

Revised 1 July 2020. For M. Alfano, J. De Ridder & C. Klein (Eds.) *Social Virtue Epistemology* (Routledge, *forthcoming*)

# Collective (Telic) Virtue Epistemology

J. Adam Carter

*University of Glasgow*

[adam.carter@glasgow.ac.uk](mailto:adam.carter@glasgow.ac.uk)

*Abstract.* A new way to transpose the virtue epistemologist's 'knowledge = apt belief' template to the collective level, as a thesis about *group knowledge*, is developed. In particular, it is shown how specifically *judgmental belief* can be realised at the collective level in a way that is structurally analogous, on a telic theory of epistemic normativity (e.g., Sosa 2020), to how it is realised at the individual level—viz., through a (collective) intentional attempt to get it right *aptly* (whether *p*) by alethically affirming that *p*. An advantage of the proposal developed is that it is shown to be compatible with competing views—viz., joint acceptance accounts and social-distributive accounts—of how group members must interact in order to materially realise a group belief. I conclude by showing how the proposed judgment-focused collective (telic) virtue epistemology has important advantages over a rival version of collective virtue epistemology defended in recent work by Jesper Kallestrup (2016).

## 1. INTRODUCTION

A familiar theory of individual propositional knowledge holds that propositional knowledge is type identical with *apt* belief. A belief is apt if and only if it is successful (i.e., accurate) because competent.<sup>1</sup>

<sup>1</sup>The first use of the term 'apt' in the theory of knowledge to designate (roughly) 'success because of ability' is due to Sosa (2007); Sosa's 2007 view is a development on his earlier bi-level virtue epistemology (Sosa 1991) which predated his defence of epistemic normativity as performance normativity. Other pioneers of the 'success because of ability' model of knowledge include Zagzebski (1996) and Greco (2003).

When suitably fleshed out, this view has a lot of explanatory power in individual epistemology. It can, among other things, help us navigate the Gettier problematic<sup>2</sup>, the Pyrrhonian problematic<sup>3</sup>, radical sceptical challenges<sup>4</sup>, the value problem<sup>5</sup>, and more recently, epistemological problems related to the suspension of judgment.<sup>6</sup>

What's much less clear—and almost entirely unexplored<sup>7</sup>—is whether the 'knowledge = apt belief' (K=AB) template view is applicable only in individual epistemology, or whether some version of it can be made to work in *collective epistemology* as well, as a thesis about *group knowledge*.

One way in which things get messy here is as follows. The (K=AB) view is able to get all of the pleasing results above only when it is understood against a particular background view of epistemic normativity as *telic normativity*—viz., the normativity of attempts *as* attempts.<sup>8</sup> According to this background view, X-attempts can be evaluated along three dimensions: for success (did it the attempt succeed in attaining X), for adroitness (did the attempt manifest a competence to attain X reliably enough), and for aptness (was the attempt successful because adroit?). *Believing* is a kind of attempt, one that aims constitutively at getting it right (i.e., at truth<sup>9</sup>). Accordingly, the (K=AB) thesis goes hand in hand with the idea that only those beliefs that enjoy a certain kind of normative assessment

<sup>2</sup>Important to this proposal's strategy for navigating the Gettier problematic (e.g., Sosa 2010a)—given that apt belief can be unsafe—is the *bi-levelled* distinction between animal and reflective knowledge, where the latter is (in short) apt belief aptly noted (e.g., Sosa 2007). Versions of view that knowledge is equivalent to apt belief have also been defended, outwith the context of a bi-levelled view, as a solution to the Gettier problem, by Greco (2009, 2010), Turri (2011), Carter (2016a), Littlejohn (2014), Broncano-Berrocal (2018), and Zagzebski (1996). For criticism of the 'knowledge=apt belief' view, conceived of as a univocal (i.e., non-bi-levelled) proposal aimed at resolving the Gettier problem, see, e.g., Pritchard (2015, 2012), Kallestrup and Pritchard (2012, 2013), Lackey (2007), and Kelp (2013).

<sup>3</sup>See, in particular, Sosa (1997a, 2009, Ch. 2). Cf., Carter (2020b).

<sup>4</sup>Defences of performance-theoretic virtue epistemology in response to sceptical challenges are found in Sosa (2007 Ch. 2), Sosa (1999), Sosa (1997b), and Carter (2020c).

<sup>5</sup>E.g., Sosa (2010b).

<sup>6</sup>See here, in particular, Sosa (2020 Chs. 3–6).

<sup>7</sup>One exception is Kallestrup (2016), whose work will be discussed here.

<sup>8</sup>See Sosa (2020, 20–23). A forthcoming articulation of this view, due to Sosa (2020), uses the terminology 'telic' normativity to supplant what was previously (Sosa 2015) termed 'performance' normativity. In some places, where substantive differences introduced in the new proposal are not merely terminological, I will note this.

<sup>9</sup>Though, for some kinds of beliefs (e.g., judgments), not *merely* at truth. This point is taken up in §4.

*qua* attempt—the status of being not merely successful and adroit, but also *apt*—qualify as known.

Already, though, with above qualifications in play, there is a serious barrier to getting any kind of ‘knowledge = apt belief’ view off the ground in collective epistemology, as a view of collective knowledge. First, as recent literature in the epistemology of groups suggests<sup>10</sup>, even if groups can possess *knowledge*, it is much less clear that groups can have *beliefs*, and thus, that they can make the *kind of attempt* that, on the (K=AB) view, aspires to knowledge.<sup>11</sup> Secondly, even those collective epistemologists who *do* countenance collective beliefs often (though not always) take such collective beliefs to be a function of *joint acceptance* of a proposition—viz., whereby a group jointly accepts a proposition,  $\langle p \rangle$ , if and only if its members commit to acting as if  $\langle p \rangle$  is true in their capacity as group members.<sup>12</sup>

But here is the fly in the ointment: individual beliefs and collective joint acceptances seem, *prima facie*, like very different *kinds of attempts*. The latter, after all, involves intentional action in a way that is broadly analogous to how individual-level acceptances are intentional.<sup>13</sup> But, individual-level acceptances are not governed by a truth aim. Accepting a proposition, in individual epistemology, is not unsuccessful if untrue<sup>14</sup>; put another way, acceptance is not a constitutive attempt at truth. The problem now sharpens: if individual belief and collective joint acceptance are not normatively constrained by the same aim, then (K=AB) will not be extensionally adequate at the collective level as a theory of group knowledge, even if it delivers the goods at the individual level.

<sup>10</sup>See, e.g., Hakli (2006), Meijers (1999, 2003), Wray (2001). For an overview, see Simion, Carter, and Kelp (2020).

<sup>11</sup>One complication here concerns the extent to which ‘belief’ ascriptions track ‘consciousness’ ascriptions. For some experimental discussion here, see Knobe and Prinz (2008).

<sup>12</sup>For proponents of this kind of view, see, e.g., Gilbert (1987, 2013), Tuomela (1995), De Ridder (2018), Tollefsen (2015).

<sup>13</sup>For discussion on this point, see Alston (2007), Cohen (1992), Van Fraassen (1980), and Schwitzgebel (2019). For a notable argument that belief and acceptance can come apart in the other direction—viz., where one believes a proposition but does not accept it—see Bratman (1999).

<sup>14</sup>See Van Fraassen (1980) and Elgin (2017) for perspectives on the permissible acceptance of untrue propositions in the philosophy of science.

## 2. (K=AB) AND SUBSTANTIVE SYMMETRY

In light of the above, let's consider now the salient option space that the (K=AB) proponent at the individual level has for doing collective epistemology.<sup>15</sup>

View	Individual view / Collective View
<i>Collective Knowledge Scepticism</i>	K=AB (Non-sceptical) / K=AB (Sceptical)
<i>Non-Sceptical Hybrid Traditionalism</i>	K=AB (Non-sceptical) / not K=AB and not K-First (Non-sceptical)
<i>Non-Sceptical Hybrid Knowledge-First</i>	K=AB (Non-sceptical) / not K=AB and K-First (Non-Sceptical)
<i>Non-Sceptical Symmetric View</i>	K=AB (Non-sceptical) / K=AB (Non-sceptical)

Each of these four options comes with some substantial *prima facie* costs, though some to a greater extent than others. Option 1, *collective knowledge scepticism*, submits that if there is collective knowledge, (K=AB) would be a correct theory of that knowledge; however, since 'AB' is on this view held to be not realisable at the collective level, there simply is no collective knowledge. *Collective knowledge scepticism* effectively throws in the towel in collective epistemology, relegating knowledge (and thus epistemology) entirely to the individual arena. The cost here is significant, given, for instance, the prevailing view in collective epistemology that there are at least some *bona fide* cases of distributed knowledge, viz., cases where a group knows something that is not reducible to any proposition known by any of the individual members of the group.<sup>16</sup>

Option 2, the *non-sceptical hybrid traditional view*, cobbles together the (K=AB) view at the individual level with some other reductive—viz., *non*-'knowledge-first'—theory of knowledge at the collective level (e.g., perhaps collective knowledge = justified, true, collective acceptance), and maintains that the conditions specified by this other theory are (ordinarily enough) collectively realised. This

<sup>15</sup>Note that I'm not considering all possible combinations here that pair (K=AB) with a position type at the collective level. For example, combinations that result from pairing the (K=AB) individual view with a sceptical knowledge-first view or a sceptical traditionalist (JTB-style) view at the collective level are probably safe to set aside as unmotivated.

