



ARTICLE

Nonreductive Group Knowledge Revisited

Jesper Kallestrup 

University of Aberdeen, Scotland, UK
Email: jesper.kallestrup@abdn.ac.uk

(Received 25 June 2021; revised 3 May 2022; accepted 12 May 2022)

Abstract

A prominent question in social epistemology concerns the epistemic profile of groups. While inflationists and deflationists agree that groups are fit to constitute knowers, they disagree about whether group knowledge is reducible to knowledge of their individual members. This paper develops and defends a weak inflationist view according to which some, but not all, group knowledge is over and above any knowledge of their members. This view sits between the deflationist view that all group knowledge is reducible to individual knowledge, and the strong inflationist view that some such knowledge even fails to supervene on features of individuals. Thus, some group knowledge is irreducible, but all such knowledge is anchored in, and so doesn't float freely from, individual features.

Keywords: Socially structured groups; group knowledge; inflationism; epistemic standards; distributed cognition

1. Introductory remarks

As a branch of social epistemology, collective epistemology is concerned with the question, among others, of whether groups can have knowledge over and above any corresponding knowledge of their individual members. *Deflationists* reject such possibility, by insisting that all knowledge of groups is reducible proposition-by-proposition to knowledge of their members. *Inflationists*, however, deny that such reduction is always possible; rather, they maintain roughly on the basis of two different kinds of models that some group knowledge is irreducible. Some inflationists focus on how certain groups conjoin epistemically relevant attitudes of their members, especially in the context of variable epistemic standards, to form irreducible group knowledge, whereas others find inspiration in the way such knowledge is generated by epistemic collaboration between members for whom the cognitive labour of the group is distributed.¹ However, neither view has typically been approached through the lens of the metaphysics of groups. That, we contend, is a mistake because such a metaphysical perspective

¹Bird (2014: 43–7) dubs these the commitment model and the distributed model. Advocates of the former include Gilbert (1989, 2004, 2010), Tuomela (1992, 2004, 2011) and Schmitt (1994). The latter is defended by Hutchins (1995), Tollefsen (2009), Bird (2010, 2014), De Ridder (2014) and Palermos (2020). We return to both models in section 4.

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

sheds new light on the debate between them. While some regard groups as non-singular pluralities of their members, others take them, correctly on our view, to be realizations by individuals of social structures. In fact, irrespective of how the question of reduction of group knowledge is answered, all groups should be understood in terms of such structured wholes. By reference to cases of both kinds, and with these metaphysical underpinnings of the disagreement over the nature of group knowledge in mind, this paper argues for a *weak inflationist* view, on which some, but not all, group knowledge is proposition-wise irreducible to knowledge of their members. We disagree both with the *deflationist* claim that all group knowledge is reducible, and with the *strong inflationist* claim that some group knowledge doesn't even supervene on individual features. On our view, no group knowledge floats freely from, but is rather anchored in, individual states, processes and relations. We shall proceed as follows. Section 2 presents a conception of groups as one rather than as many, drawing on, inter alia, the work of Ritchie (2013, 2015, 2020) and Schmitt (1994) in social metaphysics. Against said metaphysical conception of groups, section 3 develops a deflationist view of group knowledge, as distinct from aggregate knowledge which applies to mere collectives. Our aim is not to ultimately defend deflationism, but to work into shape a version of this view that avoids some fairly minor worries, e.g. about knowledge ascriptions to distinct yet coincident groups. Section 4 then argues from so-called *epistemic divergence* cases illustrating both of the above-mentioned models that, despite Lackey's (2012, 2020) recent objections, some group knowledge is irreducible and yet supervenient on individual features. As we argue against not just this amended version of deflationism, but also Bird's (2010, 2014) strong inflationist notion of "social knowledge", the upshot is a weak version of inflationism. Finally, section 5 contains some concluding remarks.

2. The metaphysics of groups

Let's begin by assuming the falsity of *eliminativism* about groups. Take mereological nihilists who deny the existence of composite objects. If nothing has proper parts, there are no humans, and so no groups either; only mereological simples exist. But arguments from efficacy, indispensability and common sense provide strong reason to believe that groups exist. The fact that we speak ubiquitously and unproblematically about groups is a reason to take their existence at face value; in particular, common-sense talk about their causal effects presuppose existence. Moreover, our ontology ought to include groups as they are arguably indispensable to our best social science theories.² So, we shall assume *realism* about groups, where we have in mind *social* groups, comprising individual human beings.

But, importantly, there are different realist conceptions of groups. We maintain that a *collective*, as a mere collection of individuals, isn't a group. Let 'X' name the 12 Swedish students in my department. Suppose they are the only EU students whose parents have the highest socioeconomic status; call those same students 'Y' under that description. Neither X nor Y forms a group, but the seven students whom we imagine happen to make up both the Library Committee and the Squash Club clearly do. Why is that? Ritchie (2013, 2015) proposes a list of important (yet non-exhaustive) constraints that groups seem to satisfy:

²The following owes much to Ritchie (2013, 2015, 2020), but see also Koslicki (2008), Haslanger (2016) and Fine (2020).

- (1) Groups may have different members at different times and across different worlds.
- (2) Groups exist at one time or at one world without existing at every time or at every world.
- (3) Non-identical groups may have extensionally coincident memberships.

While X and Y don't exist at every time and world, because those 12 students don't exist at every time and world, they fail to qualify as groups, because neither allows for the possibility of having different members at different times or at different worlds. Another reason to reject their status as groups is that they are (numerically) identical in virtue of being necessarily extensionally equivalent; after all, 'X' and 'Y' are just (rigid) names of the same 12 students. To reject (1)–(3) by insisting that X and Y are nevertheless groups is to view *groups as non-singular pluralities*: they are nothing over and above their individual members.³ True, X and Y both have *features* in virtue of which they are their members taken together, but not joined together as they furnish *no internal structure*.⁴ X and Y aren't arranged differently as a result. In fact, these students don't engage in any collaborative deliberation or enterprise. They think and act as many rather than as one. X and Y are better conceived of as *feature collectives*, or as *social kinds*, where membership requires possession of shared (clusters of) social features. Feature collectives are taxonomical classifications into which individuals may be collected on the basis of such characteristics, broadly understood.

In contrast, both the Library Committee and the Squash Club constitute groups: both have different members at different times and worlds, neither exists at every time or world, and they are (numerically) distinct despite their actual co-memberships. Satisfying (1)–(3) is, however, insufficient for group-hood. Take the British electorate as the collection of entries on the UK electoral register, and suppose that presently all UK residents over 18 are registered to vote. In that case, the British electorate and adult residents in the UK are distinct yet co-extensional. Both have different members at different times and worlds, and neither exists at every time and every world. Still, they only constitute feature collectives, because they aren't suitably structured, nor do they share objectives or means to achieve them, as witnessed by their voting behaviour. Exactly what else is needed to convert such a collective into a group is a vexed question. We shall focus on *organisational structure*, *joint intentions* and the *office and charter* of groups as what binds individuals together to form a group.⁵

First off, the requisite *cohesion* requires that groups be *structured wholes*, i.e. wholes whose parts are interdependent in that their function is constrained by their relation to other parts. More precisely, following Ritchie (2013, 2015), groups are *realizations by individuals of social structures*, in which group members occupy nodes (or play roles) connected by internal edges (or links). A social structure is a social network of

³This view also goes under the name of ontological individualism, in contrast with both group eliminativism and the ontological collectivist view which we outline below.

⁴As Gilbert (1989: 9, 2004: 96) noted, having a property in common doesn't automatically make an otherwise arbitrary collection of people a group, e.g. haemophiliacs have an inherited genetic disorder in common but constitute no "established group".

⁵Hawley (2017) argues that groups are particular, concrete wholes composed of individuals as their mereological parts, Effingham (2010) claims that groups are sets of ordered pairs of worlds, times and people, and Uzquiano (2004) holds that groups are *sui generis* entities constituted by sets. While these views can arguably accommodate (1)–(3), they offer no account of the distinction between structured groups and feature collectives in terms of the parthood relation or set membership, respectively.

functional relations, and because the edges capture such relations as holding between the nodes, they represent the functional organization of the group.⁶ Obviously, external features may also factor into social structures. So, a group exists just in case it has a social structure of nodes and edges, *and* members who occupy those nodes by bearing the functional relations that are specified by the structure. Unlike the British electorate, our Squash Club has an organizational structure, in which its members occupy specific nodes, which are defined by the functional relations that occupiers must bear to other node occupiers, as captured by the edges. Thus, the treasurer prepares the budget, which is then signed off by the chairperson before instructing the secretary to spend the funds. As regards the requirements that the nodes specify what it takes to occupy them, Shapiro (1997: 82–3) differentiates between a *places-are-offices* perspective, where the individuals that occupy the nodes (or “fill the places”) of the structures are considered, e.g. the current chairperson is more efficient than the previous one, and a *places-are-objects* perspective, where the nodes of the structures are treated as objects, e.g. the chairperson chairs the AGM. From the latter perspective, certain individual actions are required for node occupancy, which can then be evaluated as performances from the former perspective; but groups must also *act jointly* in that their members team up to do things in concert, e.g. the members of the Squash Club won the University League together.

