

⁹ See e.g. Tuomela (2007, pp. 20, 35).

which would be almost equally resistant to an account in terms of joint intentionality, could be constructed in terms of propositional knowledge). Nor can it be assumed that the biologists referred to in (g) and (h) met the requirement. While the biologists mentioned in (g) all attended to and accepted the same proposition (about the EGF receptor–ubiquitin complex), there was not necessarily any *common knowledge* (Lewis 1969, pp. 52ff.; Gilbert 1989, pp. 188ff.) about it, at least not among them all. And the biologists referred to in (h) probably did not all attend to a common target proposition. This knowledge ascription is most plausibly taken to mean that some biologists knew this and others knew that (and that therefore they, taken together, knew “something”, i.e. a complex body of propositions, which no single member had, or needed to have, attended to in its completeness). Hence there was no group belief (i), no joint commitment (ii) and no joint attention (iv) in this case.

The received view has been developed with cases like juries, boards and commissions in mind. Such collectives almost always meet the reflexivity requirement (iii), and in any case can be reasonably required to do so (if you are on the board or in the jury, you really ought to know and be aware of it, or else something is seriously wrong). They are also likely to meet the joint attention requirement (iv), though this need not always be the case, as absent-minded or uninterested members may fail to attend to the target proposition.

Juries, boards and commissions undeniably play very important roles in contemporary society, and so it is admittedly somewhat unfair to suggest that the received view is not a “real-world” account. But they are hardly the most common kind of social arrangement to be assessed in terms of its epistemic standing. And they invite a formalist and idealist approach by being entities that are given formal authority and occupy roles that generate particular ethical and legal obligations. Questions of accountability arise naturally in relation to juries, boards and commissions. Moreover, the serious practical consequences of their decisions make it natural to demand a high degree of concentration and responsibility from their members. Such concerns may be less important to the way we think about other potential subjects of collective knowledge and other epistemic task than those of settling questions of guilt or providing the knowledge basis for a controversial political decision.

What then, are the more typical subjects of collective knowledge? On reflection, lots of candidates come to mind. Project teams, working crews, office staffs, formal or informal organizational units or subsets of such units, like agencies and departments (cf. “The CIA did not know”), sports teams, research groups, scientific communities, etcetera. We do regularly attribute knowledge to such collectives; and we do care for their epistemic standing and may want to improve it.

Most of the knowledge we attribute to such collectives is *genuinely distributed* knowledge. Their members are not jointly focused on any particular proposition or task; they do not consider or even possess the same evidence; and they may not even be committed to the same procedure (*pace* Mathiesen 2006). They know things in virtue of their members’ knowing and doing different things. Hence they do not meet requirements (i), (ii) and (iv). Any comprehensive theory of group knowledge will have to dispense with these requirements.

The status of the reflexivity requirement (iii) is more open to discussion. Many of the more common collectives mentioned above do usually exhibit reflexivity. The

members of working crews consider themselves members of the crew; the same goes for members of sports teams and departments. As Pettit and List rightly points out, group membership need not be *endorsed* by the members; they can view themselves as such without wanting this to be the case (2011, p. 34). But even a member of a slave gang will be painfully aware of her gang membership.

Still, there are many cases of putative collective knowledge in which even the reflexivity requirement fails to be met. That the Russians knew how to build the bomb does not imply that any particular group of individuals at the time considered themselves members of that group. And a scientific community is not best demarcated according to the self-understanding of its members. It is more properly conceived as consisting of whoever is actually contributing to, or disposed to contribute to, some process of scientific research and communication. As with the biologists in example (h) above, a scientific community will often be delineated by the knowledge *attributor*, according to her specific interests. Often, scientists who consider themselves members of different communities—e.g. representatives of different disciplines—are in retrospect treated as part of one of the same scientific community, if their efforts were brought together and contributed to a particular scientific development.¹⁰ Some might think that a research group would be a sure example of a reflective epistemic collective, but to my experience, this is not so always; there may be diverging views about who is on or not on the group. Even the CIA-cases are far from clear. When the CIA knows something, I suppose it often does so partly—but crucially—in virtue of epistemic contributions from persons who do not view themselves as representatives (or at least not as unambiguous representatives) of the CIA.

Now it may be said in reply that the received view is motivated precisely by the kind of mess that results from taking our ordinary collective knowledge attributions at face value. I have myself pointed out that we refer vaguely and imprecisely to the subject of ordinary group knowledge attributions. The concern for reflexivity may be raised further by the fact that knowledge is bound up with moral obligations, which in turn require that we identify a subject that can be held accountable and responsible (e.g. the CIA and the government *ought* to have known; they can be blamed for not having known, etc.). If collective knowers do not know whom they represent, or what they are assumed to know, placing responsibility becomes difficult.

There is something to this reply. It does not, however, suffice to motivate a requirement on group knowledge in general, but only on knowledge (or subjects of knowledge) in connection with particular interests, for example in moral obligations or specific institutional arrangements. And inasmuch as the received view is motivated by a fear that the subject of collective knowledge may become too disintegrated, and

¹⁰ Huebner (2014) acknowledges the distributed character of scientific work and the typical disunity of scientific communities, but argues (following Kukla 2012) that because of the resulting lack of accountability, the outcome of such processes should not count as knowledge (p. 214). Tollefsen (2014) rightly points out that this argument depends on an internalist view of knowledge as requiring access to reasons. Like she, I favour a reliabilist approach; but even a more relaxed version of internalism, which allows for the distribution of epistemic factors within the group, could license the ascription of knowledge to scientific communities or research groups. Huebner does acknowledge that in some cases of actual scientific collaboration, we might correctly ascribe genuine knowledge to the group in question (2014, pp. 250ff.), but only because it meets something like a joint commitment–requirement.

a corresponding wish for tying it together by demanding strong relations between the members of the collective, it is simply misguided. Remember desideratum 4): We should avoid a double standard. It may be correct that knowledge requires the presence of appropriate mental states (though even this turns out to be less evident than is often assumed). If so, we should look to the kind of psychological requirements that are generally made for knowledge, that is, for individual knowledge, our paradigm case. And it seems that people are not very demanding, but rather go by criteria that more or less match those of their—allegedly “superficial” or “unserious”—permissive collective knowledge attributions. They do care about subjects of knowledge, about their identity and internal structure, but are not particularly picky. For example, apart from a tiny minority of formalist or radically internalist epistemologists, they do not require any general reflective awareness on part of the subject. A person can know some fact without knowing whom she is or that she is a knowing subject. And we do not normally require any higher-order knowledge. The knowing subject need not know that she knows p in order to know p , nor does she need any knowledge of epistemological principles or of the justificatory efficiency of her evidence. Sensible proponents of epistemological internalism have long since moved away from such unrealistic requirements (see for example Bonjour in [BonJour and Sosa 2003](#), pp. 65ff.; [Conee and Feldman 2004](#), p. 75).¹¹

