

Knowledge as a collective status

Jeremy Randel Koons 

Georgetown University, Doha, Qatar

Correspondence

Jeremy Randel Koons, Georgetown University, Doha, Qatar.
Email: koonsj@georgetown.edu

Abstract

While social epistemology is a diverse field, much of it still understands knowledge as an *individual status*—albeit an individual status that crucially depends on various social factors (such as testimony). Further, the literature on group knowledge until now has primarily focused on limited, specialized groups that may be said to know this or that as a group. I wish to argue, to the contrary, that *all* knowledge attributions ascribe a *collective status*; and that this follows more or less directly from an essential function of entitlement-ascriptions: Ascriptions of knowledge and entitlement serve a primarily social function in that they facilitate coordination by *maintaining consensus around true beliefs, true theories, and truth-producing methodologies*. This conclusion will shed light on ways in which traditional theories of knowledge (such as foundationalism and coherentism) fail to capture a central function of our epistemic practice.

1 | INTRODUCTION

Sellers argues that moral judgments express collective intentions. Can a similar case be made for other normative attitudes and statuses? In particular, what about knowledge attributions? I will argue that knowledge attributions ascribe a collective status; and that this follows from the function of knowledge- and entitlement ascriptions: Such ascriptions serve a primarily social function in that they

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Analytic Philosophy* published by John Wiley & Sons Ltd.

facilitate coordination by *maintaining consensus around true beliefs, true theories, and truth-producing methodologies*.¹

I hope thereby to offer a contribution to a social epistemology that is truly *social*. While social epistemology is a flourishing and diverse field, much of it still focuses on knowledge as an *individual status*. Thus, to use Martin Kusch's term, much social epistemology is “dualistic” (2002, pp. 113–15)—it assumes individual and social branches of epistemology, and asserts that knowledge can be an individual status, albeit one that crucially depends on various social factors (such as testimony). While this kind of individualism is not universal in social epistemology², it is widespread; and I shall push back against some of the assumptions underpinning it.

I hope this project also offers a helpful adjunct to some of the more radically social epistemologies, such as those of Alexander Bird, Martin Kusch, Helen Longino, and Lynn Hankinson Nelson. Bird, for example, has defended a “social-social epistemology” according to which emphasizes “the social nature of the epistemic subject” (2010, p. 23). Bird argues (convincingly) that there are many things that we as a community know, even if (given the epistemic division of labor) there is only a handful of individuals who possess this knowledge *qua* individuals—or even if no individuals at all possess this knowledge *qua* individuals. What I offer is an explanation of why knowledge attributions work like this—that is, why attributing knowledge to (say) a group of experts is *a fortiori* to attribute a status to the epistemic community as a whole. The explanation, we will see, is to be found by examining the *point* or *purpose* of making knowledge attributions in the first place—by examining their *role* within a rational community. Thus, I hope that this paper makes (at least) two novel contributions to the literature. First, a number of authors have argued that communities are the prime possessors of knowledge and that individuals can only know insofar as they are members of a community.³ I explore the consequences of this type of view, and argue—crudely stated—that knowledge attributions (whether self-attributions or attributions to another) have roughly the form “We know that *p*,” not “I know that *p*” or “You know that *p*.”⁴ Authors have not explored in-depth what the communal nature of knowledge means for how we ought to interpret knowledge attributions, and I will argue that knowledge claims attribute a collective status. I think a view like mine is implicit in or implied by these radically social epistemologies, but a more detailed account needs to be given of what we are doing when we make knowledge attributions on such epistemologies. (And, as we will see, to the extent that they do offer comments on the role of knowledge attributions, some authors—like Kusch—say things that are not consistent with their overall project; I wish to offer an account of knowledge attributions that is more in line with the communitarian epistemology of such authors.)

¹Two caveats/qualifications: First, while in the general case consensus is maintained around beliefs we take to be true, etc., as it will turn out, often—at least in scientific practice—epistemic practice will maintain consensus around a theory which practitioners hold to be best supported by the evidence, but of whose truth they are unsure. To add this qualification into the above formulation would make it extremely unwieldy, but the reader should take the qualification as implicit. Second, it should go without saying that such assessments of truth are always fallible, provisional, and made—in a sense to be articulated as we go along—according to the lights of a social practice's present theoretical and conceptual scheme.

²Some examples of epistemologists who explicitly push back against these individualistic assumptions are Nelson (1993), Kusch (2002), Longino (2002), Bird (2010), Miller (2015), and Green (2016).

³For example, “Communities, not individuals, are the primary loci of knowledge” (Nelson 1993, p. 131); “Empirical beliefs can be rational or irrational only in so far as the attributors of the respective beliefs are members of epistemic communities” (Kusch 2002, p. 87).

⁴My thesis is that all knowledge claims have the form “We know that *p*.” I believe the argument of this paper rules out the possibility of legitimate knowledge attributions that take an individualistic form (e.g., “I know that *p*”); but if for some reason I am wrong and there are such legitimate attributions, there would have to be something deeply deviant about such cases, some feature that explained their divergence from the normal case.

Second, my argument for why knowledge attributions take this form—due to the role of epistemology in maintaining doxastic consensus—offers a novel explanation for the specifically communal nature of these attributions. Thus, not only does my account of the role of knowledge attributions cohere with prominent communitarian epistemologies; but it also offers a satisfying theoretical explanation of *why* knowledge attributions take this particular form.