Secondly, as in the case of the Squash Club, the social structures are overtly created and intentionally implemented, but they may in other cases be unintentional and covert. What is imperative is that all groups involve individual as well as *collective intentionality*: members of groups typically need to cooperate in shared plans and actions, and the roles demanded by the organization of the group will require functional integration. They need to act in ways defined by the functional roles they play. That requires an individual intention to participate, if not to fully commit, as a singular agent by actively occupying one’s node, *and* a shared intention underpinning the joint action of the group where each member intends to act jointly with others. To wit, each must intend (i) that they all enact the joint performance, and (ii) to do their bit because of their belief that others intend to do their bit. There are at least three familiar ways of cashing out joint intention. Searle (1990, 1995) proposes that joint actions be based on intentions held by individuals, but with a *mode* that differs from how *I*-intentions are held. On his view, such *we*-intentions neither transcend individuals’ minds, nor are they reducible to *I*-intentions plus beliefs.⁷ Bratman (1993, 1999) places constraints on “shared intentions” to do with their *common contents* and interrelations, especially that ‘we’ shows up in the content of *I*-intentions. Gilbert (1989, 1994, 1996) suggests that joint actions involve the intentions of a *plural subject* which is brought into existence by the individuals who engage in the joint action; basically, a plural subject is constituted by those who jointly commit to some belief, as a body. We shall not here adjudicate between these views, but merely note that all we need for our purposes is Searle’s more modest account.

Thirdly, the foregoing lends itself to a conception of groups as *chartered* in Schmitt’s (1994: 272–3) sense. When a group has a *charter*, as constituted by the intentions of its (founding) members, its actions aim to fulfil its *office* in accordance with its charter.

⁶See also Haslanger (2016) and Fine (2020). Ritchie (2020: 412) allows for groups to persist through changes not only in which individuals occupy which nodes, but also, at least to some extent, in their structure, i.e. in which functional relations hold between the nodes.

⁷For more on the *we*-mode approach see Tuomela (2007).

The office of a group can roughly be understood as its unified task, goal or purpose, whereas its charter comprises the rules, norms or standards by which the group is governed,⁸ such that a chartered group wouldn't exist without either. Both charter and office are present in the case of our Squash Club, but absent for the British electorate. Particular chartered groups are characterized in terms of the specific, individual and collective, intentions of their members, as well as the actions that aim to fulfil their office as fixed by their charter. For these intentions, together with accompanying actions, determine the social structure of a group, comprising its edges and nodes. They are the glue that sticks individuals together to forge a group, without which it would crumble. A group's charter and office are sometimes formally enshrined in a system of laws, rules or regulations; other times they are evidenced by its practice. Either way, a collective is a group only if its (founding) members jointly set up common aims or objectives and agree on how to proceed in order to meet them. An adequate organizational structure in terms of nodes and edges within the group is therefore needed to facilitate the means that are carried out for the purpose of achieving its end. Only when such structure is in place can a chartered group function in (by trying to fulfil) its office.⁹

Let's summarize. When viewed as realizations by individuals of social structures, groups satisfy (1)–(3): they can vary in members across times and worlds, they don't exist at every time and world, and two distinct groups may be extensionally coincident if they differ in structure or vary in membership at other times or worlds. Moreover, for groups to sustain social structures requires collective intentionality on the part of their members. Such structure, intentionality and accompanying action evince the charters and offices of groups. We shall henceforth call these 'structured groups'.¹⁰ On the *ontologically individualist* view of groups as non-singular pluralities, a group is nothing but the individuals who belong to it, and so any group talk is ontologically non-committal. On our *ontologically collectivist* view, a structured group is irreducible to its members, but it constitutively depends on them in conjunction with their functional relations in the social structure in question.

3. Deflationist group knowledge

Equipped with the foregoing, somewhat rudimentary, conception of what structured groups amount to, let's turn to group knowledge as instantiated by such groups. This section argues that the most promising deflationist account should subsume elements of social structures, but let's first dwell on a different conception of knowledge of collectives, which has wrongly been taken by deflationists to capture group knowledge.

Let *aggregate knowledge* be a summation of knowledge by individual members of a collective, attributed by some aggregation rule, or statistical criterion. If the rule is

⁸See also Tuomela's (2011: 78) notion of a group's "ethos".

⁹Schmitt (1994) assumes that not all groups have offices since they may not be chartered, indeed such groups may still display some structure. Cases that spring to mind are people in romantic relationships, or close family members, who develop transactive memory systems by pooling their cognitive resources, or forms of rhythmic interpersonal coordination which emerge spontaneously. While the groups we consider could in principle involve either, we shall restrict attention to groups with charters, broadly construed.

¹⁰Fine (2020: 87ff) suggests the difference between structured groups and feature collectives isn't a division between kinds of groups, but rather a continuum of more or less complex groups. However, as he admits (2020: 87ff), while the structure of the former concerns internal relations between their parts, the structure of the latter is purely attributive, which in turn gives rise to importantly distinct types of attitudes, as we shall see.

majority (or supermajority), the collective has the pertinent knowledge, because *most* of its members do, if the rule is unanimity, the collective has the knowledge, because *all* its members do,¹¹ and if the rule is dictatorial, the collective has the knowledge, because a *unique* member does. We call these *majority*, *shared* and *dictatorial* knowledge, respectively, but they are all special instances of aggregate knowledge. Consider:

(AGGREGATE KNOWLEDGE) Necessarily, a collective *C* has aggregate knowledge that *p* if and only if a summary by some aggregation rule of individual members of *C* knows that *p*.

Reflect that (AGGREGATE KNOWLEDGE) involves ascription of knowledge to a collective that is reducible proposition-by-proposition to knowledge of its members. Attributing such knowledge presupposes nothing about social structure or collective intentions. It isn't action-guiding and involves no commitment to joint attitudes or endeavours. Such knowledge is therefore best classified as a property only of *feature collectives*, which have nothing in common except for demographic attributes. Still, as List (2014) notes, aggregate attitudes more generally play an important theoretical role in political science, e.g. talk about public opinion on some issue is an aggregate on elicited individual opinions, typically attributed to populations on the basis of opinion polls. As a summary of individual attitudes, aggregate attitudes are a *construct* made by an observer for some descriptive purpose that plays no direct social role within the collective to which they are ascribed.

In contrast, let *group knowledge* be a distinctive property of structured groups, to the exclusion of feature collectives.¹² If the individuals who compose a structured group are also taxonomized as a feature collective, then the latter is capable of (AGGREGATE KNOWLEDGE), but not group knowledge. Of course, one can in effect collapse the distinction by defining 'group' along ontological individualist lines without any reference to social structure, shared intentionality, common purpose or joint action. Deflationists have traditionally proceeded by doing exactly that. Thus Quinton (1975: 2) takes groups to exist as "social wholes", *logically constructed* out of individuals as their parts; hence, they are as real and concrete as their members but carry no further ontological implications. On his view, groups are defined by subject-terms of laws of social science: classes and communities in sociology, governments and parties in politics, and firms and industries in economics. That means "social wholes" illicitly lumps structured groups and feature collectives together. Still, constraints (1)–(3) are met: social wholes may change their individual parts gradually over time, and two distinct wholes may simultaneously have the same parts so long as they are held together by different aggregating relations. Quinton (1975: 22) offers the example of a village with identically membered

¹¹Also called "mutual knowledge" by Vanderschraaf and Sillari (2005), but different from *common knowledge*, which consists in everyone knowing that *p*, everyone knowing that everyone knows that *p*, everyone knowing that everyone knows that everyone knows that *p*, etc. Thus, in emperor's new clothes cases, there is mutual but not common knowledge. For simplicity we take common knowledge to include no higher than second-order knowledge.

¹²We assume that Kitcher (1994), and Giere and Moffatt (2003), are wrong to claim that group knowledge ascriptions are systematically inappropriate. An important observation concerns *ordinary parlance*. A quick Google search will confirm that we speak unproblematically about groups having knowledge, and the ubiquity and diversity of such talk is a reason to take its content at face value. See Kallestrup (2019) for more details.

parish council and church bell-ringers, where there's no incompatibility between Smith's subordination to Jones as regards parish administration, and Jones' subordination to Smith as regards bell-ringing. When described from Shapiro's (1997) *places-are-objects* perspective, Smith and Jones are simply filling two different social places. As these wholes are arranged differently (council-wise and church-wise), they are distinct. But Quinton's other examples, such as social classes or industries, involve no suitable arrangement of individuals resembling a social structure, let alone collective intentionality. Glossing over a key metaphysical distinction, Quinton winds up with a *deflationist* account of group attitudes as aggregate attitudes:

To ascribe mental predicates to a group is always an indirect way of ascribing such predicates to its members. With such mental states as beliefs and attitudes the ascriptions are of what I have called a summative kind. To say that the industrial working class is determined to resist anti-trade-union laws is to say that all or most industrial workers are so minded. (Quinton 1975: 17)

The idea is for group attitudes to be attributed to structured groups and feature collectives alike, in that any such ascription is reducible to attitudes had by most, or all, of their members, regardless of organizational structure, collective intentionality, offices or charters. Quinton's view, as it pertains to knowledge, can be articulated as follows:

(DEFLATIONISM) Necessarily, a group G knows that p if and only if all or most of the individual members of G know that p .

Now, (DEFLATIONISM) faces some smaller concerns about detail. First off, because it involves an aggregation rule of either majority or unanimity, the view rules out any instances of group knowledge where only a minority of members have the pertinent knowledge. But, say, a government may possess knowledge in virtue of a few cabinet ministers being in the know, while lacking other knowledge had by a slight majority of rank-and-file civil servants. Such cases recommend the requirement that the relevant knowledge be had by a (i) *significant* number of (ii) *operative* members.¹³ Following Tuomela (1992, 2004, 2013), an operative member is someone who has decision-making authority, and so is responsible for the content of the relevant state; and 'significant number' is deliberately vague to allow for varieties between different groups. In some cases, a single member with executive power may suffice; in others, multiple such members may be needed. Secondly, the schematic formulation "all or most" suggests a list of *named* members, in which case knowledge of a group composed of those individuals couldn't survive replacement of a majority even if the knowledge was shared by the replacing and replaced members. An easy fix is to *existentially quantify* over individuals who are group members, so as to accommodate the desiderata that groups typically neither exist, nor have the same members, at all times or worlds. For a group trivially lacks knowledge if it fails to exist through not having any members. With those points in mind, consider this amended formulation:

¹³Deflationism is also a view about doxastic justification. If such justification is non-factive, and group justification can result from justification had by less than half of its operative members, then it's possible for a group to justifiably believe p and not- p simultaneously. One way to resolve such inconsistency is to impose an ordering or weighing of importance among members vis-à-vis aggregating group justification.