It is also commonly regarded as unproblematic that the ingredients of knowledge are distributed among, *inter alia*, different mental states of the subject, which do not generally embody representations of each other.¹² And it is accepted that subjects of knowledge can be short-lived. While opinions differs at to which representational contents may be possessed by swampmen or other creatures going suddenly in and out of existence, most will be inclined to grant such transient beings at least some knowledge. Since evidence can quickly come and go, and methods and rules may only be mastered for only a brief span of time, we further accept that a subject’s epistemic

¹¹ It may be said to be more of an open question whether there are higher-order requirements on individual knowledge. Quite a few philosophers have maintained such requirements for *belief* (for a survey and convincing criticism of such views—which can be found in the work of Sellars, Davidson, Shoemaker, Haugeland, Brandom and Williams—see [Kornblith 2012](#), pp. 42ff.). Still, it seems that there is a tendency within mainstream analytic epistemology to make relatively modest requirements in terms of reflection and metacognition, in order to be able to accommodate cases of unsophisticated knowers. Besides, if one is attracted to some kind of higher-order requirement, a suitable analogue in the field of collective knowledge would be to require that within the group, *some* individual has to be able and disposed to critically monitoring the group’s first-order processes—that is, a *distributed* reflective capacity. This is still significantly weaker than the full reflexivity requirement.

¹² [Phelan et al. \(2012\)](#) claim that there *is* a significant difference between attributions of mental states to groups and to individuals, because in the latter case, people do not tend to paraphrase statements about an individual’s mental states into e.g. statements about her neurons. But this could be because the ontological dependence of a group on its members is much more obvious and straightforward than the ontological dependence of e.g. a belief state on an assembly of neurons. And the contention of Phelan, Arico and Nichols that when we attribute mental states to an individual, we take them to be “fully hers” (p. 711), requires elaboration and is far from obviously correct (we probably do *not* assume them to be particularly “pervasive”, “central”, “integrated” or whatever else could be meant by that phrase, apart from their being simply *her* mental states). Moreover, the fact that I am concerned not with mental states as such, but with *knowledge* states, makes these considerations less relevant to my proposal.

standing vis-à-vis a certain proposition or task can change swiftly. We should bear this tolerance of ours in mind when thinking about group knowledge.

5 “Only connect”: real-world group knowledge and distributed connection

The general idea behind my proposal for an alternative theory has already been stated: We should allow that the factors which, together with truth (or, in the case of knowing how, some sort of adequacy to the task in question), are necessary and jointly sufficient for group knowledge, can be distributed among the members of the group. A well-known example of genuinely distributed cognition has been provided by [Hutchins \(1995\)](#), who describes how a navy vessel crew is able to navigate successfully through the concerted efforts of many individuals, each of whom carries out a very specialized task and does not necessarily have any knowledge of the contributions of others, nor of the more general tasks or the ways in which the different contributions are merged. I suggest that this kind of example, rather than that of a jury or a board of directors facing a specific decision, should serve as a paradigm of collective knowledge.

I intend my theory to be neutral with regard to the more specific requirements and thus compatible with both internalist and externalist theories of individual knowledge and justification. Hence I will assume that in each case of group knowledge, some set of sufficiently justificatory and/or reliability-conducive factors must be present in the group. For example, some members must employ certain methods or procedures, and/or some members must possess certain pieces of evidence obtained by perception, inference, testimony and the like. For the sake of simplicity I will call the set of all such factors—conceived as epistemic factors pertaining to the *individual* members—relevant to the epistemic standing of the group J (though this should not be seen as carrying any internalist connotations; J may also consist in the use of reliable procedures or the presence of truth-tracking abilities among the members of the group).

Connections also matter, however, and so J alone may not be sufficient for turning a true belief into group knowledge (or to put it otherwise: group knowledge does not supervene on J). Consider a situation where Linda has a true belief p (but no relevant evidence for its truth), Ruth has some partial evidence e_1 for p , Martha has some other partial evidence e_2 for p , and e_1 and e_2 are jointly sufficient for turning the belief that p into knowledge. If Linda, Ruth and Martha do not interact properly, the necessary ingredients, though present *in* the collective, will fail to add up to collective knowledge (though different patterns of interaction would suffice, like Linda consulting the others in turn, or all three sharing all their information openly between them).

That connections matter more than the total amount of justification or epistemic excellence in a collective is well known in both theory and practice. Probably the CIA had both sufficient evidence and mastery of inferential rules and other relevant methods; its failure to know came from its inability to piece the elements together and/or direct the information to the relevant officials (and it may not have known even if some CIA officials did actually form a justified true belief about the matter, in case it was not adopted in the more central branches of the organization). And the members of a group may be inferior to the members of another, in terms of their individual competences, but more than compensate for that by cooperating smartly. The Condorcet Jury The-

orem shows that a group of moderately competent persons may outperform another group that bases its decisions on the verdict of a superiorly competent expert member, if they employ a voting procedure (Condorcet 1785; cf. List and Pettit 2011, pp. 86ff). While the Jury Theorem holds only under conditions that seldom obtain in the real world, it does serve to illustrate the potential for boosting collective performance by concerting the individual contributions, and thus the importance of relations within the group (a range of real-world examples can be found Surowiecki 2004).

I will call the set of epistemically relevant relations holding between the individuals in a group (and their epistemic task or output, like, for instance a target belief or action) R . Since not all of J may have an impact of the epistemic standing of the group as such—cf. the possibility of non-shared evidence or abilities that are not applied to the task in question—the total epistemic standing is not simply the product of J and R . The collectively relevant part of J , which I will call J_c , is determined by R . While there is reason to stress the importance of R over J as such (and thus the social aspect of group knowledge), the contribution of J should not be neglected, since the total epistemic standing of the group is a product of J_c and R (plus the truth of the proposition or, in cases of knowing how, the task adequacy of the ability in question). If the members of two groups, G_1 and G_2 , are similarly connected and disposed towards each other, but J_c is stronger or more extensive in the case of G_1 , then G_1 is epistemically better off than G_2 . It is also possible for a relationally inferior group (i.e. a group whose members are linked less smartly) to outperform another group simply in virtue of having more competent members—a highly competent expert may, according to the Jury Theorem, outperform a collective of very modestly competent persons. Of course a certain minimum of connectedness is required for making the individual factors matter at all. But the converse is also true: No collective procedure or social setting is able to compensate for a complete lack of evidence, competence or reliable sources of individual belief formation.