Two broadly pragmatist commitments will structure the argument of this paper (although I hope to characterize them weakly enough to make them palatable to a wide range of readers). First, I am less interested in analyzing knowledge as a property of beliefs or agents, and more interested in analyzing the purpose or point of making knowledge ascriptions.⁵ Thus, just as many claim a central purpose of our moral practice is *social coordination*, I want to uncover what some central purposes are of our epistemic practice—and what this entails for the kind of status we are ascribing when we do ascribe epistemic statuses like knowledge or entitlement.⁶

The second commitment is more methodological than substantive: Following a Brandomian pragmatist strategy, I will “prefer an order of explanation that begins with what is implicit in practice (what people do) and proceeds to an account of what they explicitly believe or say, over one taking the opposite tack” (Brandom, 1994, pp. 101). Thus, I will mostly focus on attributions of entitlement that are implicit in our various discursive practices. I will discuss explicit knowledge attributions, but the latter are (for the pragmatist expressivist) parasitic on the former, and any investigation of knowledge and/or entitlement must begin with such statuses as they are implicitly attributed by our practical activity.

Not every methodological principle can be defended in every paper; and I do not propose to defend this one here. But hopefully, I have stated these commitments in a way that the argument that follows will be of interest even to those who are skeptical of a general pragmatist outlook; and hopefully, the utility of my approach will become apparent as the argument progresses. I will note that the first commitment confers a significant argumentative advantage on the present account: Since, according to the account offered here, knowledge attributes a collective *status*, I am able to sidestep debates in social epistemology about the ontology of group minds. This is an advantage of expressivist accounts generally and was a chief motivation for Sellars's expressivism. Knowledge claims attribute a normative status; they are not descriptive, and we thus do not need to fit the “normative facts” they allegedly describe into our ontology.

I will argue my case incrementally. I begin with simpler claims (e.g., outlining the nature of collective attitudes, arguing for the role of knowledge ascriptions), and gradually build toward the more difficult-to-establish ones (e.g., that one can make a legitimate knowledge attribution even if the belief in question is not widely shared in the community). I will summarize our conclusions at the end of each section so that the reader can track our overall progress.

⁵I doubt, actually, that knowledge *is* a property of beliefs or agents; as Sellars says, “in characterizing an episode or state as that of *knowing*, we are not giving an empirical description of that episode or state” (EPM, VIII. §36/p. 169). But pursuant to my strategy of choosing my battles, I will not push the point here.

⁶This distinguishes the present project from others to which it bears some similarity, such as Alexander Bird's. Bird—while recognizing that social knowledge serves a “social function,” such as coordinated action—is more interested in analyzing the conditions under which social knowledge is appropriately attributed. I want to start with the prior question of what the function of such attributions are, and what this means for the status we are attributing when we do attribute knowledge or entitlement. Similarly, authors like Kusch and Nelson—who appropriately argue that knowledge is often appropriately attributed to communities (and not individuals), say little about the *structure* of knowledge attributions, or *why* they function the way that they do. My contribution, I hope, would be to help fill this lacuna in the literature.

2 | WE-ATTITUDES, IN THE STRONG SENSE: A SKETCH OF A PRAGMATIST ACCOUNT

Ultimately, I will argue that when we ascribe knowledge or entitlement, we are ascribing a collective status. Present accounts of collective attitudes, however, are much better developed than are accounts of collective statuses; and so I shall begin by looking at what it is to have a collective *attitude*. We can subsequently build on this account to develop an account of knowledge as a collective *status*.

Elsewhere⁷, I offer a detailed pragmatist account of collective attitudes, but here I will have space for little more than a sketch of such an account. There are many accounts of collective intentionality on offer⁸, but for the sake of brevity in exposition I will focus on just one particularly well-developed one. Thus, skipping over a lot of preliminaries, I will focus on Tuomela's account of collective attitudes. On Tuomela's account, a person has a we-attitude (such as a collective belief) in the full-blown, strong sense iff these conditions are met:

1. the person has the attitude (e.g., the belief-that-p)
2. the person believes that the others in the group also have the attitude, and also
3. the person believes that it is mutually believed among members of the group that the members have this attitude.⁹

Tuomela holds that social institutions in general constitutively depend on this mutual belief element—if not psychologically, then conceptually/presuppositionally. So if I want to speak to you using the word “bird” to mean *bird*, then:

- a. I must believe that “bird” means *bird*,
- b. I must believe that “bird” means *bird* for you, and
- c. I must believe that *you* believe that “bird” means *bird* for *me*.

Thus, this element of *mutual belief* would be a conceptual presupposition of language use, and indeed of the use of norms generally.

However, Tuomela's account of group attitudes as involving mutual belief is too strong. For many cases, this requirement is plausible—particularly for the examples that populate the literature on group action (e.g., two people carrying a table together). But when we scale up our account to involve large social collectives (like users of a particular language), the requirement is less plausible. Something *like* mutual belief is at play, but (for example) I am not sure it is plausible to attribute to me a belief that you *believe* that I use the word ‘bird’ to mean *bird* when I speak.

I would not use the word “bird” in speaking to you if I did not *in some sense* think that there was a shared understanding of what the word meant. The key is fleshing out this “in some sense.” Recall Brandom's methodological advice, noted in the introduction: On the pragmatist way of thinking, one should “prefer an order of explanation that begins with what is implicit in practice (what people do) and proceeds to an account of what they explicitly believe or say, over one taking the opposite tack” (Brandom, 1994, pp. 101). And what is implicit in our language use and other norm-governed social interactions with others is a set of *expectations* about the commitments and behaviors of other agents.

⁷For a much more detailed account of collective attitudes from within a Sellarsian and pragmatist framework, see (Koons 2019, Chapter 4; Koons 2021).

⁸See, for example, Bratman (1999), Searle (2009), Gilbert (2014), and Tollefsen (2015).

⁹See Tuomela 2002, p. 23, for his formal statement of these conditions.