(DEFLATIONISM*) Necessarily, group G knows that p if and only if there exist a significant number of operative members of G who know that p .

But there is more trouble, for (DEFLATIONISM*) predicts that a group can have knowledge even if none of the operative members know that others have that knowledge. In what Gilbert (1989: 257–67) called a *secrecy situation*, where each member believes a proposition but is afraid to tell others out of fear of being mocked, and so have no reason to believe that anyone else believes it, there is an intuitive pull to deny group belief, and hence group knowledge. In response Gilbert suggested on behalf of the deflationist that *common knowledge* of the target proposition be a requirement. To buttress this claim, consider group agency by the lights of (DEFLATIONISM*).¹⁴ As is familiar, there's a normative connection between knowledge and action, as expressed by the independently plausible *Knowledge-Action Principle*¹⁵:

(KAP) Agent S knows that p if and only if S is epistemically rational to act as if p .

To be epistemically rational in acting as if p is to treat p as a reason for action, or to use p as a premise in practical reasoning. That holds irrespective of whether S is an individual or a group, but for a group to act requires that one or more of its members perform a causally contributing act. Consider what Lackey (2020) calls the *Group/Member Action Principle*:

(GMAP) For every group, G , and act, a , G performs a only if at least one member of G performs some act or other that causally contributes to a .

The problem for (DEFLATIONISM*) arises when (KAP) and (GMAP) are applied in secrecy situations. Let a be the group-level action that is epistemically rational for group G if G knows that p . Following (DEFLATIONISM*), assume G knows that p in virtue of a significant number of its operative members knowing that p . (KAP) then has it that G is epistemically rational to perform a , but following (GMAP), G can perform a only if some such members perform causally contributing acts. While each of those members are assumed to know that p , they aren't able to perform a individually; only G can do that. But if, as in a secrecy situation, those members don't know whether any other members know that p , they don't know whether G knows that p . That makes it epistemically irrational for them to perform acts that contribute towards a , because they lack knowledge of whether G is epistemically rational in performing a , and it isn't epistemically rational for members to contribute towards a if, for all they know, a is epistemically irrational for G . But (DEFLATIONISM*) would have it that questions of group-level epistemic rationality be a matter of the rationality of enough operative members. It follows that G isn't epistemically rational to perform a , and hence that G lacks knowledge that p . Moreover, (KAP) isn't just about the availability of a reason for action; knowledge is supposed to constitute the kind of appreciation that would

¹⁴Thanks to an anonymous referee for drawing attention to the need for common knowledge on the deflationist view so as to accommodate group agency.

¹⁵Williamson (2005: 231; cf. 2017) and Hawthorne and Stanley (2008: 578) all defend a bi-conditional knowledge norm of practical reasoning, in contrast with Fantl and McGrath (2002, 2009) and Lackey (2020: 118) who merely hold the left-to-right reading of (KAP), but their worries about the other direction are orthogonal to the current topic.

convert the existence of a reason into a *personal* reason for action such that the reason also *motivates* the agent to act.¹⁶ The worry is then that nothing would motivate individual members to causally contribute towards *a* when everyone is deliberately keeping their knowledge secret from others. After all, nobody knows whether anyone else has a motivating reason to do their bit. There is thus no epistemic basis for *G* to be moved to act reasonably, rendering *G* causally inert vis-à-vis *a*. The most obvious way to resolve this tension is to demand that *p* be *commonly known*. Consider:

(DEFLATIONISM**) Necessarily, group *G* knows that *p* if and only if there exist a significant number of operative members of *G* who know both that *p* and that the other operative members know that *p*.

However, (KAP) still causes trouble, for (DEFLATIONISM**) predicts that two distinct groups share any knowledge if composed of the same operative members. This squares badly with the fact that groups with identical memberships may have different *collective* methods of deliberation, decision-making procedures, and belief-forming processes.¹⁷ Let *p** be the proposition that the University Library is short on textbooks, which our Library Committee can safely be assumed to know. Given (KAP), the Library Committee is epistemically rational to take action accordingly, e.g. to alert the management of the need to purchase additional copies. But since the Library Committee and the Squash Club are composed of the same seven students, (DEFLATIONISM**) and (KAP) together imply that it's equally epistemically appropriate for the latter to use *p** in practical rationality. But how could *p** form the basis for the Squash Club to act when *p** is *completely irrelevant* to what that group is all about? Nothing the Squash Club does would be rationalized by this knowledge: it wouldn't treat *p** as a reason for action or let *p** feature as a premise in its practical reasoning.¹⁸ For its office and charter place constraints on what course of action would be thus rationalized.¹⁹ This speaks against *p** being known by the Squash Club. Note also that (KAP) is about permission to act, and so is neutral on what the agent *ought* to do, epistemically speaking. There is no claim of being obliged by epistemic rationality to act on everything known. However, our point isn't that the Library Committee is under some obligation to act which the Squash Club isn't, but rather that there are no proper circumstances in which it's epistemically rational for the Squash Club to take the envisaged action. Moreover, bearing in mind that (KAP) is also about motivating reasons for action, the Squash Club would hardly be moved by such knowledge, which suggests it possesses no such reason.

In response to a related worry that the deflationist view has the implication that group belief, when fixed by any random beliefs of its members, "may be completely disconnected from the focus and objectives" of the group, Lackey (2020: 50–1) suggests such attributions, while strictly true, are irrelevant (unimportant or uninteresting). She

¹⁶Hawthorne and Stanley (2008).

¹⁷See also Schmitt (1994: 261), Corlett (1996: 88) and Gilbert (2004: 97–8). Gilbert (1989: 262–73) argues that such cases aren't less problematic for the deflationist if common knowledge of the target proposition is assumed among the same members of the two groups in question.

¹⁸On Schmitt's view (1994: 265), the Squash Club has a reason *q* to believe *p** just in case all its members would openly express a willingness to accept *q* as the group's reason for belief, but it's hard to fathom why they would do that when seeking to fulfil its office.

¹⁹We can ascribe knowledge of utterly pointless truths to individuals, which would probably never impact their actions. However, the disanalogy with groups arises because the epistemic life of individuals is unconstrained by offices and charters with which groups are associated.

offers the example of a philosophy department believing that the best hummus in Chicago can be found at Whole Foods, since the philosophers hold that belief individually. Her claim is that attribution of such group belief isn't peculiar in itself even though we wouldn't overtly make it.

The first point about this line, if applied to our problem about attributing knowledge to coinciding groups, is that the pertinent knowledge clearly *is* relevant. The challenge is rather to *pair* it with the right group. But if groups are feature collectives, perhaps suitably arranged, it's mysterious why the knowledge should be relevant only to one group. Clearly, the question of relevance concerns the office and charter of a group. The second point is that if the idea is that one group has relevant knowledge, while another coincident group has irrelevant knowledge, then we need an account of how to apply that distinction to (KAP). How might irrelevant knowledge rationalize and motivate action, if at all? Does it do so to a lesser extent, or in less significant ways, than relevant knowledge?

Fortunately for the deflationist, a solution is available. Basically, the current problem arises because (DEFLATIONISM**) implies that group knowledge is a mere aggregate of knowledge of enough operative members, *irrespective* of their group membership. But groups can only operate in their office and by their charter. The obvious remedy is to further constrict the individual knowledge that may feed into the group. Consider:

(DEFLATIONISM***) Necessarily, group G knows that p if and only if there is a significant number of operative members of G who qua such membership know both that p and that the other operative members know that p .

To say that individuals have *knowledge qua members of G* means at a minimum that their knowledge bears some *relevance* to the structure, intentions, office or charter of G .²⁰ They may acquire the knowledge *through* occupying nodes in the social structure of G ; or its source may originate elsewhere but sustaining it bears on such occupancy, broadly understood to include any knowledge that may further the group's office or concern its charter. Explicating such knowledge involves taking the *places-are-objects* perspective from section 2, where the nodes themselves are treated as objects. The knowledge-qua constraint is designed to be vague enough to include the different kinds of knowledge had by members of diverse groups, but precise enough to exclude the possibility that groups be ascribed knowledge of no relevance to any of their characteristic features. Hence, it blocks the attribution of irrelevant knowledge to a group that is coincident with a distinct group to which that knowledge is relevant. To illustrate, let Prof Lockdown be a world-leading epidemiologist who in her capacity of sitting on SAGE advises the British Government on COVID-19 policy and strategy. Qua member of SAGE, she has knowledge of how proposed restrictions will affect infection rates. Suppose also that Prof Lockdown, qua treasurer of her local Tennis Club, has knowledge of its finances. Since the latter knowledge isn't qua SAGE membership, there is no ascription of it to SAGE, even if she is the most significant operative member of that group.

²⁰Suppose that, while the Squash Club is meeting, the son of one of the members is playing with matches in full view, which later causes a fire. Next day a newspaper headline says 'The Squash Club knew son was playing with matches and did nothing'. On our view, the members didn't know this proposition qua members; however, the headline is a convenient way of saying that these unnamed individuals did know it, and so individually they ought morally to have acted on it. Thanks to Stephan Torre for this case.