It may be objected that sometimes all that is required for attributing knowledge to a group is that one of its members knows. In an important work that ought to be widely read by social epistemologists, economic historian Joel Mokyr stipulates that “society ‘knows’ something if at least one individual does” (2002, p. 4). Yet as Mokyr’s own subsequent, more nuanced discussion betrays, this is hardly appropriate. It is correct that *under the right circumstances*, it may suffice that only one individual knows. We may say that in 1905 (or at least shortly after), scientists knew the photoelectric effect, even if Einstein’s discovery had not yet been widely transmitted or accepted as such—because it quickly became accepted and fed into other new and fruitful lines of discovery. Because Einstein was, in spite of his marginal position as a patent office clerk, sufficiently well connected to the rest of the scientific community, and because this community was, in spite of widespread skepticism about the idea of quantized energy, sufficiently perceptive and accommodating (and thus “stood poised” to adopt the new knowledge), it makes sense to say that science and thus all of society (assuming that this was also sufficiently perceptive) knew already the moment Einstein came to know.

But under different circumstances, individual knowledge will not do. If an individual is isolated from the rest of society, or not believed by her peers (and unable to change their mind), then society does not share her knowledge. Probably some so-called

“conspiracy theories” are actually true (see [Coady 2012](#), pp. 110ff.), yet the knowledge possessed by their staunch, but ridiculed and marginalized believers does not spill over to society in general. So a single individual’s knowing p does not by itself suffice for making the group to which the individual belongs know p . Connections and social organization matter even in the cases of apparently one man-driven collective knowledge. Hence group knowledge is always a function of J and R.

The Einstein example highlights the fact that not only the *actual* connections between the members of a group (i.e. their causal interaction) matter; so do their *dispositions* to react in various ways. It matters for collective knowledge attribution that the social environment of a knowable or competent individual is perceptive; and it probably also matters that it is not *uncritically* perceptive, i.e. that it would *not* be disposed to adopt the knowledge, evidence or procedures in question if there were not good reasons for it.

This whole picture accords strikingly well with our ascription practices and criteria for individual knowledge. We implicitly (and sometimes explicitly) require appropriate connectedness for individual knowledge as well. A memory M may be epistemically relevant to (i.e. potential evidence for) a belief that p , and it may be present in the individual who believes p . But if M is hidden somewhere in the deep corners of her mind and thus very hard to retrieve, it should not be counted among the evidence for p which she possesses at the time ([Feldman 2004](#)). Similarly, though I probably know the premises for many interesting conclusions, as well the relevant inference rules, failure to connect these elements (or being disposed to so) prevents me from actually knowing these conclusions. The case of knowing p , but in an insulated way, remaining ignorant of its implications for one’s further beliefs, closely mirrors the case of a person who, because of her lack of appropriate relations to her social environment, fails to bring her epistemic assets to bear on the group as a whole.

On the other hand, standard epistemologies do allow that an individual might know even if the relevant evidence is not instantly available. A certain recall or inference time is considered permissible. It should be born in mind that knowledge is, quite generally, a dispositional notion. We can know things we are not occurrently conscious of, perhaps (though this is more controversial) things we have *never* thought of, like the number of windows in one’s house (for further discussion, see [Bonjour 2002](#), p. 30). It does not matter for my present purposes exactly where to draw the line. The important point is that most epistemologists are willing to grant that individual knowledge is to some degree, and in some respects, dispositional. This means that in the collective case, there may, similarly, be knowledge that is not expressed in any occurrent belief or actually exercised skill, and that factors that have to be retrieved or skills and competences that have to be “turned on” contribute to a group’s epistemic standing. What is required is availability, not presence or actual use.

My proposal bears some similarity to the account of extended cognition developed by [Clark and Chalmers \(1998\)](#).¹³ Consider their famous example of Otto, the Alzheimer’s patient who compensates for his failing memory by carrying a notebook

¹³ [Bird \(2010\)](#) draws the same analogy.

with him everywhere he goes, writing down new and looking up old information. Clark and Chalmers quite plausibly suggest that his notebook functions like a kind of outsourced memory. But if this is so, then cognitive functions can likewise be outsourced to other people in one's environment. Consider the case of Abby, the doctor, and Tom, her medical assistant. Tom receives the results of the patients tests from the lab, compiles and sorts the data and gives relevant information about e.g. blood values to Abby, who in turn interprets the data and arrives at a conclusion concerning the patients' health condition. As long as Tom and Abby are working in sync, with Tom providing easily accessible information to Abby, quicker and more reliably than she would be able to gather and sort it herself, they should be seen a parts of a single cognitive system and so a legitimate unit of epistemic assessment and knowledge attribution. Orthodox individual epistemology would of course describe the situation as one in which Abby alone acquires knowledge about the patients' health, based on evidence she acquires from Tom. But it seems no less—and perhaps even more—appropriate to treat Abby and Tom as a small collective subject of knowledge as long as they are occupying their specific roles.

6 Epistemic assistants and executives

What about the *belief* condition for propositional knowledge? Here we have different options, and may opt for a more or less radical position. If we take a conservative line, wanting to depart as little as possible from the traditional ways of understanding knowledge, we might require that at least one member of the group actually believes *p*. This member acts as the group's *spokesperson* (even if she doesn't have to actually express her belief. And it should be born in mind (again) that beliefs are dispositional. Hence the "spokesperson" may not even have had the occurrent *thought* that *p*.

A more appropriate and practically useful way of describing the situation may be to say that within an epistemic collective, we can distinguish between epistemic *executives* and epistemic *assistants*. The person forming the "target" or "output" belief is a clear example of an epistemic executive, as she brings the epistemic task to its fulfillment. In the case of Tom and Abby, Tom is clearly the assistant and Abby the executive. It is not required that the epistemic executive has individual knowledge, since she might not possess the relevant evidence (though as we have seen, the situation might often be analyzed as a case of knowing from testimony, or forming beliefs based on a reliable source, so the executive will tend to acquire individual knowledge as a sort of likely byproduct of the collective process).

We should, however, be open to the possibility of cases where no single person believes *p*, but the members of the group are disposed to act in a way that justifies attributing to the whole group a knowledge that *p*. Maybe the taxation agency knows where I live and how much I earn, though no single employee has the corresponding belief. If one is keen to maintain the belief requirement for propositional knowledge, this can be seen to imply attribution of a group belief as well—group belief in a weaker sense than as defined by the theories of joint intentionality, but consistent with standard theories of behavior-based belief-attribution. Alternatively, one might see it as an example of knowledge without belief.

Even if this possibility should be acknowledged, the distinction between epistemic executives and assistants remains widely applicable. In most groups, and in relation to most epistemic tasks, there will be a more or less hierarchical structure or a certain chain of supply and delivery. Though no one may be directly concerned with the target proposition or output action, some are likely to be more directly concerned with some of the subtasks than others. The roles of assistant and executive are relative. The same person may act as executive, merging the contributions of assistants into a solution to a subtask, and as assistant, passing on the solution as input to the solution of a further task. Note, however, that there is nothing in the notion of an epistemic executive that implies that such a person is any more knowable of the group's epistemic task than her assistants. She may, for example, give the right answer without knowing what question it is the answer to.