What do I mean by “expectations”?¹⁰ Imagine someone walking a meandering path that threads through a series of ornamental ponds. The person may never entertain the belief that the path will support her weight, or that the water in the pond will not, or that if she attempts to walk across the pond she will get wet. But implicit in her behavior is a set of expectations about the way the world behaves. If the path instead meandered through a series of grassy patches, the person might cut across the grass, her behavior demonstrating her differing expectations about the behavior of the world. “Expectation” might mislead—it implies an occurrent mental state where there probably is none. Again, remember our pragmatist commitment: We are making explicit commitments, expectations, and so on that are *implicit* in behavior, rather than attributing to the agent occurrent states to *explain* the behavior in question.

The case is similar with agents who share a form of life. When I utter an English sentence to a fellow English user, implicit in my behavior is the expectation that this will have a certain effect on the person in question. It might simply result in an updating of the other agent's “scorecard.” Or it might result in a certain corresponding behavior on the part of the other agent (if, for example, I have asked her for the time). But implicit in our real-time interactions with others is the *expectation* that we do share this form of life—that we share a language, or a set of norms for greeting (e.g., shaking hands), or for conducting transactions in the checkout line at the grocery store, or rowing a coxed eight, etc. Indeed, this element of mutual expectation—that each of us understands our role, that we will fulfill our part, and so on—is an ineliminable part of the form of life. (What would be the point of bringing groceries to the checkout counter if I did not expect the cashier to sell them to me? What would be the point of my ringing up your purchase if I did not expect you to pay? What would be the point of using language if I did not expect the appropriate updating and/or behavior on your part?).

One might worry that my move from belief to expectation in this account relies on a notion of beliefs as occurrent and explicit, when in fact most of our beliefs are implicit. For example, I currently believe all manner of things (that Napoleon is not hiding under my desk, that no trees are taller than the Empire State Building, etc.) that in general never find explicit formulation. Nevertheless, we should prefer expectation to belief in our account. First, even if beliefs are mostly implicit, such beliefs are usually capable of explicit formulation. However, many of our expectations are behavioral and perhaps deeply subpersonal; as such, they may defy easy formulation by the agent. As with Heideggerian skills, the agent may be unable to consciously formulate all of the elements of expectation that go into her social interactions, and so beliefs do not perfectly model expectations. Second, “expectation” is a thoroughly psychological/behavioral notion, not a normative one, in the following sense: To attribute a commitment to someone is to *prescribe* a set of behaviors to that person. To attribute an expectation to someone, on the other hand, is to describe or make a prediction about that person. Now, the expectation itself might concern someone else's norm-governed behavior—I expect the cashier to ring up my purchases—but insofar as it is appropriate to attribute an expectation to me, you are not attributing a normative status *to me*. And belief is at least partially normative (and not merely psychological): To attribute a belief to someone involves, at least in part, attributing some (unspecified, indefinite¹¹) range of *commitments* to her—and to attribute a commitment *is* to attribute a normative status. Thus, the notion of expectation cannot simply be replaced in the present account with that of belief, even on the understanding that beliefs can be implicit.

¹⁰An account of expectations similar to what follows is offered by Goldberg (2018). I am grateful to an anonymous referee for directing me to this source.

¹¹As Brandom notes, “The point is not that there is any particular set of such discriminations that one must be able to make in order to count as deploying the concepts involved. It is that if one can make *no* such practical assessments of the counterfactual robustness of material inferences involving those concepts, one could not count as having mastered them” (Brandom 2015, p. 142).

Thus, for the Brandomian, pragmatist reasons gestured at above, we should extend the mutual belief requirement to include also this element of mutual expectation implicit in practice. Naturally, some cases of group action *do* involve mutual belief. But most cases of group action will *not* involve mutual belief, and this is too strong of a requirement to impose. So from now on, let us understand the mutual belief requirement as involving mutual belief *or expectation* (in this implicit, Brandomian sense).

Thus, to re-cast our example in terms of expectation rather than belief, suppose I want to speak to you using the word “bird” to mean *bird*. Implicit in my behavior is.

- a. a *commitment* on my part that “bird” means *bird*,
- b. an *expectation* on my part that “bird” means *bird* for you, and
- c. an *expectation* on my part that you have an equivalent expectation, i.e., an expectation that “bird” means *bird* for me.

Given (a’)-(c’), my *commitment* regarding the meaning of bird is a *commitment in the we-mode in the strong sense*, as the mutual *expectation* of this commitment underlies the possibility of linguistic communication. That is, this element of mutual expectation is partially constitutive of linguistic norms, as the ability to communicate using linguistic tokens presupposes this element. Thus, the linguistic norms governing the use of “bird” must be expressed in the strong we-mode.

I am not using “commitment” or “expectation” with any particular technical definition in mind. One might give “commitment” a Brandomian gloss, in which the term has both psychological and normative dimensions. Thus, to attribute a commitment to someone is to convey certain *behavioral* expectations (say, that the person will not infer “is the largest land mammal” from “is a bird”). But it is also to convey certain *normative* expectations, expectations which are not conveyed merely by a behavioral prediction. (As I just noted, to attribute an expectation to someone is not to attribute a normative status—it is not to *prescribe* a set of behaviors to that person—but we can have expectations *about* norm-governed behavior; and this is what is at issue here.) I should emphasize here that unlike a Brandomian commitment, expectations as I understand them are not normative in any robust sense—they are descriptive (although as can be seen from the above example, they may involve a counterfactual element). In any case, as I said, I do not have a formal view on this question, but am relying on a more or less intuitive set of concepts.