Importantly, individuals' knowledge-qua is sensitive to the *epistemic standards* by which different groups are governed. What this means is not that groups impose certain limits on knowledge of their members, or that individuals have some knowledge only relative to group membership. This isn't a claim about knowledge *simpliciter*, concerning whether they know, or what they know.²¹ What it means is rather that such standards place constraints on how members may *utilize* their knowledge in deliberation and action vis-à-vis the group, with a view to potentially constituting group knowledge.²² Suppose Prof Lockdown comes to know from a newspaper report that super-spreading events drive most COVID-19 infections. Even though the content is relevant to SAGE, the quotidian manner of acquisition prevents the knowledge from being of service to SAGE. For the standards Prof Lockdown needs to meet for knowledge qua member of SAGE are more demanding, i.e. the node she occupies in its social structure requires more epistemic work, if she is to exploit that knowledge for input to SAGE. Were she to consult the peer-reviewed study to which the newspaper refers, the knowledge would be apt to be included in the decision-making processes of SAGE, and so would be knowledge qua SAGE membership.²³

A worry springs to mind at this juncture. If knowledge requires belief also at the group level, then (DEFLATIONISM^{***}) implies that groups hold beliefs, but *psychologism about belief* is the plausible claim that an agent has a belief only if it has a mind. Does (DEFLATIONISM^{***}) therefore carry an untoward commitment to *group minds* and *group agency*, involving the group itself engaging in purposeful action? One response is to allow for groups to have knowledge without belief. Hakli (2007), Wray (2007) and Mathiesen (2011) propose that knowledge be understood in terms of justified true acceptances. Alternatively, one could admit that knowledge entails belief, but then interpret group belief in terms of joint acceptance by its members.²⁴ A third option is to embrace *psychologism about groups* in the *restricted* sense of a group agent with a mind who is capable of being in a limited range of purely intentional states.²⁵ Given our characterization of groups in section 2, it should come as no surprise that goal-oriented

²¹Relatedly, Gilbert (1989: 305) contrasts believing qua group member with believing period or personally. Schmitt (1994: 269–71) argues that for a doctor the *epistemic* justification of medical beliefs doesn't depend on *occupational roles* by setting a higher evidential standard than for a layperson. Whatever evidential basis suffices to justify a layperson's medical beliefs is also sufficient for a doctor; nor is such a claim true relative to context, i.e. if a doctor is justified relative to her personal life, she is also justified relative to her professional life. That's different from saying a doctor is *medically* justified in stating or acting on her beliefs in a professional capacity. We permit ordinary standards for epistemic justification but impose restrictions on reliance on epistemically justified beliefs in occupational activities.

²²As Mathiesen (2011: 30–1), following Meijers (2003: 379), notes, one's epistemic access to evidence isn't bound to any occupied group roles, but how one proceeds to reason, speak and act on that evidence is.

²³The reason the standards for knowledge qua SAGE are higher is arguably that SAGE serves a *gate-keeping* function in Greco's sense (2020: 40), i.e. that we make it harder for information to be acquired by our epistemic community in order that only high-quality information is underwritten by a government body, as opposed to the lower standard we set for distributing information once it has officially entered our community, so as to make it readily and widely available as and when needed.

²⁴Joint acceptance can be understood in terms of it being common knowledge that each member openly expresses a willingness to let the relevant proposition *p* stand as the view of the group, or in terms of openly expressing a commitment jointly to accept *p*, conditional on others expressing the same commitment. For more details see Gilbert (1989: 306), Tuomela (1992: 286; 2011: 89–90), Schmitt (1994: 262) and Hakli (2011: 124ff).

²⁵After all, as Gilbert (1989: 238) observed, the conjunction of psychologism about belief and anti-psychologism about groups entails that no group can believe anything.

states, as well as abilities to act, are ascribable to them. Just as we explain or rationalize individual patterns of behaviour in terms of fulfilling desires if beliefs are true, we can explain or rationalize patterns of group behaviour as being belief-desire-driven in a sufficiently systematic fashion.²⁶ As long as any pairs of group belief-desire are reducible to pairs of individual belief-desire, and observing (GMAP), such limited psychologism about groups is fully compatible with a deflationist outlook.²⁷

Lackey (2020: 71ff) raises other problems for deflationism, when understood as a view about doxastic justification to do with, for instance, conflicting epistemic bases for individual beliefs resulting in justified group belief in incoherent sets of propositions. The cases she adduces are complicated, as are the details of her Group Epistemic Agent Account (GEAA), but let's briefly highlight a few similarities and differences between (DEFLATIONISM^{***}) and (GEAA), assuming (GEAA) can be extended to knowledge. Both have it that a group has knowledge that *p* only if a significant number of operative members know that *p*, but the qua-membership and common knowledge conditions are absent on (GEAA). Moreover, unlike (DEFLATIONISM^{***}), (GEAA) imposes certain evidential constraints on group knowledge to ensure that adding bases of justified beliefs of a significant number of operative members issue in coherent sets of beliefs. On (GEAA), groups are thus epistemic agents in their own right, as there is no straightforward reduction of group knowledge, and yet such knowledge is firmly constrained by epistemic features of group members; or as Lackey (2020: 110) puts it, (GEAA) "neither inflates nor deflates group epistemology".

Let's take stock. (DEFLATIONISM^{***}) incorporates elements of our preferred metaphysical conception of groups, which means that group knowledge, by deflationist lights, isn't merely a question of aggregating any random knowledge by any majority of group members. Still, (DEFLATIONISM^{***}) preserves the tenet that group knowledge is reducible proposition-wise to individual members' knowledge. There is thus nothing about knowledge of structured groups per se that prevents such reduction; nor does (DEFLATIONISM^{***}) entail any strong views about groups having *sui generis* minds, or "minds of their own".²⁸

4. Inflationism about group knowledge

Section 3 argued that deflationists should adopt (DEFLATIONISM^{***}), as it avoids several relatively minor worries, but the question of its truth when confronted with more difficult problems wasn't addressed. *Some* group knowledge clearly meets the embedded conditions. SAGE possesses knowledge of UK infection rates, because (i) Prof Lockdown and the other epidemiologists on SAGE count as a significant number of operative members, (ii) they have that knowledge qua SAGE membership, and (iii) they know that the other operative members also have it. So, no inflationist should claim that *all* knowledge of groups is irreducible to knowledge of their members, nor therefore that no group knowledge even supervenes on individual features. Such

²⁶See also List and Pettit (2011: Ch. 1).

²⁷Nor is the deflationist committed to the claim that group minds necessarily satisfy any rationality constraint. List's (2014) notion of a group agent attitude (what he calls "corporate attitude") isn't reducible to individual attitudes on each proposition, in fact even proposition-wise supervenience fails, but that is because he imposes a stringent requirement of robust rationality on groups according to which groups must form consistent and complete attitudes across variations in underlying individual attitudes.

²⁸Pettit (2003).

extreme views are implausible. Still, (DEFLATIONISM^{***}) says that as a matter of necessity all group knowledge is reducible. This section offers counterexamples to this last claim, which establish an *epistemic divergence* between groups and their members.²⁹ The resultant view is inflationism, which rejects the *necessity* of the bi-conditional in (DEFLATIONISM^{***}), and so is consistent with some group knowledge being reducible:

(INFLATIONISM) Possibly, group *G* knows that *p* even if there isn't a significant number of operative members of *G* who qua such membership know both that *p* and that the other operative members know that *p*.

But even if some group knowledge is irreducible, we maintain that all such knowledge (strongly) supervenes on *individual features* by which we include the totality of individual states, individual processes of deliberation and state formation, as well as the functional relations of the social structure into which individuals enter:

(SUPERVENIENCE) Necessarily, if group *G* knows that *p*, then there are individual states *S*, processes *P* and relations *R* such that a significant number of operative members of *G* instantiate *SPR*, and necessarily, if a significant number of operative members of any other group *G** instantiate *SPR*, then *G** knows that *p*.³⁰

While (INFLATIONISM) represents the *minimal* commitment of all inflationists, we defend a *weak inflationist* view which combines (INFLATIONISM) with (SUPERVENIENCE). On this view, some group knowledge is irreducible and yet all such knowledge supervenes on individual features. To be clear, the aim is not to vindicate a specific (commitment or distributed) model of weak inflationist group knowledge, but merely to show that any viable account of group knowledge must respect both (INFLATIONISM) and (SUPERVENIENCE). In contrast, *strong inflationism* is the conjunction of (INFLATIONISM) and the negation of (SUPERVENIENCE). On this view, not only is some group knowledge not reducible to individual knowledge, but some such knowledge even fails to supervene on individual features. As we shall see, a familiar case that is alleged to support strong inflationism is found wanting.

We showed in section 3 that a group is confined to its office and charter, in that it cannot step completely outside its office or dispense with its entire charter without ceasing to exist. Consequently, a group may have sub-groups as members, indeed distinct groups may share the same sub-group, but only if group and sub-group function in *roughly* the same office and by the same charter. The Squash Club could join the World Squash Federation, and SAGE could join an international network of advisory bodies on public health, but not the other way around. In contrast, individuals can serve as members of groups with dissimilar offices and charters; indeed, group

²⁹For other divergence cases, see Hutchin's (1995) ship crew, Tollefsen's (2009) NAS panel, Hakli's (2011) weather forecast agency and Mathiesen's (2011) hiring committee.