In a very insightful recent treatment of group knowledge, Jennifer Lackey discusses an example that nicely illustrates both the idea of an epistemic executive and the point that there need not be any such person (though this is not quite how Lackey herself describes it)¹⁴:

DISTRIBUTED INFORMATION: The UN Population Commission, which is comprised of forty-seven individual members, issues a report entitled *Charting the Progress of Populations*. Each member of the group was responsible for collecting information about a different segment of the population represented in the document, and their respective work was done entirely independently from one another. The information contained in the report is, then, widely distributed across the members of the group. Sam, who is not a member of the UN Population Commission, was hired to interpret and compile all of the data contributed by the members of this group into the published report and to serve as the group's spokesperson. One of the statements in this report is, "the birth rate of Latinos in the US is on the rise," of which not a single member of the UN Population Commission is aware (Lackey's 2012, p. 256).

Lackey rightly argues that the epistemic credit should go to the (members of the) commission, though I find the exclusion of Sam from the group rather arbitrary, since he obviously functions as an epistemic executive. He is more than a mere spokesperson, as he also interprets and compiles the data and draws the conclusion. Still, even if we leave Sam out of consideration, and even if we grant that not a single member of the group is aware of the target statement, it is not implausible (at least on some possible interpretations of the situation) to ascribe knowledge of it to the group. Since the information contained in the report was widely distributed among the members of the group, they might have been sufficiently *disposed* to form the corresponding belief. If they were *not* sufficiently disposed to so—if too much effort was required in order for them to actually form the target belief, or to act in a way that would justify attributing it to them—then the group cannot be said to have known. But in any case, the epistemic credit for the statement should go to all of those who contributed to the epistemic

¹⁴ It is a version of an example first given by Tollefsen (2007).

task in question, regardless of their official roles and regardless of their position in the chain of epistemic supply.

7 Group identity and cohesion

The received view is strongly focused on the *group* as such—on its identity and the ties that bind its members together. In contrast, I have as yet made no specific requirements on groups as such; I have simply taken their existence and potential for acting as bearers of knowledge for granted. But what makes a group a group in the relevant sense?

I see no need to add any further constraints in order to secure group identity or otherwise knit the members of the group more firmly together. Accordingly, I have rejected both the reflexivity and the joint intention requirements. My contention is that *collectively contributing* to the state of knowing in question suffices to delineate a group in the relevant sense, i.e. an *epistemic collective*. The members of the group can be united solely by the fact that they all contribute to solving the same *epistemic task* (knowing p , knowing how to Φ , justifying p , testing p etc.). They are often united in all sorts of other ways, but they need not be so, and their further social bonds are not directly relevant to the group's identity. Though it may seem controversial in the present context, individuation of groups and organizations on the basis of tasks or goals is very common in organizational theory, where it appears to be more or less the standard view (see e.g. [Scott 1992](#), p. 10; [Donaldson 1995](#), p. 135; [Kieser and Walgenbach 2010](#), p. 6).

This view of group identity does generate a problem, however. In my discussion of the importance of factors J and R, I considered cases where members of a group had epistemic assets, which, due to their disconnectedness, did not matter to the standing of the group as such. This description presupposed that such epistemically inert persons were nevertheless members of the group. But if only epistemic contributions determine group membership, then they should not be counted among the members of the group at all.

However, I have only claimed that epistemic connectedness *suffices* for group affiliation. It is not necessary: An ignorant or ignored employee may belong to a group due to the organizational role she is assumed to occupy. It is an important point, and a corrective to the currently dominant view, that legal and official status do not in themselves matter to the demarcation and epistemic assessment of collectives. Yet our practical interests, and the ways in which we have organized society often make it natural for us to associate or even identify the subject of group knowledge with some “official” collective entity.

Again, there is a parallel to individual knowledge. Barack Obama's office as president of the United States does not in itself have any bearing on his epistemic standing on any issue. This is purely a matter of the nature of and connections between a subset of his mental states and the world. Yet we do ascribe his epistemic accomplishments and shortcomings to him both as a person—an entity comprising more than the mental states that constitute his knowledge—and as the occupier of an official role (cf. “the president did not know”). And Barack Obama's knowing or not knowing certain things obviously acquires a special importance as long as he holds the office as presi-

dent. Likewise, it is of special interest to us whether governments, agencies, firms etc. know.

In principle, we can single out any group of individuals, according to whatever criteria (or even at random, though there would be no practical point in that) and then evaluate it epistemically. Yet in practice, we go either by existing labels and boundaries, or we simply “follow the epistemic relations” even where these go across conventional borders, e.g. by including lay scientists or representatives of other disciplines into scientific communities.

It should be noted that R does not only comprise positive relations, e.g. “epistemic contributions” in the genuine sense, persons doing or possessing something that helps to improve the epistemic standing of a group. Negative relations can be included as well. If a member or a group is disposed to disapprove of some epistemically advantageous process, or to bring in irrelevant considerations or misleading evidence, then this, too, will be a part of R. The ignorance of a scientist’s peers with regard to her discovery contributes to the negative epistemic standing of her scientific community. It might be expected that negative relations will only be taken into account in cases where a group is demarcated according to other criteria than mere epistemic connectedness. For when negative relations are included in R, everyone becomes, per definition, related to all epistemic tasks, and thus unknowingly a member of all epistemic collectives, simply by being ignorant or inactive. Still, there may be some actual cases where negative epistemic relations to a target belief are used as the sole criteria for group demarcation, e.g. when we say that “most people do not know *p*”.

The highly permissive view of group identity may be the most provocative feature of my account. I will address some of the worries it engenders below. But it is supported both by the fact that quite a few of our group knowledge attributions are used in such a permissive way (e.g. scientific communities are demarcated according to the relevant contributions a certain task) and by the practical benefits of taking an inclusive stance on group membership. We are often interested in getting whoever is able to contribute to actually do so. A case in point is *Wikipedia*, which does not require any formal status or qualifications of its contributors, who even remain anonymous—all that matters is that they deliver and do it reliably. We may be interested in a more precise and rigorous subject identification for legal or moral purposes. But from a purely epistemic view, all that matters is that the job is done, not who does it.

8 Return of the summative account

Treatments of joint intention and group knowledge usually depart from a brief consideration of the summative account, which is quickly dismissed as being more or less inadequate.¹⁵ And clearly there are many cases of putative group knowledge for which it is not appropriate [like the examples (a) and (h)]. Hence it must be admitted that it cannot be a necessary requirement for group knowledge in general.

¹⁵ Corlett (1996) takes the opposite line, arguing—quite implausibly, because overly restrictive—that collective knowledge requires that each member of the collective knows the target belief individually.

Notice, however, that the summative account still does fit quite a few cases. For example, cases where a state of knowledge stems directly from sense perception¹⁶ obviously call for a summative account. Linda, Peter and Ruth saw the accident if and only if Linda saw the accident, Peter saw the accident and Ruth saw the accident. Nothing less will do, and nothing else is required. Example e) seems to be of this kind: Each of the Apollo crewmembers observed a set of facts, which made each of them believe that the landing had to be aborted. Here, “transversal” connections did not matter, because of the special—perfect—pattern of distribution. Without any coordination, the crewmembers were directed at the same proposition, and each of them fulfilled the justification requirement (or whatever else may be required) for knowledge of this proposition.