There is another way Tuomela’s account is too strong. Tuomela insists that acting in the we-mode in the strong sense requires acting for a group *reason*. But as I argued above, action in the we-mode in the strong sense can instead involve mutual *expectation*, where these implicit expectations are *presuppositions* of the action rather than *reasons for* the action. Here is another analogy: Crucial to Sellars’s anti-foundationalist picture of knowledge is that a bit of language A can be conceptually and epistemically dependent on some theoretical commitments T, without A being inferred from T. Nevertheless, this dependence makes T epistemically prior to A in a way that keeps A from being foundational; the dependence is crucial to both the epistemic and conceptual status of A. This is the role, for example, played by knowledge of standard viewing conditions in Sellars’s account of perceptual knowledge. Knowing that one is in standard viewing conditions is a *condition* on an observation belief being justified, but is not a *premise* from which the observation belief is inferred. Ergo, for Sellars, there is *epistemic priority* that is not necessarily *inferential dependence*. This is also how we should understand the role of the above implicit expectation condition in group attitudes, such as in the deployment of linguistic norms, or the utilization of social institutions like money. (In the language of defeasible inferences, these expectations are “enablers”; they should not be seen as “premises” that would be appealed to if we were to reconstruct the agent’s reasons for action.) Thus, acting in the we-mode in the strong sense does not seem to require either mutual belief, or acting on a group reason. Rather, it seems to require the more modest requirement of mutual expectation or joint commitment, which is a less intellectualist, more *pragmatist* view of collective action.

2.1 | *Intermediate conclusions/looking ahead*

What I have tried to do in this section is—very briefly—establish a pragmatist account of collective *attitudes*, one that I will press into service in developing an account of collective *statuses*. Much has been written about group commitment or group belief. It is plausible that in general, social practices—in particular, linguistic activity, which is arguably necessary for sapient cognition—involve commitments in the we-mode in the strong sense; the mutual expectation of this commitment underlies the possibility of linguistic communication.

An advantage of the present account is that while acting on norms (such as linguistic norms) presupposes—as a *conceptual* matter—the existence of group attitudes (such as collective commitments), purely as an *ontological* matter such collective commitments presupposes nothing beyond individually held attitudes of commitment and expectation. Thus, we are able to offer a pragmatist account of collective attitudes that does not commit us to anything ontologically suspect, such as group minds. (I argue for this at greater length in [Koons, 2019, chapter 4].) This should be kept in mind as we continue to develop and rely on this account of collective attitudes.

But what I need to establish, in order to introduce a social epistemology modeled on these lines, is not merely we-mode commitment. We need we-mode *entitlement*. Thus, I need to start by showing that at least the standard case of knowledge involves *a mutual undertaking and attribution of entitlement*. To be sure, there are special cases of group knowledge, and these have been the subject of much literature. I want to see how far we can extend the argument beyond these special cases. But I will start small, with common knowledge.

3 | EPISTEMIC PRACTICE AS COMMITMENT COORDINATION

Much has been written about morality as a tool for social coordination. Curiously, less has been written about epistemology as serving this same purpose. This, despite the fact that epistemologists of various stripes have long recognized some version of Clifford's exhortation:

And no one man's belief is in any case a private matter which concerns himself alone. Our lives are guided by that general conception of the course of things which has been created by society for social purposes. Our words, our phrases, our forms and processes and modes of thought, are common property, fashioned and perfected from age to age; an heirloom, which every succeeding generation inherits as a precious deposit and a sacred trust, to be handed on to the next one, not unchanged, but enlarged and purified, with some clear marks of its proper handiwork. Into this, for good or ill, is woven every belief of every man who has speech of his fellows. An awful privilege, and an awful responsibility, that we should help to create the world in which posterity will live. (Clifford, 1876, p. 292).¹²

Of course, morality cannot aid in social coordination unless people overwhelmingly begin with the same beliefs. Thus, convergent behavior *presupposes* convergent belief.¹³ Thomas Kuhn wrote about the

¹²However, Clifford thought of epistemology in largely individualistic terms and thought that holding to evidence would serve the purpose of maintaining this consensus. (He does, though, in the second—and seldom-cited—half of the article, emphasize the indispensable role of testimony in informing us about the world.)

¹³Point of clarification: It is of course widely recognized that coordination required shared beliefs or commitments. But there seems to be much less recognition that *epistemic evaluation*, specifically, is the social tool used to maintain convergence in belief and commitments.

importance of convergent thinking for the special case of the sciences and noted that the pre-convergent thinking history of virtually every scientific discipline is a period characterized by “very little progress” (Kuhn, 1977, p. 231).

I want to extend the argument beyond the special case of the sciences and argue that a chief job of epistemic evaluation is to maintain consensus by sorting beliefs and theories into those that *are to be believed*—because they are *held to be true*—and by differentially allocating epistemic authority—parceling it out to those whose pronouncements are *likely to be true*. There are correspondingly two interrelated dimensions to our epistemic practice. One is the *consensus-maintenance* aspect; the other is the *truth-orientation* aspect. I will discuss these in turn.

3.1 | Epistemic evaluation and consensus maintenance

It is a commonplace that knowledge can be implicit in practice, but much epistemology focuses on explicit attributions of knowledge. I shall first focus not on implicit *knowledge*, but on implicit *attributions* of knowledge, entitlement, and other epistemic statuses. A reminder that on our pragmatist order of explanation, such evaluation (like all evaluation) is primarily implicit in practice, rather than explicit. Thus, instead of beginning our examination with overt knowledge attributions, we should begin with what knowledge- or entitlement attributions look like when implicit in practice. For example, scientists practice implicit epistemic evaluation by applying a particular paradigm, and by excluding in various ways those who do not apply this paradigm. Thus, physicists working in 1800 generally believed that Newtonian physics correctly accounted for the motion of bodies, and their application of this theory to planetary motion amounted to an implicit endorsement of this theory. Further, a physicist who attempted to use a different, incompatible paradigm (say, Aristotelian physics) to account for planetary motion would have been sanctioned by partial or complete exclusion from various marks of professional recognition (publication, membership in professional societies, opportunities to address such learned societies, etc.)¹⁴. This exclusion is also a type of implicit epistemic assessment. Implicit epistemic assessment is evident in a variety of other scientific practices—e.g., conducting double-blind studies instead of consulting tea leaves, publishing studies that have used double-blind studies and not those that used tea leaf reading, only publishing studies where the regression on the data has yielded a sufficiently low *p*-value, etc. One can, of course, *make* these assessments *explicit*: An 18th-century physicist could *say* that Aristotelian physics has been refuted, or that its application is wrong or irrational; a journal editor can *say* that a *p*-value is too high to rule out the null hypothesis, or that tea leaf reading is an irrational method; etc. But in doing so, one is only making explicit normative assessments that were already implicit in the behavior of the practitioners.