³⁰To say the supervenience base consists of individual features is not to restrict it to intrinsic properties; on the contrary, not only are the properties of standing in relations of the social structure extrinsic, so are also the properties of instantiating the relevant states. Since the supervenience claim must respect the factivity of knowledge, being in some of those states will imply the truth of their contents. Whether those states are also mental is not a question we shall settle here. It depends in part on whether knowledge itself is a mental state.

membership by no means exhausts their epistemic agency as they attain a wealth of knowledge outside any group membership.³¹

Now, as mentioned in section 3, part of a group's charter concerns the possibly special *epistemic standards* it employs, e.g. standards of permissible evidence or of weighing reasons, as devolved from its office, or determined by its charter. While individuals' knowledge-qua (group membership) is constrained by the standards of the groups to which they belong, their knowledge *simpliciter* is not. And since the question of knowledge-qua only arises for such knowledge they already possess, it's possible through different standards that *defeating evidence* is available to members individually but not collectively. Such divergence cases put (DEFLATIONISM^{***}) under severe pressure.

Take criminal proceedings in a UK court of law for which the standards of evidence include that hearsay normally be excluded, and that the standard of proof be beyond a reasonable doubt. These special standards govern how the criminal court must perform as devolved from the office of such courts. Consider the following:

(CRIMINAL COURT) A defendant is on trial for the crime of careless driving. The prosecution adduces evidence from the police report, CCTV footage, etc., as well as eyewitnesses testifying in court that the defendant was driving the van that hit the victim. The jury finds the evidence beyond a reasonable doubt, and hence sufficient to validate criminal conviction. All the members of the jury have hearsay evidence from a reliable, yet as it turns out on this occasion incorrect, source that the defendant was at home at the time of the incident. But the judge instructs the jurors to ignore this evidence as it fails to meet the conditions for being ruled admissible in a criminal court. Consequently, the jury knows that the defendant is guilty, yet none of the individual jurors know this proposition.

Since the evidence available to the jurors distributively is unavailable to them jointly, (CRIMINAL COURT) poses a (left-to-right) counterexample to the bi-conditional in (DEFLATIONISM^{***}).³² The group has knowledge, but there isn't a significant number of operative members who qua members also have that knowledge, let alone common knowledge. Remember that how the jurors may deploy any knowledge vis-à-vis the jury is fixed epistemically by its legal charter, but which knowledge they possess is not. In effect, the jurors have a *rebutting mental state defeater*, namely a (justified) belief indicating that their belief in the defendant's guilt is false, which prevents them from knowing what the jury knows.³³ To say they do have the knowledge, but only qua jury members, is implausible, as it doesn't suddenly spring into existence in court only to vanish again upon adjournment. Of course, the jurors may acquire a defeater

³¹Hawley (2017: 403–5) claims that groups with co-extensional memberships is the collective version of the single person who fills more than one social role. That may be true as far as the metaphysics is concerned, but there are crucial differences in their epistemic agency.

³²Imagine instead the hearsay evidence was both reliable and correct. In that case, the jurors, but not the jury, know the defendant is innocent, but their knowledge isn't qua jury member, because it fails to meet the standard imposed by the jury. So, we don't have a counterexample in the other direction. Group members may thus know that *p* even though the group does not, which can happen when evidence for *p* is available to the members individually but not collectively. In those cases, the members don't know that *p* qua group members.

³³Carter (2015) adopts group-level epistemic defeat to advance a dilemma even for weaker versions of inflationism about group knowledge, but he assumes on their behalf the joint acceptance account of belief formation (cf. fn. 24), which isn't integral to our view.

for the original defeater if they subsequently learn that the hearsay was flawed, thus paving the way for them to attain the knowledge. The jury, however, has no such first-order psychological defeater, and so will be in the know all along.

Either way, (CRIMINAL COURT) lends support not just to (INFLATIONISM), but also to (SUPERVENIENCE) in that any other court (or group) indiscernible in terms of the instantiation of individual states, processes, and relations by a significant number of jurors (or operative members) is also indiscernible with respect to knowledge of the defendant's guilt.³⁴ That means (CRIMINAL COURT) supports weak, but not strong, inflationism.

Take instead a jury in a UK civil court, where the standards of evidence in civil litigation include that hearsay is normally regarded as admissible evidence, and that the standard of proof be by a preponderance of the evidence. Again, these special epistemic standards govern how the civil court must operate in its legal capacity. Consider now the following:

(CIVIL COURT) Just as in (CRIMINAL COURT), a van is alleged to have hit a victim in a traffic accident, and the victim is now suing the driver for damages in a civil court. A jury comprising the same jurors as in (CRIMINAL COURT) is presented with the same police report, CCTV footage, etc., and the same witnesses in court as in (CRIMINAL COURT). After all, evidence from a criminal trial is typically also admissible in a civil action about the same matter. Moreover, all the jurors have hearsay evidence from a reliable, yet on this occasion incorrect, source that the defendant isn't responsible for the accident. Since the judge rules such hearsay admissible, neither the jurors nor the jury are in a position to know the driver is guilty and so cannot find the defendant liable on balance of the probabilities.

In (CIVIL COURT), all and only the same evidence is available to the jurors distributively and jointly, and so no epistemic divergence ensues. In (CRIMINAL COURT) and (CIVIL COURT) different epistemic standards apply, which in turn lead the jurors to engage in different legal deliberations about the available evidence, and on that basis reach contrasting verdicts. Still, in neither case do the jurors possess knowledge. The upshot is therefore that such different standards, as well as the distinct individual processes and states they imply, are sufficient for establishing a divergence in knowledge between distinct yet coincident groups, as well as between a group and its members.

Lackey (2020: 69–70) objects that deeming hearsay evidence inadmissible in a court of law has *no epistemic rationale* and reveals nothing about its epistemic credentials. The problem with such evidence is *procedural* in that the opposing side is denied the possibility of cross-examining the source of the information. The evidence may be highly truth-conducive and so ought to be factored into the epistemic basis of the jury's belief. What evidence a judge deems admissible is exclusively a question about which verdicts a jury may be legally justified in making. These *merely legal standards* have no epistemic bearing.

The reply is that the objection illicitly assumes that a criminal court can operate epistemically in ways other than in its legal capacity, so as to be legally, but not

³⁴The relevant (mental) states in the supervenience base include belief, acceptance and even knowledge. Note also that just as we distinguished between an individual knowing *simpliciter* and knowing qua group member, a similar distinction can be drawn between (justifiably) believing qua group member, which could feed into group (justified) belief, and believing *simpliciter*, which is different from outright, as opposed to graded, belief.

epistemically, rational in abiding by the standard of excluding hearsay evidence. But this standard is part and parcel of the charter of a criminal court, and its epistemic agency is constrained within the bounds of its legal office and charter without which it has no epistemic life.³⁵ It cannot fail to comply with that charter while fully function in that capacity. If, nevertheless, the court were to embrace such evidence, it would thus partially contravene its constitution. Not so for individuals, but the objection better not assume that groups ought to operate on a par with (collectives of) individuals, which is what our account of groups as structured wholes opposes. Having said that, a group may break at least some of its constitutive rules without ceasing to exist. If the judge deliberately instructs the jurors to consider the hearsay evidence, then the jury would possess the same defeater as the jurors, robbing it of knowledge of the defendant's guilt. The jury remains the same, although the distinction between criminal and civil juries would begin to blur. But that doesn't show the jury lacks knowledge of that proposition in (CRIMINAL COURT) where the judge *does* instruct the jury to disregard such evidence to ensure the jury functions in its office and by its charter. Applying the standard is thus epistemically rational for the jury; indeed in (CRIMINAL COURT), it is what makes the difference in knowledge between the jury and the jurors. So, at the group level this standard can positively impact epistemic standing, whereas the standard of *including* hearsay evidence, as in (CIVIL COURT), can have negative epistemic effects.³⁶

A related worry, following Lackey (2020: 64–7), is that linking epistemic standing with the office or charter of groups makes it possible to *illegitimately manipulate* the former by randomly changing the latter in non-epistemic ways that sever the connection between truth and such standing. After all, there are supposedly no constraints in principle on which standards may govern any given group; for example, imagine circumstances in which criminal courts only dismiss hearsay evidence from left-handed speakers.

The answer begins by recognizing that while groups are created by convention or through practice, and so are by no means *natural kinds*, there's limited scope for stipulation of epistemic standards. For the component parts of groups, i.e. their realized social structures, joint intentions, offices and charters, must form a coherent whole both internally and in relation to society at large.³⁷ A legal court must conform to both statute and common law such as court procedure rules derived from the inherent jurisdiction of the court. So, in the normal run of things, there is no question of a jury arbitrarily deciding to deviate from their office or charter, including epistemic standards. (CRIMINAL COURT) and (CIVIL COURT) aren't cases in which members decide to let a view stand as the group's view through disagreement-resolving compromise, pragmatic considerations, or personal preference. Gilbert (2004: 98) envisages an appointment committee determining that a certain applicant is the most qualified for a job despite none of its members having that applicant as their front-runner. What matters on her account is that all members have "indicated their readiness to let the belief in question be established as the group's belief". Unlike Gilbert's case where the group

³⁵See also Schmitt (1994: 274).

³⁶Goldman (1999: 292–5) and Mathiesen (2011: 34–5) argue that including hearsay evidence in criminal trials actually makes the court less likely to reach the correct verdict.