It is also far from obvious that the summative account is not *sufficient* for group knowledge (though this is often taken to follow from Gilbert (1989, Chap. 5) critical discussion). If no joint commitment is required, as I have argued, then it seems that we have everything we need in cases where all the individuals concerned know individually. “They know that I have left my wife” can be correctly applied to a group of people (e.g. to my colleagues) if they have learned in parallel and the knowledge in case is not yet *common* knowledge among them. Or to take another example: Imagine that Peter shot Paul; that he was seen doing it by Jim, John and Jack, who, however, saw it from it different locations and did not notice each other, but who where all seen by Peter. In this case, it makes perfectly good sense for Peter to think that “they (Jim, John and Jack) know that I shot Paul” and also make predictions about the group’s behavior based on this knowledge attribution (e.g. “they are going to tell the police”, “they are going to testify against me in court” etc.).

There may be cases where a lack of appropriate connections between the members of a group leads to a lack of “executive function”—that is, the group might not be disposed to act on, or otherwise manifest, the knowledge which each and every of its members possesses. Maybe a lack of common knowledge inhibits its actions. It is possible that all members of a certain branch of the CIA knew of a terrorist threat, but due to fear of the reactions, expectations of incredulity etc. failed to communicate their knowledge further upwards. However, the natural way of responding to such a case would be to say that although the CIA subgroup in question did *know*, it failed to *act* on its knowledge due to other deficiencies than purely epistemic ones.¹⁷

We should not be overly behaviorist in our theorizing about group knowledge, as compared with individual knowledge, even if we may resort to behavioral criteria in the cases where no individual knowledge of the target belief is present in the group. A group’s disposition to act vis-à-vis a proposition may be sufficient ground for

¹⁶ Dretske (1969, Chap. 3) and Cassam (2007, pp. 27–50) have championed a notion of *epistemic perception*, arguing that seeing that *b* is *F* is itself a way of knowing that *b* is *F*. Williamson (2001, pp. 33ff.) takes knowledge to be the most general factive mental state, making it encompass states likes seeing or feeling something.

¹⁷ If we expand the unit of assessment to include the higher CIA authorities, we must of course say that the CIA did *not* know. But in that case, the summative condition is no longer met.

attributing knowledge to it, but it is conceivable that the knowledge could be present even in the absence of such dispositions. (Surely there will still be plenty of dispositions grounding the knowledge attribution; only these will be dispositions of the individuals).

I am not sure how widely the summative account can be applied.¹⁸ But the considerations in its favor provides further evidence that collective knowledge comes in many different forms, and that some of them are rather simple and require little in terms of collective metacognition and other possible binding relations.

The wide applicability of the summative account moreover supports my contention that individuals can become members of the same group simply by doing the same job (e.g. knowing the same proposition). Consider also the phrase “We all know that ...”. While some non-epistemic specification of the extension of the “we” in question may be implied (for example, a restriction to a certain television audience, to adult Westerners or the like), it seems that in such cases the individuals referred to have little more in common than knowing the fact in question.

9 Objections and scruples

The more liberal account has definite advantages over the received view. It does come at a price, however, and I consider it an open question how far we should go in the direction I have suggested.

Among the less serious worries is that the liberal account might seem to be really an anti-theory. I have made the most of what is *not* required for group knowledge, and though it is possible to state my account in terms of necessary and sufficient conditions, they have to be formulated in a highly general, abstract and also somewhat vague manner, with little prospect of further substantial specification.

While it must be admitted that my proposal is in some respects less precise than the received view, this worry can easily be met by pointing out that the same sort of indeterminacy pertains to almost any epistemological analysis. Standard analyses of knowledge and justification do not say what it takes for a belief to be justified in concrete cases (i.e. how the abstract justification requirements actually have to, or can, be met). Reliabilists leave it open exactly how reliable a belief-forming process must be (Goldman 1986, p. 51). Counterfactual analyses do not tell us exactly what makes a possible world “sufficiently nearby” to be relevant to our assessments of specific beliefs. And it is taken for granted that infinitely many constellations of specific justificatory or reliability-conducive factors may satisfy the same general epistemic requirements. The appeal to the no double standard-principle should thus be more than sufficient to allay this worry.

More serious worries concern the nature of groups and collectives. Philosophers appear to maintain strong intuitions that groups must be unified by something like

¹⁸ While the summative account obviously fails as an account of group knowledge in the strong, demanding sense associated with the standard view, it does not only fail in such cases. It is even less appropriate for cases of genuinely distributed knowledge. Hence the summative-non-summative distinction does not match the distinction between the liberal and the received view (thanks to an anonymous reviewer for pressing me on this point).

joint intentionality. Some of their intuitions may be accommodated by devising a division of labor between the liberal and the received view and making some further distinctions. First, it might be said that the received view is about collective epistemic *agency*, whereas my account is rather about collective epistemic *behavior*. A group in my liberal sense is presumably not an agent, since it does not act on the basis of deliberation or awareness of its goals and activities.¹⁹ This is quite innocuous, as long as it is not also asserted that agency is a prerequisite for knowledge in general. Secondly, it might be said that the received view is a view about *knowing collectively*, whereas my theory is rather about *knowing collectives*. A group that exhibits reflexivity and joint attention can be said know *qua* group, whereas the looser collectives do indeed manage to know, but the members do not *share* the knowledge.²⁰ Again, it might well be so; but both common usage and practical concerns support the inclusion of mere knowing collectives among the legitimate subjects of group knowledge.

My most serious worry has to do not with the cohesion of the group or its possible transience or instability, but with its demarcation. My analysis makes crucial use of the notion of an *epistemic contribution*. While we might easily marshal an array of paradigm cases of such contributions—possessing evidence and making it available to other members of the group, providing and inferring from testimony, standing poised to check or criticize results or procedures—it is much more difficult to draw the outer line between merely marginal and completely irrelevant contributions. Moreover, with nothing but epistemic connectedness being required for group membership, there is a serious risk that many groups turn out to be extremely large and within any clear boundary.

Part of the problem lies in the infinitely many ways in which actions can be epistemically relevant. What about the people who help keeping other people sharp and efficient? Surely the controller in an organization, who monitors its epistemic output, and makes recommendations for improvements, should be counted among the members of the group. But what, then, about the guy who gets coffee for the office staff? What about the builders, craftsmen or indoor architects that help improve the cognitive niche of a group? What about the musicians whose recorded music helped me relax and concentrate on writing this paper? Must they be given epistemic credit and included into a comprehensive collective subject of knowledge?