<sup>16</sup>For some representative discussion on this point, see, e.g., Brady and Fricker (2016), Palermos (2020), Bird (2010), Tollefsen (2015), Lackey (2014), Carter (2015), Hutchins (1995).

view seems to do better than *collective knowledge scepticism* in one sense but much worse in another. It does better in the sense that it does not throw in the towel at the collective level. But the cost of the non-sceptical result here is foregoing an important theoretical desiderata, *substantive symmetry*, which we can define as follows:

**Substantive symmetry:** A theory of individual and collective knowledge, *T*, is *substantively symmetrical* if and only if *T* posits a description of knowledge conditions at the collective level that matches *T*'s description of knowledge conditions at the individual level; otherwise, *T* is substantively asymmetrical (viz., 'hybrid').

Notice that Option 2, the *non-sceptical hybrid traditionalist view* is clearly substantively asymmetrical. This substantive asymmetry across the individual/collective divide invites a metatheoretical objection: when a view is substantively *asymmetrical*, this substantive asymmetry is (defeasible) evidence that the view is identifying something other than knowledge at at least one of the two levels. And this general metatheoretical objection (which requires further explanation on the part of the view to address it) gains additional traction once attention is drawn to the *prima facie* putative differences between individual and collective believing, and in particular, to the intentional character of the latter and the non-intentional character of the former.

Option 3, the *non-sceptical hybrid knowledge-first view*, avoids a sceptical implication in collective epistemology. That's good. And it also sidesteps entirely complications that arise for views—like Option 2—that attempt to vindicate collective knowledge by showing how it can be dismantled into a collective belief (or acceptance) condition plus other epistemic conditions that are satisfied if and only if the collective has knowledge. In recent co-authored work<sup>17</sup>, I've argued that—as a view of collective knowledge—a K-first view<sup>18</sup> has much to recommend it over 'justified true collective belief' and 'justified true collective acceptance' accounts of collective knowledge.<sup>19</sup>

There remains, though, the lingering issue of substantive symmetry. Going 'K-first' at the collective level achieves substantive symmetry in individual

<sup>17</sup>See Simion, Carter, and Kelp (2020).

<sup>18</sup>Such a view would be a kind of collective analogue to the kind of view championed by Williamson (2000) in individual epistemology.

<sup>19</sup>For examples of the former, see, e.g., Bird (2010), Palermos (2020), and De Ridder (2014). For an example of the latter, see Hakli (2006).

and collective epistemology *only if* it is paired with a K-first view at individual level. My co-authors with whom I've recently written on behalf of K-first collective epistemology—Christoph Kelp and Mona Simion—are happy to embrace knowledge-first epistemology at the individual level.<sup>20</sup> What this means is that by going K-first at the collective level, they maintain substantive symmetry, whereas I would maintain this only by then jettisoning (K=AB) at the individual level. As the reader will have gleaned from §1, I think it is right to be impressed with what (K=AB) can do at the individual level<sup>21</sup>, and more so than with what K-first can do.

For those epistemologists who like me favour (K=AB) at the individual level, then, the prospects don't look initially very good to get 'everything we want' at the collective level, including substantive symmetry. *Unless* of course, there is some way to make Option 4, the *non-sceptical symmetric view* work, *despite* all of the worries noted in §1 that seem to stand in the way of defending (K=AB) at the collective level.

In the remainder of this paper, I am going to argue that—despite things looking very bleak initially—Option 4 really *is* defensible. (K=AB) can be made to work not only as a theory of individual knowledge but *also* as a theory of collective knowledge. In order to see how the view works, it will be important to distinguish, following Ernest Sosa (2015, 2020), two importantly different kinds of beliefs: mere *alethic affirmations* and *judgmental beliefs*. While there is (perhaps) no collective analogue to individual alethic affirmations<sup>22</sup>, there is a collective analogue to individual judgmental beliefs. And—here is the second part of the argument—the collective analogue of individual judgmental beliefs is performatively analogous with individual level judgmental beliefs, and in this respect, it is analogously knowledge-apt. By drawing a parallel between individual judgment and collective judgment, we can see how (K=AB) is not merely a serious option for individual epistemology but for collective epistemology as well.

<sup>20</sup>See, e.g., Kelp (2018) and Simion (2019).

<sup>21</sup>For a recent defence of a new variation of this view in individual epistemology, see Carter (2020a).

<sup>22</sup>To be more precise, there is plausibly no collective analogue to at least one species of alethic affirmation, namely, functional alethic affirmation. See fn. 20.

### 3. GRADES OF KNOWLEDGE

Following Sosa, the (K=AB) slogan can be unpacked in individual epistemology in different ways—corresponding with different *grades of knowledge*—depending on the kind of *attempt* at getting it right that the relevant belief is.

Suppose, for example, that you are taking an eye exam, and you begin losing confidence as you get closer to the bottom row. But you read out the bottom row anyway. It turns out your lack of confidence on those bottom-row letters wasn't warranted, as you were actually perfectly reliable at the bottom row, despite the shaky confidence.<sup>23</sup>

When you affirm a given letter,  $\langle p \rangle$ , on the bottom row, it is an attempt to get it right (on whether  $\langle p \rangle$ ) by affirming that  $\langle p \rangle$ —viz., it is, in Sosa's terminology, a kind of *alethic affirmation*. In the above case, your alethic affirmation that  $\langle p \rangle$  is adroit (more so than you recognise); and even more, your alethic affirmation that  $\langle p \rangle$  is *apt*. As such, it constitutes a kind of subcredal *animal knowledge*, where animal knowledge is type-identical with *apt alethic affirmation*.<sup>24</sup>

So: 'animal knowledge = apt alethic affirmation' represents one kind of substantive gloss of the (K=AB) template. But, importantly, not *all* beliefs are *mere* alethic affirmations, viz., attempts to simply get it right by affirming. Some beliefs—*judgmental beliefs*—are attempts to get it right *aptly* by *alethically affirming*. In a bit more detail for now (see §4 for further elaboration): in judging something to be so, a thinker aims intentionally to get it *aptly* right (by alethically affirming that  $\langle p \rangle$ ) on a given question. *Apt judgment*—i.e., when one's attempt to get it right aptly by alethically affirming is *itself* apt—is knowledge. But it's not merely the kind of knowledge you get when a mere alethic affirmation is apt. Apt judgment is—on the telic virtue epistemologist's framework—knowledge *full well* (alternatively: judgmental knowledge).

<sup>23</sup>Though, not a *sheer* guess, as one might take when simply flipping a coin and affirming, before the coin drops, 'heads' or 'tails'. For discussion, see Sosa (2015 Ch. 3) and Carter (2016b).

<sup>24</sup>Note that while some alethic affirmations, *qua* attempts, are consciously intentional, others are what Sosa calls functionally teleological (2020, 50, fn. 27). On this latter point, consider that right now, for example, your belief—sitting bored in a room—that a clock is ticking, and that the light is on). As Sosa puts it, such 'functional' beliefs: '[...] guide everyday action entirely below the surface of consciousness. Unlike judgmental beliefs, these are not accessed through a simple conscious response to the relevant "whether" question' (2015, 67). But they have nonetheless a teleological aim *qua* functional representation. As such, such functional beliefs are subject to assessment within the kind of 'triple A' framework (accurate/adroit/apt) characteristic of telic virtue epistemology. For an extended discussion on this point, see Sosa (2015 Ch. 3, fn. 5).

In sum, then, the (K=AB) template account of knowledge can be glossed—within telic virtue epistemology—in the following two ways, which correspond with two different kinds of ‘attempts’ to get it right through affirmation.