We can also see the role of epistemic evaluation in maintaining consensus by looking at cases where such evaluation has failed to be fully effective. In recent years, we commonly decry political polarization that exists in democratic countries, for example, in the US. Discussions of this polarization often take one of (at least) three explicitly epistemological guises.

In the first guise, this discussion ties polarization to an insensitivity to certain kinds of evidence. Thus, liberals might accuse conservatives of failing to respond appropriately to evidence of racist policing, or the ineffectiveness of supply side economics, or the evidence of lower rates of criminality among immigrants, etc. (Or scientifically minded individuals might accuse anti-vaxxers of

¹⁴Some disagreement is inevitable and allowed within research paradigms; see, for example, Miller (2013) and Dang (2019). The point, however, is that if scientists do not impose broad conformity around the paradigm itself, science does not progress. Thus, the consensus-maintaining role of epistemic evaluation is essential to scientific progress; this is, I take it, one of the essential lessons to be gleaned from Kuhn (1977).

insensitivity to the evidence that vaccines are not tied to conditions such as autism.) Proposed remedies to these problems often take the form of suggestions for how to make people more epistemically virtuous—that is, how to correct for known psychological processes (negativity bias, the halo effect, confirmation bias, etc.) that lead to epistemically vicious thought.

In its second guise, this discussion accuses individuals of failing to practice the epistemic virtues by failing to seek out a variety of sources for their information and opinions, and only getting information from sources that reinforce their pre-existing judgments (epistemic bubbles and echo chambers). The proposed remedy, of course, is to urge people to practice Millian epistemic virtue—to seek opinions from a variety of sources, including (perhaps especially) those that conflict with one's own existing judgments and opinions.

The third guise is the most interesting, because it is the least individualistic. Instead of focusing on the failings of individual agents, it focuses on our failure to maintain a healthy epistemic environment. Recognizing that belief formation and -maintenance are affected by various social conditions, critics often point to various disinformation campaigns (e.g., climate change denial campaigns funded by fossil fuel companies and other groups) whose purpose is to undermine consensus and muddy the epistemic waters. On this view, it is a mistake to tie the current lack of consensus wholly (or even primarily) to lack of good epistemic practice by individuals.

These explicitly epistemic projects are motivated by practical concerns about democratic society. Whatever people's normative judgments or political views, constructive discussion about public policy or coordinated action can only proceed on the basis of general *factual* agreement—e.g., agreement over the health benefits and risks of vaccination, the economic costs and benefits of various levels of immigration, the economic consequences of this or that taxation policy, and so on. Thus, this epistemological project is motivated by a concern that this polarization, with all of its negative social consequences, results largely from a failure of society's larger *epistemological* practice. Again, we see that one of the things—perhaps the most important thing—we are trying to do when we hold each other accountable to publically shared standards of evidence and rationality is to *maintain the system of shared commitments/entitlements that underlie the very possibility of group action*—and, at the end of the spectrum, the possibility of human society. Ironically, I believe that a philosopher who saw this the most clearly—and defended this the most eloquently—was Clifford, whose famous essay is mostly remembered as an individualistic defense of the obligation to collect evidence. More interestingly, though, the essay defends the claim that the duty to believe responsibly is *a duty we owe to each other, and to society at large*. Clifford writes, for example:

Belief, that sacred faculty, which prompts the decisions of our will, and knits into harmonious working all the compacted energies of our being, is ours not for ourselves but for humanity. It is rightly used on truths which have been established by long experience and waiting toil, and which have stood in the fierce light of free and fearless questioning. Then it helps to bind men together, and to strengthen and direct their common action. It is desecrated when given to unproved and unquestioned statements, for the solace and private pleasure of the believer...It is not only the leader of men, statesman, philosopher, or poet, that owes this bounden duty to mankind. Every rustic who delivers in the village alehouse his slow, infrequent sentences, may help to kill or keep alive the fatal superstitions which clog his race. Every hard-worked wife of an artisan may transmit to her children beliefs which shall knit society together, or rend it in pieces. No simplicity of mind, no obscurity of station, can escape the universal duty of questioning all that we believe. (Clifford, 1876-7, pp. 292-3)

Thus, you can see present demands for people to abide by shared and appropriate epistemic standards as a response to the breakdown in *consensus around the truth* required for a flourishing society and the kind of communal action made possible by (and largely constitutive of) such a society; and as an attempt to regain this consensus.

These considerations should be regarded as the opening salvo in an argument for the conclusion that coordination of doxastic commitments is a central function of our epistemic practice. The full case will be made in sections IV and V, hand-in-hand with the case that knowledge is a collective status.

3.2 | Epistemic evaluation as truth oriented

We can cast our eyes back over centuries, and millennia, of inquiry, and conclude that for most of human history, humans knew very little; for most of what they thought they knew was false. This tells us very little, however, about what role epistemic attributions (both explicit and implicit) played *within* these historical communities.

Perhaps, it is helpful to start off by considering our own practice of knowledge attribution. When, for example, scientists say that we *know* some fact (say that the Earth is roughly 4.54 billion years old), what do they mean by this? At the very least, they are saying that this claim is *true*.¹⁵

But what is the nature of the connection between truth and our epistemic practices? If our present scientific view of the world turned out to be largely false (as has been the fate of previous generations' scientific world views), it would be a mistake to say that our present inquiries had not all along been guided by epistemic norms *merely because* our beliefs are factually wrong. This is to confuse truth with entitlement. But this also suggests that a practice is defined by the end which it pursues, not by whether some portion of the practice (or even by whether some historical segment of the practice) fails to achieve this end.