³⁷Following Goldberg (2018), the *epistemic expectations* we have of individuals also apply to groups, both basic ones about reliability and coherence, as well as non-basic ones that stem from the agent's social position. In particular, assuming groups are capable of partaking in the same testimonial practices as individuals, similar epistemic standards must apply to group testimony; otherwise, contamination of individual knowledge ensues. In so far as a group's charter pertains to its involvement in such practices, it therefore has to prescribe belief-forming processes that meet the epistemic expectations we have of individuals.

holds a belief that isn't shared by any of its members, the jurors in (CRIMINAL COURT) believe the defendant is guilty; they just have a mental state defeater for that belief which prevents knowledge. So, group belief in (CRIMINAL COURT) is down to individual members' belief.³⁸ But if they were to abandon their belief upon receiving the hearsay evidence, one might naturally fall back on a Gilbert-style account of group belief in terms of joint acceptance as an explanation of how the jury forms its belief.³⁹

Our second case, following Bird (2010: 34; 2014: 57–8), exploits the division of cognitive labour in research teams⁴⁰:

(RESEARCH TEAM) Dr X, a physicist, and Dr Y, a mathematician, collaborate on a project to demonstrate the truth of the conjecture that q . Their project can be broken down into three parts. Part one is a problem in physics, the problem of showing that p , which will be the work of Dr X alone. Part two is a problem in pure mathematics, that of proving $p \rightarrow q$, for which Dr Y takes sole responsibility. Part three is an application of *modus ponens* to the results of parts one and two. They arrange for an assistant to publish the paper if and only if the assistant receives from Dr X the demonstration that p is true and from Dr Y the proof of $p \rightarrow q$ (the brief final part with the application of *modus ponens* has been prewritten). We can imagine that Drs X and Y have no other communication with each other or with the assistant and so do not know at the time of publication that q has been proven.

Four observations are in order. First, the team can be said to *dispositionally believe* q in the sense of having information q readily available for endorsement. What is needed for knowledge is precisely such dispositional belief,⁴¹ which sits between occurrent belief,

³⁸Note also that the group belief isn't "base fragile" in Lackey's sense (2020: 45) of the bases of the jurors' beliefs being in conflict, or irrelevant to the team's belief. Nor is the group belief riddled with what she calls (2020: 41–4) "judgment fragility" in the sense that the jurors would easily have reached a different judgment, were they to deliberate about the same body of evidence with no relevant change in the information that emerges via the deliberation. As long as the jurors function in the legal roles they occupy, not easily would they have come to a different judgment on the basis of the same body of evidence.

³⁹See fn. 24. Lackey (2020) is worried that joint acceptance accounts of group justification make it too easy to come by, in that a group can ignore or fabricate evidence through what it chooses to accept or reject, and thereby wind up with beliefs that wrongly count as justified. The problem with these accounts is supposedly that joint acceptance can be guided by factors that are disconnected from the truth, such as practical interests. After all, unlike belief, acceptance is under voluntary control. One response is to let the nodes that individuals occupy in the social structure of a group constrain how joint acceptance of a reason as the group's reason by their members may constitute group justified belief. Joint acceptance by jurors of a reason wouldn't then qualify for justified jury belief unless that reason pertained to the office or charter of the jury. Furthermore, following McMahan (2003) and Mathiesen (2011), joint acceptance can be restricted to *epistemic* reasons, i.e. for the purpose of attaining truth and avoiding error. Alternatively, Schwengerer (Forthcoming) argues that even though joint acceptance accounts allow for the inclusion of arbitrary prima facie reasons for group belief, such reasons can be normatively defeated by considerations about what evidence a group should have had, and thereby preventing ultima facie justified group belief.

⁴⁰(RESEARCH TEAM) is of course a toy example for illustrative purposes, but science is replete with division of labour on a massive scale. For example, see Knorr Cetina (1999) for discussion of the High Energy Physics experiments at CERN, which preceded the Large Hadron Collider experiments. See also Hutchins (1995) on socially distributed cognition.

⁴¹Rose and Schaffer (2013: 21ff).

i.e. a thought currently endorsed, in effect the mental counterpart of overt assertion, and a disposition to believe, i.e. information which would be endorsed, were it available to mind through a process of belief formation. Whereas dispositional beliefs manifest upon considering whether q is true, dispositions to believe also require a recognition of an appropriate relationship between q and available reasons.⁴² Secondly, the team's knowledge that q is *through* or *because of* the skills, action and knowledge of its members. When combined adequately, the team possesses no skills and performs no action in addition to what the members contribute, and yet the team still ends up with knowledge that none of its members have.⁴³ Thirdly, while (RESEARCH TEAM) relies on a closure principle of deductive knowledge to pose a counterexample to (DEFLATIONISM***), similar cases could *mutatis mutandis* make use of inductive or abductive knowledge. Fourthly, the team's knowledge isn't proposition-wise reducible to knowledge of its members, but it does supervene on the totality of individual features, encompassing individual states, individual processes of deliberation and belief formation, as well as the functional relations of the social structure into which individuals enter. In that sense, the team's epistemic agency is anchored in individual features.⁴⁴ Consequently, (RESEARCH TEAM) is evidence in support of (INFLATIONISM), but since the case conforms to (SUPERVENIENCE), it provides reason to accept weak, rather than strong, inflationism. Bird (2014: 58) is right that this case refutes the supervenience of group knowledge on individual *mental states*; not because group knowledge mysteriously emerges given a certain complexity at the individual level, but because the supervenience base extends beyond such states. As per (SUPERVENIENCE), individual features *addition* to mental states must be included in the base, but, pace Bird (2014: 58), "non-human" vehicles for representing information that aren't integrated into human cognition should be left out.⁴⁵ On his take on (RESEARCH TEAM), the venue of publication is one such feature, as it makes the information q available for access. However, what matters for group knowledge is that the assistant goes through the pre-arranged process of deliberately applying modus ponens to p and $p \rightarrow q$ upon receiving the instructions from Drs X and Y, irrespective of whether he proceeds to actually publish the result.⁴⁶

Let's briefly ponder Bird's notion (2010, 2014) of "social knowledge". On his view, *social knowers* are social structures which have characteristic propositional outputs and mechanisms which function to ensure that those outputs are true and inputs for social action. More precisely, Bird (2010: 37) takes social knowers to be "organic groups",

⁴²See also Audi (1994), Rose and Schaffer (2013) and Bernecker and Grundmann (2019). Note that (RESEARCH TEAM) can be tweaked to involve individual occurrent belief. Suppose the assistant forms such belief in q on the basis of competent deductive inference from p and $p \rightarrow q$ of which the assistant has prior knowledge. Still, the assistant may fail to know q ; perhaps some defeater destroys the knowledge, or perhaps p and $p \rightarrow q$ are too complicated to remember while drawing the inference to q .

⁴³See also Kallestrup (2020).

⁴⁴Schmitt (1994), Gilbert (1996, 2004), Pettit (2003) and Tuomela (2004) all hold that social states are distinct from, but supervene on, individual mental states. These inflationist views are thus more restrictive than (SUPERVENIENCE), which also includes non-mental individual features in the supervenience base for group knowledge.

⁴⁵The best way to demarcate cases of extended cognitive processes from mere background or auxiliary resources that causally influence a cognitive system without constituting parts of it is to impose Clark's (2010: 46) "trust and glue" conditions for a non-biological entity to count as integrated within such a system, thus ensuring functional isomorphism between, say, biological memory and an extended memory system. See also Carter and Kallestrup (2020).

⁴⁶See also fn. 50.

involving bonds that arise out of the mutual interdependence and cognitive cooperation brought about by the division of cognitive labour. These bonds are the social glue that joins components of such groups together to form a social epistemic subject. Examples include the scientific community or wider science, which are *distributed cognitive systems* with goals and functions, although not by design, and without significant collective commitment. These systems are bearers of “social knowledge”, which is defined in terms of information accessible to their members as and when they need it. “Social knowledge” draws upon a functional analogy with individual knowledge, yet it doesn’t supervene on individual mental states, or indeed on any other features of individuals, and so represents a form of strong inflationism.

By our lights, entire scientific communities form no structured group, and so have no capacity for group knowledge. These are at best feature collectives to whom aggregate knowledge can be ascribed. Pace Bird (2010), they merely exhibit Durkheim’s (1893/1994) “mechanical solidarity”, by which their members share certain views, values and attitudes. Even though the division of cognitive labour brings about interdependence and cooperation among working scientists, a scientific community doesn’t form a unified whole with a shared goal.⁴⁷ But specific research teams clearly do. Our notion of social structure, as exemplified by such teams, are more akin to Durkheim’s (1893/1994) idea of “organic solidarity”, owing their cohesion to functional dependence and joint commitment, which draw or bind researchers together who fill places in social structures of scientific groups.⁴⁸

Let’s dwell on two objections to the foregoing, both inspired by Lackey (2020: 116–27). First off, in the case of individuals, there’s a familiar distinction between *knowing* and *being in a position to know*, which pertains to the difference between accessed information and accessible information. If in (RESEARCH TEAM) nobody has actually read the article demonstrating the truth of the conjecture that q , then, by analogy, it’s more natural to describe the case as one in which the group is merely in a

⁴⁷Tollefsen (2009, 2011) argues that the Wikipedia community constitutes a group, capable of engaging in joint deliberation. Its members play designated roles in a democratic structure, e.g. editor, author, administrator, and are subject to guidelines, policies and decision-making processes developed by contributors, via the consensus approach. However, many Wikipedia articles are cognitively unstable in a way the outputs of typical epistemic agents are not, and this community is often unwilling to accept epistemic responsibility. Thus Tollefsen (2009, 2011) takes Wikipedia to be a somewhat immature yet developing group. More generally, Greco (2020: 25) suggests that an epistemic community be defined as a “group of cognitive agents engaged in shared information-dependent tasks and sharing norms for evaluating information associated with those tasks”, but such a community clearly doesn’t exhibit the requisite social structure to be classified as a group in our sense.