<i>What kind of attempt?</i>	<i>What does it constitutively aim at?</i>	<i>What is it when apt?</i>
alethic affirmation	to get it right (whether $\langle p \rangle$ ) by affirming that $\langle p \rangle$	<b>animal knowledge</b>
judgment	to get it right (whether $\langle p \rangle$ ) <i>aptly</i> by alethically affirming that $\langle p \rangle$	<b>judgmental knowledge</b> (i.e., knowledge full well)

As the reader might have anticipated, judgmental knowledge is of particular interest for the telic virtue epistemologist who wants to embrace a (K=AB) view at both the individual and collective level, and thus, to retain substantive symmetry at the individual and collective levels. And this is because judgment, *qua* attempt, involves an intention—where the intention here is to get it right *aptly* by alethically affirming. Intention, recall, seemed *ex ante* to be a property of collective belief (or collective acceptance) not shared with individual belief. Even more, and unlike mere acceptance, the constitutive aim of judgment is such that, when aptly attained, what results is a kind of knowledge.

The take away point, then, seems to be the following: a proponent of (K=AB) at the individual level has a real shot at vindicating a *non-sceptical symmetric view* (§2) despite what looked like initial barriers. The key will be to set aside the collective analogue of mere alethic affirmation and to instead zero in on the collective analogue of individual intentional judgment.

Here is the plan for the remainder of the paper. §4 will unpack some of the key features (glossed over so far) of judgmental belief at the individual level, and in doing so clarify the sense in which judgmental belief is a species of intentional action on the telic VE framework. §§5-6 then transpose individual judgment and judgmental knowledge (and its key telic normative features) to the collective level, establishing a proof of concept of how the (K=AB) template can be vindicated as a theory of (judgmental) knowledge at the collective level, and one whose conditions are ordinarily enough met so as to be a non-sceptical theory. §7 then shows that the view has important advantages over a rival version of collective virtue



epistemology defended in recent work by Jesper Kallestrup (2016).

#### 4. JUDGMENT

On telic virtue epistemology, a judgment that  $\langle p \rangle$  is an intentional constitutive attempt to get it right aptly by alethically affirming that  $\langle p \rangle$ . In order to bring in to view what a collective analogue of an *apt* judgment would be, let's first sharpen some of the components of individual judgment itself, by briefly answering some key 'FAQs':

- (a) In what sense is a judgment an *attempt*?

There are two ways one might attempt to attain an objective, corresponding with a distinction between *instrumental attempts* and *constitutive attempts*.<sup>25</sup> In the former case, one makes an attempt, at an objective,  $O$ , by implementing means that are both preliminary and viewed as such. Inquirers, for example, might instrumentally attempt to know whether there is a chaffinch in the garden by implementing the preliminary means of finding a good vantage point from which to spot the bird. *Constitutive attempts* are different. In the case of a constitutive attempt at an objective  $O$ , one implements means aimed at  $O$ , but not means which are regarded as preliminary and viewed as such. Rather, one—in making a constitutive attempt at attaining  $O$ —implements means that are aimed at grounding one's success in attaining  $O$ . (In an athletic case, compare: instrumentally attempting to hit a hole-in-one by practicing hard for months, and constitutively attempting to hit a hole-in-one by swinging the club.)

The idea that a judgment is a constitutive attempt to get it right whether  $\langle p \rangle$  aptly by alethically affirming that  $\langle p \rangle$  registers this constitutive rather than instrumental character of the attempt.

- (b) So judgments are constitutive attempts. But if the idea is that they are constitutive attempts at *apt* alethic affirmation (and that they are not merely constitutive attempts at *successful* alethic affirmation), then does that mean that judgments are attempts to do more than just 'get it right' (whether  $\langle p \rangle$ ) by affirming' that  $\langle p \rangle$ ?

<sup>25</sup>See Sosa (2020, 24–26).

Yes, that's right. And a helpful way to think about this will be to consider two ways you might constitutively attempt to make a basketball shot. You might on the one hand shoot in the endeavour of *making* it by shooting it. On the other hand, you might shoot in the endeavour of *making it aptly* by shooting it. A reckless shot from too far out—viz., beyond your threshold for sufficient reliability—that happens to go in is successful relative to the first kind of attempt. It is not successful relative to the latter kind of attempt, even though it goes in. The basketball coach will advise a player to take only the second kind of shot (unless the clock is running out). The second kind of shot, which aims not just at success, but at aptness, is analogous to a judgment, which aims not just at getting it right, but at getting it right aptly.<sup>26</sup>

- (c) Even if all of the above about judgment is granted, there remains an elephant in the room. Why should we think that the kind of constitutive attempt at apt alethic affirmation that judgment is should be understood as a constitutive *and intentional* attempt at attaining this aim? After all, judgments, like many other beliefs we have, surely aren't voluntary?

To say that judgment is a (constitutive) attempt, with *intention*, to attain a given aim (to wit, the aim of apt alethic affirmation, or animal knowledge) does not imply that how we judge is thereby under the sort of voluntary control whereby we could judge directly through arbitrary choice.<sup>27</sup>

In order to see why, let's consider what suffices for the kind of intentional action that judgment is. To do this, we can distinguish between a basic action and a 'simple' intentional action. Suppose I intentionally move a finger. Or think of a triangle. These are both deeds I do intentionally, but in each case, there's no other deed I do in the endeavour to do these things, either the physical basic action of moving a finger or the mental basic action of thinking of a triangle.

Affirming that  $\langle p \rangle$  is a basic action. When one affirms that  $\langle p \rangle$ , one doesn't do this partly by doing something else in the endeavour to affirm that  $\langle p \rangle$ . Basic action can be distinguished from 'simple' intentional action, where an agent aims to perform a deed (at  $t$ ) at least partly by performing a basic action,  $B$ , at  $t$ .<sup>28</sup>

<sup>26</sup>See also Sosa (2015 Ch. 3) for a helpful development of this idea. For criticism see Schechter (2019).

<sup>27</sup>For a discussion on this point, see Sosa (2020, 29) and also (2015, 166–7).

<sup>28</sup>(Sosa 2015, 166).

Whereas affirming is a basic action, judgment is not; it is a simple intentional action. Here's Sosa (2015):

In a judgment, the agent affirms in the endeavor to (thereby) affirm aptly. If the agent does attain that objective, then, we have the following structure: the agent affirms aptly that *p* at *t* partly-by affirming that *p* at *t* (2015, 166).

The above sense in which judgment is a kind of (simple) intentional action, thus, doesn't imply at all that judging is voluntary in the sense that it could be reversed arbitrarily—viz., the sense in which researchers widely deny that belief, more generally, is voluntary.<sup>29</sup> Moreover, given that not all beliefs are judgmental, not all beliefs are intentional even in the above sense that does not imply voluntary control.

- (d) But if judgment is intentional even in the sense described, then couldn't it potentially be redirected towards practical ends in a way that is analogous to how individually accepting a proposition can be practically aimed? If the answer is 'yes', then wouldn't 'apt judgment' be a candidate for knowledge if and only if apt acceptance is a candidate for knowledge?

Careful! There's a sense in which judgments are beholden to practical factors: we might let practical values dictate which inquiries we take up in the first place—the kind of normativity here (whereby our doing so is better or worse) is the (broader) normativity of intellectual ethics.<sup>30</sup> For example, the lawyer—in the course of considering the case of a potential client—might make various judgments about how the law applies in that particular case, all broadly for the sake of a moral objective: to provide assistance to a vulnerable client. But, in making these judgments, the lawyer is nonetheless making a constitutive attempt at knowing with one's alethic affirmation. (Furthermore, the telic assessment of a judgment—the kind of assessment that matters for whether the judgment is knowledge—is not, as Sosa (2020) puts it, 'properly affected by extraneous objectives that the agent may also be pursuing through the same means.'<sup>31</sup>

<sup>29</sup>For additional discussion, see Sosa (2020, 98, fn. 54).

<sup>30</sup>See, for discussion, Sosa (2007, 89) and especially Sosa (2020 Ch. 2).

<sup>31</sup>(Sosa 2020, 36). As a point of clarification, Sosa adds here that 'extraneous aims may define attempts of their own, and these would be properly assessed by reference to those aims. But

The above kind of case is very different from constitutively attempting to do something *else* (other than to get it right aptly) by affirming, even when that something else is epistemic. For example, a contestant on a game show who affirms that  $\langle p \rangle$  in the endeavour to get it right whether  $\langle p \rangle$ , by affirming, fully cognisant *that* they are guessing, is not thereby judging. They are merely alethically affirming, making a constitutive attempt to get it *right* (not: aptly right) by affirming. Likewise, one is not judging if one is engaging in wishful thinking—affirming what one merely hopes is the case, or, when affirming through speech, what one merely hopes others believe that one believes is the case.