Thus, for example, the practice of medicine is defined by (say) preserving and restoring health. Identifying a practice by reference to its ends means, for example, that someone who "treats" patients with the goal of shortening their lives or decreasing their health simply is not practicing medicine (although they may be making use of medical knowledge in so "treating" their patients). Crucially, this also means that a practice retains its identity even if the practice fails to achieve its ends. For example, until recent centuries, physicians could do little to improve the health of their patients, and many of their therapies (such as bloodletting) actually harmed, rather than improving, patients' health. Does this mean that these physicians were not practicing medicine, and indeed that they were not even physicians in the first place? Of course not; the practice of medicine existed in (say) the 16th century, but it was not particularly successful in achieving its ends—namely, the promotion of health.

Similarly, our practice of evaluating beliefs, arguments, evidence, and so forth counts as an epistemic practice because it is defined by the goals of understanding, explanation, truth, and instrumental control over nature. So one who aims to spread disinformation simply is not being governed by epistemic norms, but by political or propagandistic ones. Further, a scientific practice that once upon a time maintained consensus around Aristotelian physics, or the theory of spontaneous generation, or caloric theory, or what have you, was still guided by epistemic norms in that the practice as a whole was (presumably) directed toward uncovering the truth, understanding and explaining the natural world, and so on.

¹⁵An anonymous referee has challenged whether this is true, asking whether scientists do not sometimes mean by this that the theory is the best supported by the available evidence (while remaining agnostic as to the truth). I reply that scientists are often fully committed to theories and their investigation—witness the Standard Model in physics—but I do not think that scientists would assert that we *know* such theories are true before they have been experimentally confirmed.

This account has the advantage of making the connection between our epistemic practice and truth conceptual (in that the practice is defined in terms of its pursuit of the truth) while allowing us to say that inquirers in disciplines that for most of history were simply wrong about most things (as in many scientific disciplines) were still guided by epistemic norms.

Given the two dimensions of our epistemic practice, there are two distinct ways in which our epistemic practice can fail. First, our epistemic practice can fail to maintain consensus. I have discussed this already in section III.1. Epistemic practice can fail along the second dimension, however: It can maintain consensus, but in the wrong direction. Consider the example of Ignaz Semmelweis. While working at the Vienna General Hospital in the late 1840s, he instituted a policy of rigorous handwashing for doctors and medical students who worked on cadavers before they were permitted to assist with childbirth. This resulted in a 90% reduction in maternal mortality. However, his handwashing theory was rejected by the medical establishment and did not become standard practice until more than 20 years later, when Pasteur's germ theory of disease provided a theoretical explanation for the effectiveness of good sanitation. This is a clear example where consensus was maintained—but around a bad theory or practice, suppressing the good theory/practice.

That such examples exist, is undeniable. But they do not demonstrate that the present account of the function of our epistemic practice is mistaken. I have already noted the commonplace observation that morality is a tool for social coordination. Of course, it cannot *just* be a tool for social coordination; ideally, we coordinate around some kind of optimal strategy. So a moral practice—like an epistemic practice—idealizes along two dimensions. In the moral case, these are (a) the *coordination dimension* and (b) the *optimization dimension*. A moral practice can fail along either of these dimensions. It can fail along the first (consensus can break down), or it can fail along the second (as when there is consensus, but people converge on a suboptimal strategy). None of this shows that the point of morality is not coordination *and* optimization; it merely gives us a tool for diagnosing what is wrong with a malfunctioning moral practice.

Similarly, an account of the “function” of a healthy epistemic practice gives us the tools to diagnose an unhealthy one. An epistemic practice also has *both* a coordination dimension *and* an optimization dimension. A practice that fails to produce consensus (i.e., fails along the coordination dimension), or which produces consensus around false beliefs while suppressing true ones (i.e., fails along the optimization dimension), or which allocates epistemic authority in ways that are not connected to actual epistemic privilege—a topic which has been much discussed since Miranda Fricker's landmark (2007) work—is an epistemic practice that is, to that extent, malfunctioning. (This latter example would perhaps be failure along a third, meta-dimension, regarding allocation of epistemic power and who exerts force in the coordination dimension.)¹⁶

At any rate, it cannot be doubted that consensus maintenance has often been sought for non-epistemic reasons. No doubt, for example, much of the Inquisition was prosecuted as much to maintain

¹⁶Similar comments can be made about the possibility of bias—for example, when physicians use the male body as the model for human physiology and ignore differences with the female body. Feminist epistemologists, for example, have often used bias as an argument in favor of understanding knowledge as communal, for the simple reason that bias shows how evidence is not self-interpreting, but can only be interpreted in light of communal standards, which may embody (for example) androcentric biases. Thus, as Nelson writes, “Viable theories, like evidence, are not self-announcing... The standards and knowledge that underwrite the acceptability of androcentric and feminist assumptions are communal” (Nelson 1993, pp. 146 and 147). Again, though, the fact that an ineliminable feature of our epistemic practice is open to abuse is not an objection against understanding our epistemic practice as involving this feature. For example, complex human culture necessarily involves a division of epistemic labor, and therefore a differential distribution of epistemic authority. Epistemic authority can, unfortunately, be unjustly misallocated—resulting in epistemic injustice—but the possibility of abuse does not change the fact that no society can grow to complexity without a complicated division of epistemic labor (and a corresponding parceling out of epistemic authority). I am grateful to an anonymous referee for raising this issue of bias.

the hegemony of the Catholic Church as it was to stamp out heresy *qua* false belief. But that just makes such pursuit of consensus politics or propaganda, and not strictly epistemology. Or it makes it deviant epistemology—much like the epistemic injustices described by Fricker involve a malfunctioning epistemic practice.