Adopting an internalist and/or deontological epistemology, which distinguishes “intrinsically” epistemic factors sharply from everything else (and does not make them dependent on actual truth-conduciveness), would make the demarcation easier, and generally tend to make group membership more exclusive. But it would detract from

¹⁹ However, there is little consensus about the definition of agency or the distinction between agency (or action) and behavior. See e.g. Dretske (1988, pp. 3ff.) and List and Pettit (2011, pp. 19ff.).

²⁰ In a strong sense of “sharing”. Of course they do share it, if this means simply having the same items of knowledge. But they do not know *qua* group. I do not think, however, that this gives us reason to not treat such a group as a genuine subject of knowledge ascriptions (*pace* the suggestions of an anonymous reviewer). Apart from the evidence from linguistic usage, the no double standard-principle also tells against such a move, since we do not require a high degree of reflexive consciousness in cases of individual knowledge attribution. Note, however, that my theory is can accommodate the point that knowing collectively is a distinctive kind of collective knowledge, as it acknowledges the standard view as a fitting description of an important *subspecies* of group knowledge.

the meliorative potential of the theory, since improvement of the external conditions of knowledge production will often be more efficient than enforcing the use of a restricted class of principles and methods.

The problem has affinities to other problems that bedevil reliabilism, like the “generality problem” (Conee and Feldman 1998), and, especially, the problem of fixing the extent of belief-forming processes (Goldman 1979). And some of the typical moves on part of reliabilism are not available in this case, for example a restriction to internal processes (Goldman 1979, pp. 12ff; Alston 2005, pp. 120ff.). For the whole point of my account is that it should be permissive and include interactions between the members of a group among the epistemically relevant factors.²¹

The solution has to be pragmatic. As already stated, there is in principle no end to the range of persons that may be counted among the members of a group. Likewise, there is in principle no limit to the types of contribution that may be considered relevant to the epistemic assessment of a group. But in practice, we will likely focus on factors that are either integral, or at least fairly close, to the paradigm cognitive processes involved in the task solution process, and treat everything else as belonging to the background conditions, assuming these to be fairly normal. (There is a close analogy to our practices of *causal selection*, which divide the extremely complex and wide-ranging set of factors on which an event depends into genuine causes and mere *enabling conditions*, but do so in a highly context- and interest-sensitive manner (Broadbent 2008)).

We can add that groups should normally be assessed in their natural environment, since features of it may be crucial to R. The interaction between the members will often depend on technological means of communication, and the physical layout of an organization might matter as well. “The Russians knew how to build the bomb” may well be taken to contain an implicit assumption that the Russians (in the relevant sense) were also maintaining an appropriate infrastructure. “Softer” factors of the environment, like a certain “climate of trust” (see Hardwig 1991) or specific “evaluation culture”, may also help support or strengthen R and so bring the justificatory (and/or other epistemically relevant) factors to bear on the collective task solution.

Sometimes, however, there may be a point in assessing the epistemic standing of a group independently of its environment or certain environmental features. The ability to form true beliefs or act intelligently under differing circumstances, in the absence of specific physical or social conditions, may be an epistemic asset in itself. The epistemic virtue of *cognitive flexibility* has been defined as the ability to adapt the cognitive processing strategies to face new and unexpected conditions in the environment (Cañas et al. 2003). But it should be understood more broadly, as also including the ability to adapt to whole new environments. Hutchins’ navy vessel crew is bound very tightly to its niche, as it works in a kind of close cognitive symbiosis with the ship and depends crucially on some very specific technology. Medical teams and laboratory staffs might likewise become epistemically inept once they are outside their special-

²¹ Note, however, that my account does not really commit one to accept the extended mind-hypothesis, even though it is inspired by it. One might maintain that only internal processes count as *cognition*, but deny that epistemic states supervene on cognitive processes.

ized, technology-rich environment. Task forces and consultant teams are, by contrast, expected to be able to operate reliably under differing and unforeseen conditions.

The answer to the demarcation problem is, in short, that we just have to live—and should be able to live well—with the fact that there is no principled limits to the range of potentially relevant, membership-conferring relations. Getting coffee or decorating the meeting room may sometimes count as epistemic contributions, though we will mostly be able to ignore such activities and stick to the usual suspects like inference procedures, methods of deliberation, processes of communication and information storage, monitoring etc. (though the range of such “paradigm” examples may increase steadily, as we learn more about the details and mechanisms of knowledge production. Hence the answer to the demarcation problem will also depend on empirical investigations).

A further likely worry concerns the *redundancy* (and apparent dispensability) of group knowledge attributions. The same complex states and actions can be described in both collective and completely individualist epistemological terms. Instead of ascribing knowledge to the collective consisting of Tom and Abby, we might just ascribe *testimony-based* knowledge to Abby and describe Tom as her reliable source of evidence. So the collective knowledge attribution appears superfluous and should be eliminated by Occam’s razor. Even in cases where no single member has full knowledge of the target proposition, we might explain the complex of actions and dispositions of the collective in purely individual terms, with reference to various items of partial knowledge or belief and their causal consequences.

However, as I noted at the outset, reducibility is not unreality. I have no pretense that collective knowledge is in any sense fundamental. All that I insist on is that collective knowledge attributions are getting at real and significant states and processes, and that they are conceptualized at a level that is suitable for theoretical understanding and practical manipulation, inasmuch as they are able to support explanations, predictions and generalizations.²² That complementary descriptions are available is an expected and welcome result. It is, for example, completely obvious in the cases where the summative account applies.

Moreover, even though there are clear examples of redundancy or reducibility, the possibility of a complete reduction to individual states remains an open question. There may be some group knowledge attributions that cannot, event in principle, be substituted by complex descriptions in individualist terms. Bird (2010) presses the case against reductionism by including physical features of the environment in his account

²² Huebner (2014) seem to agree with these requirements, but is more skeptical about the extent to which they are met. Part of this disagreement stems from Huebner’s being concerned with collective *mentality*, whereas I am concerned with collective *knowledge* (and do not think the latter needs any very robust psychological foundation). Part of it stems from Huebner’s making, apparently, stricter requirements for explanatory usefulness than I would like to do. This points to the issue of different levels of reality (i.e. is it still legitimate to speak of “mental” or “collective” reality, even if the entities in question are wholly reducible to “lower-level” entities?), the relationship between fundamental and special sciences etc., which I cannot go into here. For relevant criticism of Huebner, see Tollefsen (2014). To resolve the disagreement conclusively, I would have to provide a detailed analysis of concrete cases. As long as it has not been shown that the requirements are widely met, even my permissive account may not suffice to fend off skepticism. This, however, is also a common predicament of theories of individual and group knowledge. But in both cases, softening the requirements will of course make it more likely that they are actually met.