## 5. COLLECTIVE JUDGMENT

With the key contours of judgment, at the individual level, in mind—let's now transpose things to the collective level. As was noted in §1, it is contentious in collective epistemology, and more widely in collective intentionality, whether groups can have genuine *beliefs*, or whether they can merely *accept* propositions.

A quick clarification is in order. There is no dispute whatsoever that groups can have beliefs in an uncontentious *summativist* sense. A group summatively believes a proposition,  $\langle p \rangle$ , if and only if most or all of its members believes that  $\langle p \rangle$ . For example: the group that is 'Swedes' believes the proposition  $\langle \text{Volvos are safe} \rangle$  in a summative sense if and only if all or most individual Swedes believe this.

The more contentious, and more philosophically interesting, notion of group belief is a non-summativist, or alternatively an 'inflationist', notion of group belief, according to which it's false that the group belief is reducible to an aggregate of the individual beliefs of group members.

As Alexander Bird (2019) notes, the rationale that has persuaded those collective epistemologists who have embraced an inflationist view of group belief is that the conditions specified by the summativist account seem in certain cases neither sufficient nor necessary for a group to believe something. Here is an example that challenges sufficiency. Suppose all members of a group (say, a city council) believe  $\langle p \rangle$  but, in their capacity as council members behave as though they do not believe  $\langle p \rangle$  (and, where each member thinks they are the only one to

---

their assessment would remain irrelevant to gnoseology, the part of epistemology that concerns assessment of judgments and beliefs that are truth-directed, and more specifically knowledge-directed (as has been postulated by telic virtue epistemology....)' (2020, 36). For some related discussion, see also (2015, Ch. 9).

have this ‘strange’ belief that  $\langle p \rangle$ .) It seems like, in this case, it would be a mistake to say the city council believes that  $\langle p \rangle$  even though all its members do.<sup>32</sup> Likewise, regarding necessity—suppose a group of jurors is evaluating the guilt of an immigrant whose ethnicity is different from their own.<sup>33</sup> The evidence for innocence is overwhelming, and the jurors—each of whom is privately racist and believes the defendant is guilty—recognise this, and publicly affirm in line with the evidence (rather than in line with their racist beliefs) by voting for his innocence. Here, it looks like the *jury* believes the client is innocent, even though its individual members do not. Likewise, cases of distributed cognition speak against necessity: a group might come to endorse a viewpoint, as a group, by way of a division of cognitive labour dispersed across group members, who have different cognitive roles and then ‘feed’ different bits of information to a centralised database that collates the information.<sup>34</sup>

As was noted in §1, the matter of whether we ought, in doing epistemology, to regard these kinds of cases as ones that feature (non-summativist) group *belief* has been deeply divisive. The divisiveness has centred around two claimed disanalogies between belief on the one hand, and so-called inflationary group ‘belief’ on the other. These disanalogies concern (i) automaticity; and (ii) involuntariness. Regarding automaticity: As Hakli (2006) reasons: beliefs are paradigmatically formed in an automatic and involuntary manner (consider your belief ‘There’s a knock at the door!’ which you form automatically after hearing a knock at the door), while whatever mental states groups are capable of hosting are not automatic, but rather, the result of careful deliberation. And, regarding voluntariness: beliefs seem to be *essentially* involuntary; voluntary affirmations are ‘make beliefs’, not genuine beliefs. But, group inflationary ‘belief’ comes apart from ordinary individual-level belief in both of these ways: group ‘beliefs’ are never automatic, and always voluntary. So, we should reject that group ‘beliefs’ are beliefs.

The above line of argument against group inflationary beliefs unhelpfully runs together a number of things that should be kept apart. By being more careful here, we can make some concessions to the rejectionist about group belief while upholding a version of group belief inflationism, one on which we can unproblematically attribute judgments to groups.

<sup>32</sup>See Gilbert (1992, 257–8); cf., Bird (2019, 3).

<sup>33</sup>This is a variation on a case from Kallestrup (2016), though for other versions, see also Carter (2015, 718, fn. 8), Gilbert (1992, 187), Bird (2019, 3–4), and Simion, Carter, and Kelp (2020).

<sup>34</sup>See Bird (2019, 3); for related discussion of distributed beliefs attained by aggregating information into a central database, see Lackey (2018).

A first concession to the rejectionist is that anything that is a belief (whether the host of the belief be an individual or a group) is essentially non-voluntary.<sup>35</sup> A second concession is that at least some individual beliefs are clearly automatic, and, *if* there are group inflationary beliefs, such group beliefs would never be automatic.<sup>36</sup> What follows from these concessions, about group beliefs, is just that: groups can't have beliefs that are either voluntary or automatic.

Does the postulation of a group *judgment* violate either of these conditions we're conceding to the rejectionist? The answer, I want to suggest, is 'no'. Just consider the following minimal statement of what a group must *do* in order to make a judgment, as construed within a telic virtue epistemology:

**(Non-summativist) collective judgment:** A group *G* judges that  $\langle p \rangle$  if and only if the *G* constitutively attempts, with intention, to get it right (whether  $\langle p \rangle$ ) aptly by alethically affirming that  $\langle p \rangle$ .

Such a collective judgment will not be 'automatic'; it is (like an individual judgment) an intentional action. But does it *matter* that we've conceded to the rejectionist that *no* collective beliefs can be, like at least some individual beliefs, automatic? No; a collective being unable to produce automatic beliefs, as individuals can, is of course compatible with it being able to make judgments. For a thought experiment: imagine an individual who, operated on by neuroscientists, was unable to form beliefs in any other way than by intentional judgment. This individual would be unusual, disadvantaged even. But this inability does not call in to doubt the individual's capacity to make judgments, not in the least. Such an individual would therefore hardly be incapable of having beliefs, and even knowledge (whenever these judgments are apt), and *mutatis mutandis*, for groups—analogously unable to believe automatically while capable of judging.

Likewise, the concession to the rejectionist that beliefs can't be voluntary is not problematic for the prospects of group judgment. It would be if and only if group judgments are voluntary in virtue of being intentional. But the kind of intention characteristic of judgment at the individual level was already shown not to imply a kind of voluntariness incompatible with any sort of knowledge-apt belief. The same holds at the collective level, when a group (collectively) makes a constitutive attempt, with intention, to get it right aptly by alethically affirming. Granted, a

<sup>35</sup>For a detailed argument on this point, see Hieronymi (2009).

<sup>36</sup>To be clear, I'm just granting this here to the rejectionist for the sake of argument. There are actually some reasons, which I won't get in to here, for thinking that this should not be granted. See, e.g., Simion, Carter, and Kelp (2020) for discussion.

group could very well—for example, by jointly committing to a proposition—make a (collective) attempt to get it right (as a group) in the endeavour to do something *else*—for example, to let a proposition stand as the group’s view if and only if doing so would be strategically wise from a marketing perspective. (For example: imagine Philip Morris’ board jointly committing to the proposition that deaths from cancer are unfortunate.) Even when such propositions are true—and indeed, even if the individuals of the group know this proposition to be true, individually—the collective attitude would not be a judgment, as the collective aim is something other than getting it right aptly. Likewise—and analogously at the individual level—a group could (if competing in a game show) jointly commit to an arbitrary answer to a game show question, and in doing so, jointly accept a proposition by merely (jointly) alethically affirming the proposition, though without judging.

Let’s briefly take stock. What the foregoing shows is that the basic ‘template’ for an individual judgment, on telic virtue epistemology, can be transferred to the collective level as a view of (inflationary) group judgment without falling foul of the kinds of objections that are so often raised to so-called group beliefs. And this is the case even if these objections (or some version of them) are applicable to some *non-judgmental* accounts of inflationary group belief.<sup>37</sup>

Moreover, the basic view—that a group judgment is a constitutive attempt, with intention, to get it right aptly via alethic affirmation—is, conveniently, compatible with different views of the metaphysical nature of group belief, that is, different views about what kind of mechanisms have to take place among group members to materially realise a group belief. Let’s briefly now consider how the template proposal could potentially be glossed differently on two leading proposals types: (i) Margaret Gilbert’s (1987, 2013) *joint commitment account*; and (ii) *social-distributed accounts*, a Durkheimian functionalist version of which is defended by Alexander Bird (2010, 2019) and a cognitive integrationist version of which is defended by S. Orestis Palermos (2016, 2020).