I will say no more about this now; I merely wanted to sketch the two dimensions of our epistemic practice, and their relation. It is important to gesture, at least, to the connection between entitlement and knowledge, on the one hand, and truth, on the other. The present account shows how an epistemic practice has a conceptual connection to truth.

3.3 | *Intermediate conclusions/looking ahead*

What I have tried to establish in this section is that our epistemic practice—like our moral practice—has not one, but (at least) two intimately interrelated functions. The first function is commitment coordination/maintenance of consensus. Just as coordination problems in morality can only be solved by securing consensus around a set of moral practices, so can human society only be formed by the maintenance of doxastic consensus; and a chief role of our epistemic practice is the maintenance of this consensus. Of course, maintenance of consensus cannot be the only function of our epistemic practice, any more than coordination can be the only function of a moral practice. Such practices must also be evaluated along the optimization dimension. In the case of epistemology, this dimension assesses the extent to which the theories, practices, and so on around which consensus is maintained conduce toward truth, understanding, explanation, and so on. Establishment of the first function (consensus maintenance) will be foundational in my argument that knowledge attributions attribute collective statuses. Establishment of the second function is crucial to arguing that epistemic practice remains an *epistemic* practice, and not merely a play of power.

4 | COMMON KNOWLEDGE AND COMMITMENT COORDINATION

The literature on we-attitudes is filled with examples of limited, specialized groups that may be said to know this or that as a group. However, the overwhelming majority of our common knowledge¹⁷ subsists as we-attitudes, because our common knowledge serves as the *common medium* in which our behavior and interactions can occur. You know that birds fly, that water is wet (and drinkable), that the sun is hot and ice is cold, that it is dark at night and light during the day that rain will make one wet, and so forth. Second, you actually *know* these things. Further, others in your community also know these things. Finally—and here is the crucial piece—implicit in our practice is a *mutual recognition* that these things are mutually known. I will offer you water to drink. I will mention that you might need an umbrella, if it is raining. We will plan a sporting event (or other event requiring good visibility) during the day (or where there is good lighting). As Alexander Bird writes, “Common knowledge plays a useful social function, because everyone can rely on other people having the same knowledge, which in turn means that it can guide group action without further discussion” (2010, p. 29). This, I think, is radical understatement: The background of mutual knowledge provides the common medium

¹⁷Like Jane Heal, I am using “common knowledge” in a very non-technical way, to denote widely shared commitment/entitlements, without regard to whether they are actually true. Thus, as Heal writes, “The account we seek is to be of a structure which can accommodate the mediaeval ‘common knowledge’ that the Sun circles the Earth” (Heal 1978, p. 116).

enabling any sort of rational social activity—communication, coordinated action, any type of action that is characteristic of a human community. And in general, this background knowledge is “ready-to-hand”—it exists as assumed, unconscious, as a generally ignored presupposition to anything else we want to accomplish.

4.1 | Shared commitment vs. shared entitlement

A natural objection arises: Cannot the work of social coordination be done by shared *commitment*? Do we really need shared *entitlement*—and therefore shared *knowledge*? If a community shares a commitment to the drinkability of water, the wetness of rain, the coldness of ice, etc.; then will not that produce the proper coordinated behavior, regardless of entitlement to those claims?

No; mere commitment will not reliably produce stable coordination. First, notice how ascriptions of entitlement structure the behavior of agents—and compare this with mere ascriptions of commitment. For example, scientists working at CERN take themselves to be entitled to various commitments concerning the existence and behavior of subatomic particles (photons, neutrinos); thus, such claims form an assumed background for the experiments that are conducted there. By contrast, the scientists there are *committed* to—but do not take themselves to be fully *entitled* to endorse—the Standard Model of particle physics, and this is reflected in their continued efforts to verify this model (e.g., in their efforts to confirm the existence of the Higgs boson).

Shared entitlement differs from mere commitment. If I take you to be committed to the accuracy of a data set, but not entitled to it—perhaps I doubt your methods, but not your conviction—then I will not myself rely on that data set. But since the scientists at CERN take each other to be *entitled* to various shared commitments, they can jointly act on these shared commitments in a way that allows them to test further shared commitments—ones to which they do not yet take themselves to be fully entitled. The scientists at CERN taking themselves and each other to be entitled to certain claims about subatomic particles underlies and enables their cooperation in verifying the Standard Model. Thus, *shared, we-mode entitlement* underlies and enables such cooperation.

This section also marks a crucial step in the argument in moving from talking about collective *attitudes* to collective *statuses*. As noted in my brief commentary above on the notion of a commitment, the concept (in the Brandomian usage) has a psychological-cum-normative sense; it is thus closer to an attitude. But entitlement is clearly a *status*, a status that (say) a commitment can possess. Crucially, as noted above, it is a status that can also be collective or communal.¹⁸ What I wish to argue is that knowledge attributions *always* attribute communal or collective entitlement. Let us continue with the argument that maintenance of consensus requires entitlement (and not merely commitment).

4.2 | Common knowledge: Institutional facts

But even *shared* commitment cannot reliably produce coordination; we need shared *entitlement*. This is because the mechanisms that produce reliable coordination and those that produce entitlement turn out to be identical. This is easiest to demonstrate for institutional facts. (And indeed, most work on common knowledge, and its role in coordination of behavior, focuses on common knowledge of conventions.) Recall our earlier example of linguistic communication using the word “bird.” Our mutual

¹⁸Cf., for example, Kusch (2002, Ch. 5).

commitment that “bird” means *bird*—and the mutual expectation it gives rise to—would be mysterious and unstable if it did not arise out of a shared social practice. But the shared social practice both (a) explains our coordinated behavior and (b) entitles us to thinking that “bird” means bird, since it is in virtue of this social practice that “bird” means *bird* in the first place. Thus, with institutional facts, the mechanism that produces reliable coordination and the mechanism that produces entitlement are one and the same mechanism. Thus, reliable coordination requires—but also generates—*mutual entitlement*. For identical reasons, in ordinary cases of linguistic communication, the following claims also seem plausible—and notice how an account of collective *statuses* (like knowledge) begins to emerge naturally from an account of collective *attitudes*:

- a. I *know* that “bird” means *bird*.
- b. There exists an *expectation* on my part that you *know* that “bird” means *bird*.
- c. There exists an *expectation* on my part that you have an equivalent expectation, i.e., an expectation that I *know* that “bird” means *bird*.