of social knowledge, arguing that it does not supervene on the mental states of individuals. This is very much in line with my idea that group knowledge does not supervene on J; that that we have to take R into account as well. The reductionist, however, might respond that this does not really take us beyond individualism at the fundamental level. All that matters is, basically, that the individuals in the group have access to the right information, instantiate the right processes etc. (apart from having certain individual epistemic properties). Though such a situation may sometimes, as a matter of empirical fact, only be realized in an environment that contains certain pieces of technology, material inscriptions etc., it can still be described in terms of the states and propensities of individuals. Were the same propensities based not on e.g. the presence of books in library or information processing done by computers, but by individual memory or inference processes, the epistemic standing of the group would be the same. Whereas it is possible that two groups can be alike with respect to their *occurrent* mental states, but only one of them have knowledge, they can hardly be alike with respect to all their *dispositional* mental states without having the same epistemic standing. Features of the environment that are not reflected in the dispositions of individuals cannot matter for knowledge. Hence I will not rule out that everything that is covered by my account of group knowledge could be analyzed in individualist epistemological terms as well, though it would be an extremely laborious and trivial exercise.²³

Finally, what are we to make of the psychology requirement? Exploring the forms and range of collective knowledge obviously puts it under pressure, making it at times look like little more than a sentimental relic.²⁴ It may be argued that knowledge must, intuitively, embody some kind of representation, if not of the target proposition or task, then at least of some of the steps in the acquisition or exercise of knowledge (as is arguably the case with knowing-how, which is assumed to involve some kind of control over one's actions, a sensitivity to changed circumstances (Ryle 1949, p. 47) or even some kind of understanding (Bengson and Moffet 2012)).

It is tempting to make some kind of “system requirement”—that is, accept that mental states may not be involved in all (or even most) states of knowledge, but still insist that the system to which such knowledge is ascribed must be able to have such states (this rules out libraries and computer networks as genuine subjects of knowledge). This way of thinking is analogous to denying that mental states must themselves be conscious, but maintaining that they can only be ascribed to creatures that are able to have conscious states (e.g. Searle 1992, pp. 155ff.; Burge 1997). The trouble with such system requirements is that it is hard to see what difference the presence of the allegedly crucial extra factor does. The epistemic (or, in the analogous case, the mental) job would be done just as well in its absence. Perhaps the intuition rests on the idea that a genuine subject of knowledge must be able to sometimes “awaken” and perform some kind of metacognitive assessment or acquire some kind of awareness of what is going on. But tempting though this idea might seem, it is

²³ Simon (1982, p. 43) likewise notes that by including social relations and dispositions among the explanatory factors, methodological individualists are able to make their theories compatible with non-reductionist views of social phenomena.

²⁴ Note, however, that if one opts for an internalist version of the theory, J will necessarily comprise mental states, and so the psychology requirement will be fulfilled.

hard to find a rational justification for it. As noted earlier, we do not generally take first-order knowledge to depend on higher-order processes; and even if such processes were necessary, they could probably be implemented in such a way that we would not necessarily categorize them as mental.

Of course, one might think that the dispositions underwriting knowledge attributions must themselves provide sufficient basis for the ascription of mental states. But although I understand that unconscious mental states do come cheap (this is not the place to argue for any necessary connection between consciousness and mental states), I would like to leave open the possibility that there could be more to genuine mentality than the mere presence of such dispositions. One might choose a path of more direct stipulation. Williamson, for instance, considers knowledge a paradigm mental state on par with belief (2001, p. 6). Bird tends to follow him, but cautiously adds that one may read “mental” as shorthand for “mental or epistemic” (2010, p. 5), effectively leaving the question unanswered.

The issue is complex and difficult—and, I suspect, partly a matter of terminology and taste. But fortunately it is also peripheral, as my account makes no crucial reference to mental states as such, even though there are plenty of such states involved in the processes and situations it is aimed at.

10 Conclusion

I have argued for a liberal view of collective knowledge, according to which it requires merely (1) a true belief (or adequate task solution or other practical achievement) and (2) the presence of a sufficient amount of sufficiently related positive (e.g. justificatory or truth-conducive) epistemic states, processes and capacities among the members of a group. I have departed from the received view by not requiring anything in terms of reflectivity (or other specific conditions on group membership) or joint attention to a target belief or common procedure.

To end on a conciliatory note, let me emphasize that there is still much to be said in favor of the received view, even from the point of view of practical social epistemology. It has a considerable meliorative potential. Though groups need not exhibit reflexivity or joint attention, it will often be a useful ideal to which they should aspire, or which they should be helped or forced to attain or at least approach to a certain degree. Making epistemic contributions recognizable and transparent might help to increase epistemic performance.

Nevertheless, the received view remains severely limited. It is not just that it is descriptively inadequate. It is also questionable how far it should count as a practical ideal. The cost of implementing reflectivity, joint attention, and, more generally, higher-order cognitive processes and processes of explication in social settings may be too high.²⁵ It can even be counterproductive, by slowing down processes or creating unnecessary worries and uncertainties. This is a collective analogue to the well-known

²⁵ Though of course it may be argued that a high price can be worth paying for other than epistemic—e.g. moral or political—reasons.

cases where the exercise of a skill is impaired by attention, reflection or verbalization (see e.g. [Beilock et al. 2002](#)).

As is often said, but still seldom done in social epistemology, the theoretical work should be followed up by applied and empirical studies. The distribution of cognitive processes and the interaction between individuals obviously matters; but we need to know more about the vices and virtues of specific patterns of distribution and interaction, the benefits and costs of more or less centralized (or automatic or explicit) procedures, and the right balance between “nice” and “retaliatory” behavior²⁶ for the attainment of epistemic goals. Existing empirical research contains a welter of valuable findings, but is marked by a lack of clear focus on the epistemic dimension. Social epistemology should inform and be informed by future social studies.

Acknowledgments Thanks to Peter Graham, Eric Schwitzgebel, Esben Nedenskov Petersen, Kay Mathiesen, Don Fallis, Bryce Huebner audiences at the University of Arizona and University of California, Riverside, and two thorough and unusually constructive anonymous reviewers for helpful suggestions and comments.

Work on this paper was supported by a grant from the European Union Interreg 4a Project NexusPro2017.