Consider first Gilbert’s *joint commitment* model.<sup>38</sup> On this view, it is necessary and sufficient for a group to believe a proposition,  $\langle p \rangle$ , that the group *jointly*

<sup>37</sup>This point applies, principally, to voluntariness-based objections. The reasons noted above in this section for why an ‘automaticity’-style objection to inflationary group belief would not threaten the prospects of group judgment (on a telic VE model) would, as far as I can tell, apply *mutatis mutandis* to other accounts of inflationary attitudes.

<sup>38</sup>While Gilbert is arguably the leading proponent of the view (in various forms over her career), other philosophers have also defended variations on this kind of proposal, including Tuomela (1995) and Tollefsen (2015).



*accepts* that  $\langle p \rangle$ . Further, the members of a group jointly accept that  $\langle p \rangle$  when the members conditionally commit to accept that  $\langle p \rangle$ , which they do if and only if each is committed to acting as if  $\langle p \rangle$  provided the others do. Thus, on this view, it is the obtaining of this joint commitment, by individual members of the group, to act as if  $\langle p \rangle$  is true provided others do, that gives rise to, or ‘realises’, the group belief.

Here is how a non-summativist account of collective judgment could be given a metaphysical gloss along the lines of the above Gilbert-style account, according to which joint commitments are the realisers of group belief. First, we assimilate a group’s joint acceptance that  $\langle p \rangle$  with (mere) affirmation that  $\langle p \rangle$ , such that a group affirms that  $\langle p \rangle$  if and only if the group members jointly accept that  $\langle p \rangle$ . An affirmation is an *alethic affirmation* if and only if it is an attempt to get it right (on whether  $\langle p \rangle$ ) that  $\langle p \rangle$ . Now, if the individuals in a group privately desire that the group jointly accept that  $\langle p \rangle$  just in case by doing so they would be (jointly) affirming truly, would *this* suffice to make the kind of attempt the group is making (when jointly committing that  $\langle p \rangle$ ) an *alethic* affirmation, and not only an affirmation?

The answer to this question is ‘no’. The best way to interpret the above scenario is as one of a *mere* collective affirmation, given that the group itself is not affirming in any endeavour (for truth, aptness or anything else) *as* a group, even if individuals are so endeavouring when jointly accepting (and in doing so, affirming) the proposition as a group member. Rather, a group *alethically affirms* a proposition only if the group has already *collectively* established (viz., through a prior joint commitment) the scope of their endeavour to be an attempt to get it right (on whether  $\langle p \rangle$ ). With this idea in hand, we can now characterise a group *judgment* (on the joint acceptance model) in terms of a group alethic affirmation. To a first approximation, a group judges that  $\langle p \rangle$  only if, antecedent to affirming whether  $\langle p \rangle$ , the group jointly commits to (i) *alethically* affirm whether  $\langle p \rangle$ ; (ii) to get it right  $\langle p \rangle$  *aptly* through (i). The combination of (i) and (ii) establish that the kind of attempt at getting it right the group makes when jointly accepting (i.e., affirming) that  $\langle p \rangle$  lines up with collective judgment, rather than, say, with *mere* collective affirmation or *mere* collective alethic affirmation.

Our (non-summativist) collective judgment template can just as easily be given a very different gloss, if paired with a different view about what realises group belief. Consider, for example, *social distributed* views of group belief, according to which cognition involves relatively tightly integrated groups working together, with scientific research teams being the classic example (e.g., Bird 2010; De Ridder 2014; Palermos 2016). On these views, it is in virtue of the social relations at



work between group members that different parts of the system contribute to the generation of the system's collective mental state that  $\langle p \rangle$ , and even if (though not only if) no individual in the group actually hosts the belief that  $\langle p \rangle$ .<sup>39</sup> For example, suppose a scientific research team—inquiring into whether there are over 1000 species of a certain kind of bird,  $B$ —divides up tasks with the plan of having each individual input their own data (on the basis of performing certain individual epistemic tasks corresponding with their roles in the group) in a centralised database, which then combines the data, spitting out a collective result in the affirmative if and only if the aggregated data compiled through the shared database identifies over 1000 species of  $B$ . On this kind of view, the *group* takes a representational stance on the matter of whether  $\langle p \rangle$  (whether the number is above or below 1000)—a belief according to social distributed views. And while the individual members may permissibly also hold the group belief at some point in the process (e.g., by consulting the database, conferring with each other, etc.) their doing so isn't what realises the group belief. Rather, it's the distributed contributions of the individuals to the collective result in accordance with their roles that is doing the work.

One internal dispute among proponents of social-distributed models of group belief concerns the matter of how to delineate who exactly should 'count' as part of the group that has the belief it has in virtue of the distributed epistemic contributions of its members. Consider—to use a case discussed by Simion, Carter, and Kelp (2020)<sup>40</sup>—the mailperson, whose job it is to bring the mail to the scientific research team—including some important research documents—and in doing so, causally contributes to the research team's group belief  $\langle p \rangle$ . Is the mailperson thereby a *member* of the group that believes that  $p$ ?<sup>41</sup> Intuitively, it seems the answer here should be 'no', despite the mailperson's making an epistemic contribution to the group belief in light of occupying a social role. Whereas a social-role functionalist account like Bird's has a difficult time explaining why the mailperson should be ruled out, S. Orestis Palermos's (e.g., 2016) cognitive integrationist version of a social-distributed account of group belief can do so easily. For Palermos, the relevant social interactions between group members that serve to collectively realise, via distributed cognitive tasks, a group belief, must include *feedback*

<sup>39</sup>For discussion, see Simion, Carter, and Kelp (2020, sec. 3).

<sup>40</sup>The original example of the mailperson is due to Mona Simion.

<sup>41</sup>As Mark Alfano has pointed out to me, a variation on this kind of philosophical issue is on display in cases of adjudicating CERN authorship, especially where research teams are especially when authorship is large. For an overview of CERN's practical guidelines for establishing authorship of large papers, see <http://library.cern/cern-author-guidelines>.

*loops*—viz., two-way causal interactions with other contributing group members, of the sort that are, on dynamical systems theory (DST), the mark of dynamical systems.<sup>42</sup> Because the causation in the case of the mailperson is asymmetrical (from the mailperson to the research team, but not vice versa), the mailperson is not a part of the group that believes despite the causal epistemic contribution made.

That said, the feedback loop requirement itself may be too strong, as it's not obvious that all *bona fide* members of, e.g., a scientific research team that produces a result must interact via feedback loops with each other.<sup>43</sup> For our purposes, we remain neutral about how a social-distributed model might 'thread the needle' to accommodate the above kind of dilemma. Rather, it suffices to register that an account of collective judgment could be glossed on a social distributed model, regardless of whether the relevant social relations are social-functional contributions (Bird) or reciprocal-causal contributions (Palermos). To a first approximation, the idea is as follows: A group  $G$  judges that  $\langle p \rangle$  if and only if the  $G$  constitutively attempts, through *distributed* individual contributions to the group attempt, to get it right (whether  $\langle p \rangle$ ) aptly by alethically affirming that  $\langle p \rangle$ .

In order to make this idea concrete, it will be helpful to compare a group judgment, glossed on the social-distributed model, with a *mere* group affirmation, also glossed on the model. For ease of reference, let's suppose we have two scientific research teams, 'Team Affirmation' and 'Team Judgment', each aiming to establish—to continue here with our previous example of distributed cognition—whether there are over 1000 species of bird  $B$ . Team Affirmation, low on research funding, is trying to cut corners. The interactions between group members reflect this—none of the individuals are accountable for the accuracy of their results (as long as they register a plausible enough looking number indicating how many species they have found of bird  $B$  in their designated sector). With this kind of epistemic laxity characterising the social norms in play, the group eventually affirms a collective answer (fewer than 1000 species)—and suppose even that this is true. Even so, we have here a mere (collective) affirmation, and (depending on

<sup>42</sup>See, e.g., Beer (1995, 2014).

<sup>43</sup>An example case features in Kallestrup (2016), which involves a single mathematician and a physicist who do not communicate with each other, but who both—according to an agreed plan—send their results to a third team member, who applies the modus ponens inference rule to generate the group position. See also Simion, Carter, and Kelp (2020) for explicit discussion of the idea that a feedback loop condition will be too strong such that it will have the result of ruling out individuals (who ought to be ruled in) as part of a group that believes something.

how the details are further filled out) at the very most, an *alethic* affirmation. But Team Affirmation falls short of making a judgment.