⁴⁸See also Wray (2007). Habgood-Coote and Tanswell (2021) take the classification theorem (for finite simple groups) to show that when its proof was finished in 2004, the group theory community had objectual knowledge thereof, which involves understanding the proof as well as (tacitly) knowing how to carry out its steps. In fact, they claim the community knew the theorem by exercising their know-how in having carried out the proof. While individual mathematicians contributed sub-proofs, none of them possessed the connective knowledge of how these fit together to form the entire proof. Its immense complexity, running to thousands of pages, means no single mathematician understood all its details. The knowledge is therefore not “distributive”; nor is it “collective” since the community lacks the required internal structure and collective intentions to form a robust group. Rather, it is “cumulative” in that *between them* members of the community have knowledge of the proof by having knowledge of its parts. By our lights, the knowledge obviously falls short of group knowledge too; however, since it supervenes on the totality of individual mental states, it doesn’t qualify for Bird’s “social knowledge” either. After all, the proof arguably provided undefeated epistemic support for individual members to know that the theorem is correct.

position to know q , since not a single member has actually accessed q . Group knowledge on this model is fundamentally a matter of information being accessible to group members; or so the thought goes.

In reply, we appeal to the distinction between *dispositional* and *occurrent* belief. For while the research team doesn't occurrently believe q , because q isn't presently considered, let alone endorsed, the team can be said to dispositionally believe q , because that proposition is readily available to Drs X and Y for endorsement. Having considered q carefully, together they decide to try to prove that conjecture. Were they to correspond after the proof, they would together easily be able to retrieve q from biological memory for endorsement without recognition anew that available reasons epistemically support q . Keep in mind that a dispositional belief doesn't require a prior occurrent belief, e.g. if a car passes by while absorbed in deep conversation, you may form the dispositional belief that a car drove by even if you aren't consciously aware of that proposition.⁴⁹ What matters for dispositional belief is that the proposition be preserved in biological memory, or similar place, from which it can effortlessly be recovered for active deployment in reasoning or planning. Since dispositional belief is exactly what is needed for knowledge, the team knows that q .

Lackey (2020: 116–27) discusses another of Bird's cases (2010: 32, 2014: 57) which is suitable to illustrate his notion of "social knowledge" of a distributed cognitive system. Consider:

(DR N) Dr N is working in mainstream science, but in a field that currently attracts only a little interest. He makes a discovery, writes it up and sends his paper to the Journal of X-ology, which publishes the paper after the normal peer-review process. A few years later, Dr N has died. All the referees of the paper for the journal and its editor have also died or forgotten all about the paper. The same is true of the small handful of people who read the paper when it appeared. A few years later yet, Professor O is engaged in research that needs to draw on results in Dr N's field. She carries out a search in the indexes and comes across Dr N's discovery in the Journal of X-ology. She cites Dr N's work in her own widely read research and because of its importance to the new field, Dr N's paper is now read and cited by many more scientists.

Bird (2010: 32, 2014: 57) claims that Dr N's contribution to what is known in wider science survived his death in virtue of the publication being accessible throughout, and so (DR N) supposedly illustrates the importance of non-human forms of representation in distributed cognition cases. What matters, on Bird's view, for wider science to know Dr N's discovery is its availability in a library or on the web, rather than occupying the minds of scientists. So, (DR N) seems to show that "social knowing" conforms to strong inflationism, keeping in mind that social knowers are equated with so-called "organic groups".⁵⁰

⁴⁹Audi (1994: 421).

⁵⁰There is a way of interpreting Bird's talk of "non-human entities" so as to make his view compatible with (INFLATIONISM), namely if the supervenience base for "social knowledge" includes individual dispositions to be able to access information within a scientific community. Thanks to an anonymous referee for this suggestion. However, unlike (RESEARCH TEAM) where Drs X and Y consider the conjecture that q , (DR N) is such that no individual scientists would seem to have such disposition vis-à-vis Dr N's results during the interval between Dr N's death and Prof O's discovery of those results, because no scientists are even aware of their existence, let alone their publication in the Journal of X-ology.

However, Lackey is surely right to challenge the claim that the scientific community (or wider science) knew the results of Dr N's paper all along, despite there being a period of time when everyone who knew about it was dead. For the community doesn't dispositionally believe, and hence doesn't know, those results. The reason is that the results aren't readily available for endorsement in the required sense: they aren't stored in biological memory, or an extended version thereof, or indeed in anyone's 'belief box'. The case merely shows that the scientific community has a disposition to believe, and if everything goes to plan, know, those results. By talk of "libraries as repositories of knowledge" (2010: 33), Bird's notion of "social knowledge" obscures the distinction between actually knowing and having a disposition, or being in a position, to know, by illicitly assuming that storage in a university library, or uploading to a journal website, suffices for knowledge. To illustrate the epistemic difference between (RESEARCH TEAM) and (DR N), the putative "social knowledge" attributed to the scientific community is lost if the library burns down, or the journal webpage is deleted, without any individual members undergoing a change, but the group knowledge attributed to the research team can survive either, so long as relevant individual features remain fixed. The upshot is that "social knowledge" fails to constitute group knowledge, and so there is no reason to include "non-human entities" in the supervenience base for group knowledge. That means (DR N) provides no support for strong inflationism; rather, as (SUPERVENIENCE) dictates, and as (RESEARCH TEAM) illustrates, all knowledge of groups is fixed by states and processes of, and functional relations between, their individual members.

The second objection, also due to Lackey (2020: 115–23), revolves around *group agency*. The worry is how a group can act on knowledge, if none of its members have it, and the only way for a group to act is through the contributing actions of its members. Let's spell this out in more detail. According to the *Knowledge-Action Principle* from section 3:

(KAP) Agent *S* knows that *p* if and only if *S* is epistemically rational to act as if *p*.

So, if our team in (RESEARCH TEAM) knows that *q*, it is epistemically rational for it to act accordingly, but that in turn requires that at least one of Drs X and Y perform a causally contributing act. Consider again the *Group/Member Action Principle*:

(GMAP) For every group, *G*, and act, *a*, *G* performs *a* only if at least one member of *G* performs some act or other that causally contributes to *a*.

The alleged problem arises when (KAP) and (GMAP) are combined with the team knowing *q* despite *none* of Drs X or Y having that knowledge. For the team to act on *q* requires some of its members to perform causally contributing acts, but how is that possible if none of them share the group's knowledge that *q*? Since they trivially cannot act on knowledge they lack, what *else* could explain their contributing act?

The answer is that the division of causal labour latches onto a corresponding division of cognitive labour. In (RESEARCH TEAM) Drs X and Y acquire knowledge that *p* and that $p \rightarrow q$, respectively, on which they each act by forwarding their demonstration and proof on to the assistant. Not only do such acts explain, by causally contributing to, the group act of submitting the article for publication, they are also based on their respective knowledge, which together explain how the team ends up knowing that *q*. None of these contributing acts are accidental, in that Drs X and Y have carefully planned in advance who is responsible for which sub-tasks and how to coordinate

their completion.⁵¹ In contrast, to ascribe “social knowledge” to a scientific community on the basis of Dr N’s article being filed in a library or repository would make a mystery of how its members could possibly contribute to some act allegedly performed by the entire community as they have no knowledge (during the relevant interval) that would non-accidentally have any bearing on the findings in Dr N’s study. So, in short, this objection from group agency would apply to a scientific community having “social knowledge”, but not to a research team having group knowledge.

5. Concluding remarks

Let’s sum up. We began by drawing a metaphysical distinction between structured groups and feature collectives, in order to show that while the latter can be attributed aggregate knowledge, only the former can enjoy group knowledge. Pace traditional deflationism, there is nothing about this view per se which prevents groups, when understood as realizations by individuals of social structures, from having knowledge which is proposition-wise reducible to knowledge of their members. (DEFLATIONISM^{***}) is exactly this more sophisticated view, which has the added advantage of being able to accommodate group agency and epistemic differences between coincident groups. We proceeded to argue that even though (DEFLATIONISM^{***}) does capture some instances of group knowledge, this view is ultimately unviable as it faces knockdown counterexamples, either because the epistemic standards of groups differ from those of their members, or because the division of cognitive labour implies that groups can possess knowledge that their members lack. That means we should adopt (INFLATIONISM), which is exactly the view that groups can have knowledge over and above their members. However, in neither case does (SUPERVENIENCE) fail, i.e. the claim that group knowledge supervenes on individual states, processes and relations; nor do putative cases of so-called “social knowledge” demonstrate the failure of (SUPERVENIENCE). Consequently, none of the cases considered support *strong inflationism*, i.e. the view that some group knowledge is not only non-reductive but also non-supervenient, thus understood. What they recommend is rather *weak inflationism*, according to which some group knowledge is irreducible to individual knowledge, yet all group knowledge supervenes on, and so in that sense is firmly grounded in, features of individual members. We defended this view against a number of objections, which helped further characterize both the epistemic profile of groups as functioning in their offices and by their charters, and the agency of groups as depending on their members knowingly contributing relevant acts.⁵²

References

- Audi R. (1994). ‘Dispositional Beliefs and Dispositions to Believe.’ *Noûs* 28(4), 419–34.
- Bernecker S. and Grundmann T. (2019). ‘Knowledge From Forgetting.’ *Philosophy and Phenomenological Research* 98(3), 525–40.

⁵¹In particular, the research team’s knowledge doesn’t display the kind of base or judgment fragility, mentioned in fn. 38. For example, the bases of Drs X’s and Y’s beliefs aren’t in any way in conflict, or irrelevant to the team’s belief. Nor are they in the kind of secrecy situation we mentioned in section 3, which by deflationist lights requires common knowledge of the target proposition. Rather, in (RESEARCH TEAM) the common knowledge concerns at most their joint (conditional) commitment to intend as a group to prove the conjecture that *q*, and thus being ready to share in acting together to try to demonstrate that *q*.