References

- Alston, W. P. (2005). *Beyond “Justification”*. Ithaca: Cornell.
- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
- Beilock, S., Carr, T., MacMahon, C., & Starkes, J. (2002). When paying attention becomes counterproductive: Impact of divided versus skill-focused attention on novice and experienced performance of sensorimotor skills. *Journal of Experimental Psychology: Applied*, 8(1), 6–16.
- Bengson, J., & Moffet, M. A. (2012). Nonpropositional intellectualism. In J. Bengson & M. A. Moffet (Eds.), *Knowing how. Essays on knowledge, mind and action* (pp. 161–195). Oxford: Oxford University Press.
- Bird, A. (2010). Social knowing: The social sense of ‘scientific knowledge’. *Philosophical Perspectives*, 24, 23–56.
- Bonjour, L. (2002). *Epistemology*. Lanham: Rowman & Littlefield.
- BonJour, L., & Sosa, E. (2003). *Epistemic justification*. Oxford: Blackwell.
- Broadbent, A. (2008). The difference between cause and condition. *Proceedings of the Aristotelian Society*, 108, 355–364.
- Burge, T. (1997). Two kinds of consciousness. In N. Block, O. Flanagan, & G. Güzeldere (Eds.), *The nature of consciousness* (pp. 427–433). Cambridge, MA: MIT.
- Cañas, J. J., Quesada, J. F., Antolí, A., & Fajardo, I. (2003). Cognitive flexibility and adaptability to environmental changes in dynamic complex problem-solving tasks. *Ergonomics*, 46, 482–501.
- Cassam, Q. (2007). *The possibility of knowledge*. Oxford: Clarendon Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 10–23.
- Coady, D. (2012). *What to believe now? Applying epistemology to contemporary issues*. Oxford: Wiley-Blackwell.
- Condorcet, J. A. N. d. C. (1785). *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: Imprimerie royale.
- Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, 89, 1–29.
- Conee, E., & Feldman, R. (2004) [2001]. *Internalism defended* (pp. 53–82). (Reprinted from *Evidentialism*, Oxford: Oxford University Press.)
- Corlett, A. J. (1996). *Analyzing social knowledge*. Lanham, MD: Rowman & Littlefield.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Clarendon Press.

²⁶ Famously emphasized by [Axelrod \(1984\)](#) as opposite tendencies that are both required for successful cooperation.

- Donaldson, I. (1995). *American anti-management theories of organization: A critique of paradigm proliferation*. Cambridge: Cambridge University Press.
- Dretske, F. (1969). *Seeing and knowing*. Chicago, IL: Chicago University Press.
- Dretske, F. (1988). *Explaining behavior*. Cambridge, MA: MIT.
- Feldman, R. (2004) [1988]. Having evidence. (Reprinted from *Evidentialism*, pp. 219–241, by E. Conee & R. Feldman, Eds., Oxford: Oxford University Press.)
- Foley, R. (2012). *When is true belief knowledge?*. Princeton: Princeton University Press.
- Gilbert, M. (1989). *On social facts*. London: Routledge.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel. (Reprinted from *Liaisons: Philosophy meets the cognitive and social sciences*, by A. Goldman, Ed., 1992, Cambridge, MA: MIT.)
- Goldman, A. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (2004). Group knowledge versus group rationality: Two approaches to social epistemology. *Episteme*, 1(1), 11–22.
- Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy*, 82(7), 335–349.
- Hardwig, J. (1991). The role of trust in knowledge. *Journal of Philosophy*, 88(12), 693–708.
- Hetherington, S. (2011). *How to know: A practicalist conception of knowledge*. Oxford: Wiley-Blackwell.
- Huebner, M. (2014). *Macro-cognition: A theory of distributed minds and collective intentionality*. New York: Oxford University Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT.
- Kieser, A., & Walgenbach, P. (2010). *Organisation*. Stuttgart: Schäffer-Poesel.
- Klausen, S. H. (2010). Kollektiv viden og læring—myte og realitet [Collective knowledge and learning—myth and reality]. In M. Paulsen, S. H. Klausen et al. (Eds.), *Filosofiske perspektiver på kollektiv læring* (pp. 15–39). Aalborg: Aalborg University Press.
- Kornblith, H. (2012). *On reflection*. Oxford: Oxford University Press.
- Kukla, R. (2012). “Author TBD”: Radical collaboration in contemporary biomedical research. *Philosophy of Science*, 79(5), 845–858.
- Kuklick, H. (1983). The sociology of knowledge. Retrospect and prospect. *Annual Review of Sociology*, 9, 287–310.
- Lackey, J. (2012). Group knowledge attributions. In J. Brown & M. Gerken (Eds.), *New essays on knowledge ascriptions* (pp. 243–269). Oxford: Oxford University Press.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.
- List, C. (2005). Group knowledge and group rationality: A judgment aggregation perspective. *Episteme*, 2(1), 25–38.
- List, C. & Pettit, P. (2011). *Group Agency*. Oxford: Oxford University Press.
- Mathiesen, K. (2006). The epistemic features of group beliefs. *Episteme*, 2, 161–175.
- Mokyr, J. (2002). *The Gifts of Athena*. Princeton: Princeton University Press.
- Pavitt, C. (2003a). Colloquy: Do interacting groups perform better than aggregates of individuals? *Human Communication Research*, 29(4), 592–599.
- Pavitt, C. (2003b). Why we still have to be reductionists about group memory. *Human Communication Research*, 29(4), 624–629.
- Phelan, M., Arico, A., & Nichols, S. (2012). Thinking things and feeling things: On an (alleged) discontinuity in folk metaphysics of mind. *Phenomenology and the Cognitive Sciences*, 12(4), 703–725.
- Quinton, A. (1975). Social objects. *Proceedings of the Aristotelian Society*, 75, 1–27.
- Rupert, R. (2005). Minding one’s cognitive systems: When does a group of minds constitute a single cognitive unit? *Episteme*, 1, 177–188.
- Rupert, R. (2011). Empirical arguments for group minds: A critical appraisal. *Philosophy Compass*, 6(9), 630–639.
- Rupert, R. (forthcoming). Individual minds as groups, group minds as individuals. In B. Kaldis (Ed.), *Mind and society: Cognitive science meets the philosophy of the social sciences*. Synthese Library Special Volume. Berlin: Springer.
- Ryle, G. (1949). *The concept of mind*. Chicago: University of Chicago Press.
- Schmitt, F. F. (1994). The justification of group beliefs. In F. F. Schmitt (Ed.), *Socializing epistemology* (pp. 257–287). Lanham, MA: Rowman & Littlefield.
- Scott, W. R. (1992). *Organizations: Rational, natural, and open systems* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Searle, J. R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT.

- Simon, M. A. (1982). *Understanding human action: Social explanation and the vision of social science*. Albany: SUNY Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Clarendon Press.
- Stanley, J. (2011). *Know how*. Oxford: Oxford University Press.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98, 411–444.
- Surowiecki, J. (2004). *The wisdom of crowds*. New York: Doubleday.
- Tollefsen, D. P. (2002). Challenging epistemic individualism. *Protosociology*, 16, 86–117.
- Tollefsen, D. P. (2007). Group testimony. *Social Epistemology*, 21(3), 299–311.
- Tollefsen, D. P. (2014). Review of Huebner, M.: *Macro-cognition*. *Notre Dame Philosophical Reviews*. Accessed April 09, 2014, from <http://ndpr.nd.edu/news/47449-macro-cognition-a-theory-of-distributed-minds-and-collective-intentionality/>.
- Tuomela, R. (2004). Group knowledge analyzed. *Episteme*, 1(2), 109–127.
- Tuomela, R. (2007). *The philosophy of sociality. The shared point of view*. Oxford: Oxford University Press.
- Williamson, T. (2001). *Knowledge and its limits*. Oxford: Oxford University Press.