Team Judgment—true to their namesake—approaches things differently. The social norms governing the individual contributions to the collective output, as well as the interactions between individuals in ensuring reliability in reporting and collating the individual epistemic contributions, are knowledge directed. Suppose even that these norms are deeply internalised as well as explicit in the research team's manifesto: individual contributions to the group output are accepted only if reliable methods are used, and individual epistemic contributions are even cross-checked by other team members as a matter of policy to minimise epistemic risk. Suppose finally that the result is the same (fewer than 1000 species).

Given that both of our two research teams generate, through organised distributed efforts, a collective representational output (i.e., that there are fewer than 1000 species of bird *B*), both (on a social-distributed model) count as at least affirming this, perhaps even as both *alethically* affirming this. But only Team Judgment is actually aiming at *aptness*, and not merely at correctness, in a way that is characteristic of judgment. Team Judgment is affirming with the aim of (through the distributed intellectual contributions) getting it right *aptly*, viz., through not just any kind of way of organising and aggregating individual contributions, but—as the social norms governing their inquiry prescribe—through a reliable, knowledge-conducive way of doing so.

## 6. COLLECTIVE JUDGMENTAL KNOWLEDGE

Let's take stock. We've seen how judgmental belief can be realised in a promising way at the collective level such that it is structurally analogous, on a telic theory of epistemic normativity, to how it is realised at the individual level—viz., through a (collective) intentional attempt to get it right *aptly* (whether  $\langle p \rangle$ ) by *alethically* affirming that  $\langle p \rangle$ . Moreover, we've seen that an advantage of the proposal is that it is in principle compatible with competing views—viz., joint acceptance accounts and social-distributive accounts—of how group members must interact in order to materially realise a group belief.

But two residual questions remain.

*Question 1:* First, if a collective judgment is to result in collective judgmental *knowledge* (alternatively: in knowledge *full well*), the collective judgment must itself be *apt*. What, then, is required in order for a collective—as opposed to merely an individual—judgment to attain the status of aptness?

*Question 2:* How much collective judgmental knowledge *is* there? If none, or very little, then the proposed view (paired with a K=AB view at the individual level) secures substantive symmetry, but not full-blown *non-sceptical* substantive symmetry of sort desired.

Let's now answer these two questions in turn.

(a) *An answer to Question 1*

Any judgment—just like any (constitutive) attempt, more generally—is apt if and only if it is successful and the success is through *competence*. A  $\phi$ -competence, generally speaking, is a disposition of a subject to succeed reliably enough, whenever one makes a  $\phi$ -attempt and is in proper shape and properly situated. (In the case of performing a triple axel, for example: if you were unable to succeed when you attempt to perform the jump while drugged and strapped inside an airplane—away from an ice rink—this would *not* count against your competence to land a triple axel. What matters is whether you'd succeed reliably enough if you tried while in proper shape (undrugged) and properly situated for such an attempt (equipped with skates, on an ice rink, plenty of ambient oxygen, etc.).

'Successful' judgment is equivalent to *apt alethic affirmation*—viz., that which judgment, as such, is a constitutive attempt to bring about. And so: a *competent judgment* is a judgment that manifests a disposition (on the part of the judging subject) to (reliably enough) succeed at doing *that*—viz., that which constitutes successful judgment. An individual thinker, for example, would possess such a competence only if disposed to reliably enough affirm alethically only if she would do so aptly. In light of such a competence, the thinker *not easily* would fail to affirm with alethic aptness.

In sum, then: a judgment is *apt* just in case its success, that is, its securing the aim of apt alethic affirmation manifests (or: is because of) one's disposition to (reliably enough) affirm alethically only if one would do so aptly. We can now—drawing from our template account of collective judgment from §5—extend this

account of apt judgment to the special case of where the judgment is a collective judgment as follows:

**(Non-summativist) collective apt judgment (i.e., collective judgmental knowledge):** A group  $G$  judges that  $\langle p \rangle$  aptly if and only if  $G$  constitutively attempts, intentionally, to get it right (whether  $\langle p \rangle$ ) aptly by alethically affirming that  $\langle p \rangle$ ; (ii)  $G$  secures this aim; and (iii)  $G$ 's securing this aim manifests (or: is because of)  $G$ 's disposition to (reliably enough) collectively affirm alethically whether  $\langle p \rangle$  only if  $G$  would do so aptly.

Important here are two key points, the first of which has to do with clause (iii)—viz., the idea that  $G$  must have a disposition—one that it manifests in judging successfully—to (reliably enough) collectively affirm alethically only if  $G$  would do so aptly. This disposition, it should be emphasised, is a disposition of the judging *group*, one the group has in an inflationary sense, such that it's false that the group possesses the disposition if and only if individuals in the group possess this disposition.

The second idea concerns the 'because of' locution. Even if a group judgment is both *successful* (i.e., it results in apt alethic affirmation) and *competent*, it might still fall short of being *apt, qua* judgment, if the group's success is not *because* competent.<sup>44</sup> Suppose, for example, that a research team *competently* judges that  $\langle p \rangle$ , and further, that  $\langle p \rangle$  is true. However, here's the twist: due to a fluke error in some of the computer equipment that the group uses to take the measure of its members' individual contributions to the group judgment (whether it be a joint-commitment-registering machine, in a Gilbert-style model, or a data aggregating machine, on a social-distributed model), the machine first (i) incorrectly registers the group's view as not- $\langle p \rangle$ ; but, then, (ii) a second computer glitch occurs, causing the machine to issue an arbitrary result that happens to be  $\langle p \rangle$ . In this case, the group lacks judgmental knowledge, (and even apt alethic affirmation) as the fact that the group affirms what it does rather than something else is not due to any kind of competence but just to dumb luck.<sup>45</sup>

<sup>44</sup>Compare here with the 'double gust of wind'-style cases that are familiar in virtue reliabilist epistemology to illustrate aptness (e.g., Sosa 2007, 2010a; Pritchard 2012; Greco 2010, Chs. 5-6).

<sup>45</sup>Couched in slightly different language: the structure of this case involves *intervening* luck (see, e.g., Pritchard 2005). Whereas intervening luck is incompatible with aptness, more controversially, *environmental luck* is compatible with it. Though this point will not concern us here. For recent discussion on full aptness and its relationship with risk, see Carter (2020a).

(b) *An answer to Question 2*

In answering Question 2—about whether group judgmental knowledge is, on the proposal advanced here, about as common as we'd expect—I want to draw a brief parallel to the individual level. Is *individual* judgemental knowledge rare? Why would we think it is? Three kinds of arguments that might try to establish this (none of which is very persuasive) go as follows: group judgmental knowledge would plausibly be rare, on the telic theory proposed, if (i) *all* knowledge—judgmental or otherwise—is rare due to the epistemic hostility of our environment (e.g., if we are being often deceived); (ii) if there are barriers to individuals performing the act of judgment given what this involves on telic virtue epistemology; (iii) if judgment itself (as it is understood on telic virtue epistemology) is—despite there being no barriers to performing it—rarely performed nonetheless.

I'm going to set aside (i) out of hand, as it has very little bearing on what's of interest here, given that its implications concern much more than just judgmental knowledge. That said, (ii) and (iii) concern judgmental knowledge more directly. Regarding (ii): We've seen in §4 that typical arguments against the possibility of voluntary beliefs cut no ice against the idea that judgments are a species of intentional action. However, with this point established, it's unclear why there should be any barrier to individual judgment, or (iii) for that matter, why—given our interest in getting it right knowledgeably (rather than just getting it right any way) when we inquire—judgment is something an individual would rarely perform.

The point of working through the implausibility of any of (i-iii) as reasons to doubt that judgmental knowledge at the individual level is about as widespread as we'd expect is that *the same* rationale extends—*mutatis mutandis*—to the collective level. Again, setting aside (i), let's consider (ii). §5 showed in some detail why typical objections to the countenancing of inflationary group beliefs do not carry over to (inflationary) group judgment, at least as it is construed within a telic framework. Regarding (iii): there is no reason at all to think that such collective judgments are rare. On the contrary: collective judgment (as opposed to collective *mere* alethic affirmation) plausibly best describes what scientific research teams often do when endeavouring not merely to get it right any old way on a particular inquiry, but to get it right knowledgeably, and on this point juries are no different.<sup>46</sup>

<sup>46</sup>For a more general defence of idea that knowledge (rather than truth) is the norm of science, see Bird (2007a); Bird (2007b). For criticism, see Rowbottom (2008). Likewise, for de-

## 7. A COMPARISON

There is one other attempt in the literature to extend the (K=AB) model from individual to collective epistemology, due to Jesper Kallestrup (2016), in his paper ‘Group Virtue Epistemology’. There is a lot to like about Kallestrup’s project. For one thing, Kallestrup maintains that there is ‘nothing in [Sosa’s] framework precludes ascriptions of knowledge to group agents.’ Agreed. Even more, Kallestrup accepts inflationary (i.e., irreducibly collective) epistemic properties in his concession that it is ‘perfectly possible for groups to instantiate epistemic properties none of their members instantiate’. Also agreed. Thirdly, Kallestrup thinks that the right way to think of group knowledge is as apt belief, and even more, his position allows that a ‘group may form apt beliefs none of its members share’. This third point, along with the other two, line up exactly with my own thinking.