⁵²Many thanks to Adam Carter, Finnur Ulf Dellsen, Fed Luzzi, Luca Morreti, Orestis Palermos, Lukas Schwengerer, Mona Simion, Martin Smith, Stephan Torre and an anonymous referee for this journal for helpful and constructive comments on an earlier version of this paper.

- Bird A.** (2010). 'Social Knowing: The Social Sense of 'Scientific Knowledge.' *Philosophical Perspectives* **24**, 23–56.
- Bird A.** (2014). 'When is There a Group that Knows? Scientific Knowledge as Social Knowledge.' In J. Lackey (ed.), *Essays in Collective Epistemology*, pp. 42–63. Oxford: Oxford University Press.
- Bratman M.** (1993). 'Shared Intention.' *Ethics* **104**, 97–113.
- Bratman M.** (1999). *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- Carter J.A.** (2015). 'Group Knowledge and Epistemic Defeat.' *Ergo* **2**(28), 711–35.
- Carter J.A. and Kallestrup J.** (2020). 'Varieties of Cognitive Integration.' *Noûs* **54**(4), 867–90.
- Clark A.** (2010). 'Memento's Revenge: The Extended Mind, Extended.' In R. Minary (ed.), *The Extended Mind*, pp. 43–66. Cambridge, MA: MIT Press.
- Corlett J.A.** (1996). *Analyzing Social Knowledge*. Lanham, MD: Rowman and Littlefield.
- De Ridder J.** (2014). 'Epistemic Dependence and Collective Scientific Knowledge.' *Synthese* **191**(1), 37–53.
- Durkheim E.** (1893/1994). 'On Social Facts.' In M. Martin and L. McIntyre (eds), *Readings in the Philosophy of Social Science*, pp. 433–40. Cambridge, MA: MIT Press.
- Effingham N.** (2010). 'The Metaphysics of Groups.' *Philosophical Studies* **149**, 251–67.
- Fantl J. and McGrath M.** (2002). 'Evidence, Pragmatics, and Justification.' *Philosophical Review* **111**(1), 6–94.
- Fantl J. and McGrath M.** (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Fine K.** (2020). 'The Identity of Social Groups.' *Metaphysics* **3**(1), 81–91.
- Giere R.N. and Moffatt B.** (2003). 'Distributed Cognition: Where the Cognitive and the Social Merge.' *Social Studies of Science* **33**(2), 301–10.
- Gilbert M.** (1989). *On Social Facts*. Princeton, NJ: Princeton University Press.
- Gilbert M.** (1994). 'Remarks on Collective Belief.' In F. Schmitt (ed.), *Socializing Epistemology*. Lanham, MD: Rowman and Littlefield.
- Gilbert M.** (1996). *Living Together: Rationality, Sociality, and Obligation*. Lanham, MD: Rowman and Littlefield.
- Gilbert M.** (2004). 'Collective Epistemology.' *Episteme* **1**(2), 95–107.
- Gilbert M.** (2010). 'Belief and Acceptance as Features of Groups.' *Protosociology: An International Journal of Interdisciplinary Research* **16**, 35–69.
- Goldberg S.** (2018). *To the Best of our Knowledge*. Oxford: Oxford University Press.
- Goldman A.** (1999). *Knowledge in a Social World*. New York, NY: Oxford University Press.
- Greco J.** (2020). *The Transmission of Knowledge*. Cambridge: Cambridge University Press.
- Habgood-Coote J. and Tanswell F.** (2021). 'Group Knowledge and Mathematical Collaboration: A Philosophical Examination of the Classification of Finite Simple Groups.' *Episteme*. <https://doi.org/10.1017/epi.2021.26>.
- Hakli R.** (2007). 'On the Possibility of Group Knowledge without Belief.' *Social Epistemology* **21**(3), 249–66.
- Hakli R.** (2011). 'On Dialectical Justification of Group Belief.' In H.B. Schmid, D. Sirtes and M. Weber (eds), *Collective Epistemology*, pp. 119–56. New York, NY: Walter de Gruyter.
- Haslanger S.** (2016). 'What is a (Social) Structural Explanation?' *Philosophical Studies* **173**, 113–30.
- Hawley K.** (2017). 'Social Mereology.' *Journal of the American Philosophical Association* **3**(4), 395–411.
- Hawthorne J. and Stanley J.** (2008). 'Knowledge and Action.' *Journal of Philosophy* **105**(10), 571–90.
- Hutchins E.** (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Kallestrup J.** (2019). 'Groups, Trust and Testimony.' In K. Dormandy (ed.), *Trust in Epistemology*. New York, NY: Routledge.
- Kallestrup J.** (2020). 'Group Virtue Epistemology.' *Synthese* **197**, 5233–51.
- Kitcher P.** (1994). 'Contrasting Conceptions of Social Epistemology.' In F.F. Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*, pp. 111–34. Lanham, MD: Rowman and Littlefield.
- Knorr Cetina K.** (1999). *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.
- Kosllicki K.** (2008). *The Structure of Objects*. Oxford: Oxford University Press.
- Lackey J.** (2012). 'Group Knowledge Attributions.' In J. Brown and M. Gerken (eds), *New Essays on Knowledge Ascriptions*, pp. 243–69. Oxford: Oxford University Press.
- Lackey J.** (2020) *The Epistemology of Groups*. Oxford: Oxford University Press.
- List C.** (2014). 'Three Kinds of Collective Attitudes.' *Erkenntnis* **79**(9), 1601–22.

- List C. and Pettit P.** (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. New York, NY: Oxford University Press.
- Mathiesen K.** (2011). 'Can Groups be Epistemic Agents?' In H.B. Schmid, D. Sirtes and M. Weber (eds), *Collective Epistemology*, pp. 23–44. New York, NY: Walter de Gruyter.
- McMahon C.** (2003). 'Two Modes of Collective Belief.' *Protosociology* **18/19**, 347–62.
- Meijers A.** (2003). 'Can Collective Intentionality be Individualized?' *American Journal of Economics and Sociology* **62**, 167–93.
- Palermos O.** (2020). 'Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism.' *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00258-9>.
- Pettit P.** (2003). 'Groups with Minds of their Own.' In F.F. Schmitt (ed.), *Socializing Metaphysics*, pp. 167–93. Lanham, MD: Rowman and Littlefield.
- Quinton A.** (1975). 'Social Objects.' *Proceedings of the Aristotelian Society* **75**, 1–27.
- Ritchie K.** (2013). 'What are Groups?' *Philosophical Studies* **166**(2), 257–72.
- Ritchie K.** (2015). 'The Metaphysics of Social Groups.' *Philosophy Compass* **10**(5), 310–21.
- Ritchie K.** (2020). 'Social Structures and the Ontology of Social Groups.' *Philosophy and Phenomenological Research* **100**(2), 402–24.
- Rose D. and Schaffer J.** (2013). 'Knowledge Entails Dispositional Belief.' *Philosophical Studies* **166**, 19–50.
- Schmitt F.F.** (ed.) (1994). 'The Justification of Group Beliefs.' In *Socializing Epistemology: The Social Dimensions of Knowledge*, pp. 257–87. Lanham, MD: Rowman and Littlefield.
- Schwengerer L.** (2021). 'Defending Joint Acceptance Accounts of Justification.' *Episteme*. doi: 10.1017/epi.2020.55.
- Searle J.** (1990). 'Collective Intentions and Actions.' In P. Cohen, J. Morgan and M.E. Pollack (eds), *Intentions in Communication*. Cambridge, MA: MIT Press.
- Searle J.** (1995). *The Construction of Social Reality*. New York, NY: The Free Press.
- Shapiro S.** (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford: Oxford University Press.
- Tollefsen D.** (2009). 'WIKIPEDIA and the Epistemology of Testimony.' *Episteme* **6**(1), 8–24.
- Tollefsen D.** (2011). 'Groups as Rational Sources.' In H.B. Schmid, D. Sirtes and M. Weber (eds), *Collective Epistemology*, pp. 11–22. New York, NY: Walter de Gruyter.
- Tuomela R.** (1992). 'Group Beliefs.' *Synthese* **91**, 285–31.
- Tuomela R.** (2004). 'We-Intention Revisited.' *Philosophical Studies* **125**, 327–69.
- Tuomela R.** (2007). *The Philosophy of Sociality: The Shared Point of View*. New York, NY: Oxford University Press.
- Tuomela R.** (2011). 'An Account of Group Knowledge.' In H.B. Schmid, D. Sirtes and M. Weber (eds), *Collective Epistemology*, pp. 75–118. New York, NY: Walter de Gruyter.
- Tuomela R.** (2013). *Social Ontology: Collective Intentionality and Group Agents*. New York, NY: Oxford University Press.
- Uzquiano G.** (2004). 'The Supreme Court and the Supreme Court Justices: A Metaphysical Puzzle.' *Noûs* **38**, 135–53.
- Vanderschraaf P. and Sillari G.** (2005). 'Common Knowledge.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2005 Edition). <http://plato.stanford.edu/archives/win2005/entries/common-knowledge/>.
- Williamson T.** (2005). 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge.' *Philosophical Quarterly* **55**, 213–35.
- Williamson T.** (2017). 'Acting on Knowledge.' In J.A. Carter, E. Gordon and B. Jarvis (eds), *Knowledge First*, pp. 163–81. Oxford: Oxford University Press.
- Wray L.B.** (2007). 'Who has Scientific Knowledge?' *Social Epistemology* **21**(3), 337–47.

Jesper Kallestrup is Professor of Philosophy at the University of Aberdeen.