But Kallestrup and I—despite traveling the same road, from individual to collective epistemology, part of the way—diverge at several important places. I want to conclude by showing why I think that the view defended here, and not Kallestrup’s, gets it right on these key points of divergence.

### (a) *First point of divergence*

According to Kallestrup, while a group may form apt beliefs that none of its members share, the ‘competence’ of a group is nothing over and above the competences of its members when suitably combined. He writes:

Novel competences of groups do not spring into existence or mysteriously emerge when conjoining existing individual ones. On the other hand, the aptness of group belief is not similarly reducible to the aptness of the beliefs of its individual members (2016, 10).

In a bit more detail, Kallestrup’s contention is that the following thesis is true of group competences but false of knowledge *qua* apt belief

**Reductive individualism:** all (or at least most) of the members of *g* having *E*-type properties is necessary and sufficient for *g* having *E*. (2016, 10–11).

---

fences of the idea that a knowledge norm, rather than any weaker norm, governs assessments of whether to convict in criminal proceedings, see Littlejohn (2017) and Blome-Tillmann (2017).

There are two problems with this view. Firstly, empirical evidence for cases of ‘Mandevillian intelligence’ indicate that reductive individualism does not hold for group competences.

Cases of Mandevillian intelligence, in short, indicate that such a reduction is problematic because some dispositions that are *unreliable* and thereby are not individual-level competences can, and reasonably often enough do, lead to knowledge-conducive dispositions at the collective level, particularly when these individual-level shortcomings play the *de facto* role of generating cognitive diversity within a group. In a series of recent papers reviewing these kinds of cases, Paul Smart (2018a, 2018b) has noted individual-level ignorance, extreme-thinking and forgetfulness as (somewhat paradoxically) among the individual level traits that have been reported as contributing to epistemic benefits at the group level.<sup>47</sup> Whereas the countenancing of Mandevillian intelligence cases is incompatible with Kallestrup’s reductive individualism about group-level competences, it is compatible with my proposal which, unlike Kallestrup’s, does not make this commitment.

A second problem for Kallestrup’s reductive individualism about group competence is that it is paired with *non-reductive individualism* (i.e., the denial of reductive individualism) about group knowledge qua apt belief. This pairing is unstable in the following way: it posits a non-reductive subject as possessing an apt belief while denying that the *same subject*, the subject to whom knowledge is attributable, can possess a competence. This pairing implies that Kallestrup must reject the following which is analytically true within the kind of virtue-theoretic framework he takes himself to be operating: a subject, S, has an apt belief only if S has a competent belief.

(b) *Second point of divergence*

A second point of divergence concerns theoretical neutrality. According to (K=AB) proposals, knowledge is a *normative (epistemic) kind*, not a psychological kind. Such theories, at the individual level, do not carry any heavy-duty commitments about how beliefs are materially realised in the mind of a knower. For this reason, a proponent of (K=AB) at the individual level could potentially be paired with various kinds of views about beliefs in the philosophy of mind, e.g., from representationalism to functionalism.<sup>48</sup> The version of the (K=AB) template opted

<sup>47</sup>For representative empirical literature describing cases of Mandevillian intelligence, see Zollman (2010) and Xu, Liu, and He (2016).

<sup>48</sup>For an overview of some additional options, see Schwitzgebel (2019, sec. 1).



for here (§5) was shown to be compatible with very different theories of what has to happen to realise a knowledge-apt collective judgment. In this respect, the proposal is, like individual-level (K=AB) views, theoretically neutral about the conditions on realising belief. Kallestrup's proposal, by contrast, weds itself to several heavy-duty commitments on this front. To give one such example, for Kallestrup, an (inflationary) group belief is brought about, through the activity of individual members, *only* if those individuals' contributions are made (i.e., only if the individuals form the beliefs they form) *because* they believe that others intend to make the certain prescribed individual contributions towards the group attitudes (2016, 13). This caveat alone will make the view off limits to social-distributed views that lack any such explicit requirement. Even more, the view seems to generate some implausible predictions. For example, the view predicts that a scientific research team fails to generate group belief (and thus group knowledge) if each individual on the research team, in short, *would have kept on working* exactly as they had even had the others bunked off. And this looks like the wrong result.

(c) *Third point of divergence*

The third point of divergence concerns the matter of 'screening off' rejectionist-style objections from voluntariness. Just as individual-level acceptance is voluntary in a way that individual-level knowledge-apt belief is not, likewise—and this is just a restatement of the kind of point that has been made by various rejectionists about group belief (see §5)—a voluntary acceptance (joint or otherwise) of a proposition is not a knowledge-apt belief.

I've shown how collective *judgments*, though intentional, are not just usually, but essentially, not voluntary, and this result falls out of the normative structure of judgment as a distinctive kind of constitutive attempt, as understood within a telic virtue epistemology.

Kallestrup's proposal, at the collective level, does not (as mine does) unpack the (K=AB) template as a thesis about collective *judgmental knowledge* specifically—knowledge that implicates an intentional (and successful) constitutive aiming at aptness of alethic affirmation rather than at any other end—but rather, as a thesis about collective knowledge more generally, knowledge generated whenever collective 'belief' is apt.<sup>49</sup> But here is the problem: there is nothing in this proposal

<sup>49</sup>As it will be, on Kallestrup's proposal, if the belief's *truth* manifests a (individually reductive) competence to reliably enough generate true beliefs.

that screens off the objection that group ‘beliefs’ aren’t genuine (knowledge-apt) beliefs, given that—in light of how Kallestrup has described their conditions of realisation—they could in principle be brought about voluntarily. After all, on Kallestrup’s proposal, group beliefs—even though they purport to describe the world and as such are attitudes with a mind-to-world direction of fit—can be generated and reversed arbitrarily by a group so long as the group satisfies (following List and Pettit (2006)) conditions whereby the group acts intentionally to have that attitude stand *as* the group’s attitude (see, e.g., 2016, 13). But then, and taking a wider view, susceptibility to this kind of criticism leaves Kallestrup’s version of collective virtue epistemology one that preserves substantive symmetry (see §2) only the cost of inviting the charge that the view is not suitably *non-sceptical* at the collective level, even if substantively symmetrical.

## 8. CONCLUDING REMARKS

In sum, the simple idea, powerful in individual epistemology, that knowledge = apt belief is not a ‘one trick pony’. It can be made to work at the collective level as well. But in order to make it work, we have to understand *how* to make it work, and that requires availing ourselves of the full theory of telic epistemic normativity, within which we can distinguish different kinds of beliefs in light of the distinctive kinds of constitutive attempts they are at getting it right. With this in mind, the view is that ‘K=AB’ is correct theory of knowledge at the collective level *as a theory of collective judgmental knowledge*, or of *apt judgment*. Understood as such, the view was shown to be insulated against standard fare criticisms of collective belief that would seem, *prima facie* to pose problems for any account of collective knowledge that is ‘built’ out of any sort of collective belief. Moreover, the view is shown to diverge from an alternative view of collective virtue epistemology defended by Jesper Kallestrup in three important ways—and I’ve argued why, at each of these three forks in the road, there are key advantages to throwing in with the theory defended here.<sup>50</sup>

<sup>50</sup>Thanks also to Mark Alfano for helpful comments on a previous version of this paper. This paper was written as part of the Leverhulme-funded ‘A Virtue Epistemology of Trust’ (#RPG-2019-302) project, which is hosted by the University of Glasgow’s COGITO Epistemology Research Centre, and I’m grateful to the Leverhulme Trust for supporting this research.

## REFERENCES

- Alston, William P. 2007. "Audi on Nondoxastic Faith." *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*, 123–41.
- Beer, Randall D. 1995. "A Dynamical Systems Perspective on Agent-Environment Interaction." *Artificial Intelligence* 72 (1-2): 173–215.
- . 2014. "Dynamical Systems and Embedded Cognition." *The Cambridge Handbook of Artificial Intelligence*, no. 812: 856–73.
- Bird, Alexander. 2007a. "Justified Judging." *Philosophy and Phenomenological Research* 74 (1): 81–110.
- . 2007b. "What Is Scientific Progress?" *Noûs* 41 (1): 64–89.
- . 2010. "Social Knowing: The Social Sense of 'scientific Knowledge.'" *Philosophical Perspectives* 24: 23–56.
- . 2019. "Group Belief." In *The Routledge Handbook of Social Epistemology*, edited by Miranda Fricker, Peter J. Graham, David Henderson, and Nikolaj Pedersen. Routledge.
- Blome-Tillmann, Michael. 2017. "'More Likely Than Not' - Knowledge First and the Role of Statistical Evidence in Courts of Law." In *Knowledge First - Approaches in Epistemology and Mind*, edited by Adam Carter, Emma Gordon, and Benjamin Jarvis, 278–92. Oxford, UK: Oxford University Press.
- Brady, Michael S, and Miranda Fricker. 2016. *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*. Oxford University Press.
- Bratman, Michael E. 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press.
- Broncano-Berrocal, Fernando. 2018. "Purifying Impure Virtue Epistemology." *Philosophical Studies* 175 (2): 385–410.
- Carter, J. Adam. 2015. "Group Knowledge and Epistemic Defeat." *Ergo* 2 (28): 711–35.
- . 2016a. "Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution." *Pacific Philosophical Quarterly* 97 (1): 140–55.
- . 2016b. "Sosa on Knowledge, Judgment and Guessing." *Synthese*, 1–20.

- . 2020a. “De Minimis Normativism: A New Theory of Full Aptness.” *The Philosophical Quarterly*. <https://doi.org/10.1093/pq/pqaa017/5822748>.
- . 2020b. “Epistemic Perceptualism, Skill and the Regress Problem.” *Philosophical Studies* 177 (5): 1229–54.
- . 2020c. “Virtue Perspectivism, Externalism, and Epistemic Circularity.” In *Knowledge from a Human Point of View*, edited by Michela Massimi and Ana-Maria Cretu, 123–40. Springer.
- Cohen, L. Jonathan. 1992. *An Essay on Belief and Acceptance*. New York: Clarendon Press.
- De Ridder, Jeroen. 2014. “Epistemic Dependence and Collective Scientific Knowledge.” *Synthese* 191 (1): 37–53.
- . 2018. “Representations and Robustly Collective Attitudes.” In *Socially Extended Epistemology*, edited by Jesper Kallestrup J. Adam Carter Andy Clark and S. Orestis Palermos, 36–58. Oxford: Oxford University Press.
- Elgin, Catherine Z. 2017. *True Enough*. Cambridge, MA: MIT Press.
- Gilbert, Margaret. 1987. “Modelling Collective Belief.” *Synthese* 73 (1): 185–204.
- . 1992. *On Social Facts*. Princeton University Press.
- . 2013. *Joint Commitment: How We Make the Social World*. Oxford University Press.
- Greco, John. 2003. “Knowledge as Credit for True Belief.” In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, edited by Michael DePaul and Linda Zagzebski. Oxford: Oxford University Press.
- . 2009. “Knowledge as Success from Ability.” *Philosophical Studies* 142: 17–26.
- . 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Hakli, Raul. 2006. “Group Beliefs and the Distinction Between Belief and Acceptance.” *Cognitive Systems Research* 7 (2-3): 286–97.
- Hieronymi, Pamela. 2009. “Believing at Will.” *Canadian Journal of Philosophy* 39 (sup1): 149–87.

- Hutchins, Edwin. 1995. *Cognition in the Wild*. MIT press.
- Kallestrup, Jesper. 2016. "Group Virtue Epistemology." *Synthese*, 1–19.
- Kallestrup, Jesper, and Duncan Pritchard. 2012. "Robust Virtue Epistemology and Epistemic Anti-Individualism." *Pacific Philosophical Quarterly* 93 (1): 84–103.
- . 2013. "Robust Virtue Epistemology and Epistemic Dependence." In *Knowledge, Virtue, and Action: Essays on Putting Epistemic Virtues to Work*, edited by Tim Henning and David P Schweikard. London: Routledge.
- Kelp, Christoph. 2013. "Knowledge: The Safe-Apt View." *Australasian Journal of Philosophy* 91 (2): 265–78.
- . 2018. *Good Thinking: A Knowledge First Virtue Epistemology*. Routledge.
- Knobe, Joshua, and Jesse Prinz. 2008. "Intuitions About Consciousness: Experimental Studies." *Phenomenology and the Cognitive Sciences* 7 (1): 67–83.
- Lackey, Jennifer. 2007. "Why We Don't Deserve Credit for Everything We Know." *Synthese* 158 (3): 345–61.
- . 2014. *Essays in Collective Epistemology*. Oxford University Press, USA.
- . 2018. "Group Assertion." *Erkenntnis* 83 (1): 21–42.
- List, Christian, and Philip Pettit. 2006. "Group Agency and Supervenience." *The Southern Journal of Philosophy* 44 (S1): 85–105.
- Littlejohn, Clayton. 2014. "Fake Barns and False Dilemmas." *Episteme* 11 (4): 369–89.
- . 2017. "Truth, Knowledge, and the Standard of Proof in Criminal Law." *Synthese*, 1–34.
- Meijers, Anthonie. 1999. "Believing and Accepting as a Group." In *Belief, Cognition and the Will*, 59–73. Tilburg University Press.
- . 2003. "Why Accept Collective Beliefs?: Reply to Gilbert." *Protosociology* 18: 377–88.
- Palermos, S. Orestis. 2016. "The Dynamics of Group Cognition." *Minds and Machines* 26 (4): 409–40.

- . 2020. “Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism.” *Erkenntnis*, 1–20.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford University Press.
- . 2012. “Anti-Luck Virtue Epistemology.” *Journal of Philosophy* 109 (3): 247–79.
- . 2015. “Anti-Luck Epistemology and the Gettier Problem.” *Philosophical Studies* 172 (1): 93–111.
- Rowbottom, Darrell P. 2008. “N-Rays and the Semantic View of Scientific Progress.” *Studies in History and Philosophy of Science Part A* 39 (2): 277–78.
- Schechter, Joshua. 2019. “Aiming at Aptness.” *Episteme* 16 (4): 438–52.
- Schwitzgebel, Eric. 2019. “Belief.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2019. <https://plato.stanford.edu/archives/fall2019/entries/belief/>; Metaphysics Research Lab, Stanford University.
- Simion, Mona. 2019. “Knowledge-First Functionalism.” *Philosophical Issues* 29 (1): 254–67.
- Simion, Mona, J. Adam Carter, and Christoph Kelp. 2020. “On Behalf of Knowledge First Collective Epistemology.” In *Doxastic and Propositional Warrant*, edited by Paul Silva and Luis Oliveira. Routledge.
- Smart, Paul R. 2018a. “Mandevillian Intelligence.” *Synthese* 195 (9): 4169–4200.
- . 2018b. “Mandevillian Intelligence: From Individual Vice to Collective Virtue.” In *Socially-Extended Epistemology*, edited by J. Adam Carter, Andy Clark, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard, 253–74. Oxford, UK: Oxford University Press.
- Sosa, Ernest. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge University Press.
- . 1997a. “Mythology of the Given.” *History of Philosophy Quarterly* 14 (3): 275–86.
- . 1997b. “Reflective Knowledge in the Best Circles.” *The Journal of Philosophy* 94 (8): 410–30.
- . 1999. “How to Defeat Opposition to Moore.” *Noûs* 33 (13): 141–53.

- . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1*. Oxford University Press.
- . 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume 2*. Oxford University Press.
- . 2010a. “How Competence Matters in Epistemology.” *Philosophical Perspectives* 24 (1): 465–75.
- . 2010b. “Value Matters in Epistemology.” *The Journal of Philosophy* 107 (4): 167–90.
- . 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- . 2020. *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains*.
- Tollefsen, Deborah Perron. 2015. *Groups as Agents*. John Wiley & Sons.
- Tuomela, Raimo. 1995. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press.
- Turri, John. 2011. “Manifest Failure: The Gettier Problem Solved.” *Philosopher’s Imprint* 11 (8): 1–11.
- Van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wray, K Brad. 2001. “Collective Belief and Acceptance.” *Synthese* 129 (3): 319–33.
- Xu, Bo, Renjing Liu, and Zhengwen He. 2016. “Individual Irrationality, Network Structure, and Collective Intelligence: An Agent-Based Simulation Approach.” *Complexity* 21 (S1): 44–54.
- Zagzebski, Linda Trinkaus. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.
- Zollman, Kevin JS. 2010. “The Epistemic Benefit of Transient Diversity.” *Erkenntnis* 72 (1): 17.