Thus, ordinary cases of linguistic exchanges involve mutual knowledge—knowledge in the “we-mode,” in the strong sense. (Indeed, in the case of institutional facts like linguistic facts, it seems like there is a relation of *conceptual priority*—I can only know the institutional fact because *we* know this fact. It is (roughly) agreement in use that institutes “bird” as meaning *bird*, and so my knowledge of the meaning of “bird” is parasitic on the knowledge of English-language users generally.) Thus, the structure of linguistic communication requires a structure of *shared statuses*—namely, *shared entitlements* regarding the meanings of various terms, and no doubt regarding other semantic, pragmatic, and conversational norms.

In general, being members of a *society* means inhabiting a common social reality. And this social reality must be intersubjective, which entails that the process of socialization involves inculcating into each of us a common store of knowledge—knowledge of institutions, roles, norms, objects, significances, and so on. Only thus are we able to inhabit a common world, coordinate our actions, and interact and communicate *as* social beings. As Berger and Luckmann write in their landmark work, *The Social Construction of Reality*, “I live in the common-sense world of everyday life equipped with specific bodies of knowledge. What is more, *I know that others share at least part of this knowledge, and they know that I know this*. My interaction with others in everyday life is, therefore, constantly affected by our common participation in the available social stock of knowledge” (Berger & Luckmann, 1966, p. 56, emphasis added). Alfred Schutz (1962), for example, writes about “typifying constructs,” which we use to “anonymize” people—to abstract from their particularity and assign them specific *roles, motives, and functions*—thereby allowing us to predict their behavior and therefore interact with them as types, rather than as individuals. Thus, we do not need to know the identity of the letter carrier to know that putting a letter in the mailbox will result in its being collected and delivered; nor does she have to know my identity (or that of the recipient) to know what to do with the letters she collects. But there are mutually interlocking attributions of role, function, and motive which allow interactions among strangers in society to take place; mutual knowledge of these “typifying constructs” underlies our division of labor in society:

I take it for granted that my action (say putting a stamped and duly addressed envelope in a mailbox) will induce anonymous fellow-men (postmen) to perform typical actions (handling the mail) in accordance with typical in-order-to motives (to live up to their occupational duties) with the result that the state of affairs projected by me (delivery of

the letter to the addressee within reasonable time) will be achieved. 2) I also take it for granted that my construct of the Other's course-of-action type corresponds substantially to his own self-typification and that to the latter belongs a typified construct of my, his anonymous partner's, typical way of behavior based on typical and supposedly invariant motives. ("Whoever puts a duly addressed and stamped envelope in the mailbox is assumed to intend to have it delivered to the addressee in due time.") 3) Even more, in my own self-typification—that is by assuming the role of a customer of the mail service—I have to project my action in such a typical way as I suppose the typical post office employee expects a typical customer to behave. Such a construct of mutually interlocked behavior patterns reveals itself as a construct of mutually interlocked in-order-to and because motives which are supposedly invariant. The more institutionalized or standardized such a behavior pattern is, that is, the more typified it is in a socially approved way by laws, rules, regulations, customs, habits, etc., the greater is the chance that my own self-typifying behavior will bring about the state of affairs aimed at. (Schutz 1962, pp. 25-6)

Thus, an indefinitely large range of shared statuses (knowings-that and knowings-how) underlie the very possibility of human community.¹⁹

4.3 | Common knowledge: Non-institutional facts

Further, these claims are also plausible with respect to a large range of non-institutional facts. Many such facts are generally known, but are known (again) due to our general inculcation into a social practice—not due to first-hand experience or observation. Our knowledge that the earth is round, that dogs are mammals, that there are many countries in the world (which we have never visited), that humans have walked on the moon, and so on—these are claims that we generally know, and which can be assumed as mutual knowledge in a conversation (special circumstances aside), but which generally we know as a result of evidence that is possessed by the community at large. Thus, again, the mechanism that produces entitlement (that these commitments are produced and sustained by the larger community, which obtains and possesses evidence which we as individuals do not have the resources to acquire) is also the same mechanism which maintains consensus, thereby enabling

¹⁹An anonymous referee has raised the concern that this move makes knowledge relative to societies. I would argue that my approach—which is a thoroughly "social practice" approach to knowledge—has the resources to resist accusations of relativism. Nelson, however, admirably summarizes the argument as follows: "There is an additional and important benefit to construing evidence broadly and focusing on communities: namely, that these preclude the claim that cases like that which we have considered are examples of incommensurability—or, in some other way constitute 'evidence for' judgmental relativism. Advocates of man-the-hunter theory and feminist critics disagree about many things... But they do not disagree about everything; they share a larger body of knowledge and standards that includes physical object theory, a heliocentric view of the solar system, and the view that humans evolved and that their activities were factors in that process. Hence, members of these groups can discuss (and disagree about) the significance of 'chipped stones' without any lapse in conversation and use other aspects of the knowledge and standards they share to evaluate the conflicting claims. The flip side of the point is this: although the knowledge and standards currently at issue are community specific, feminist communities and science communities both overlap (consider feminist primatologists) and are themselves subcommunities of larger communities—a fact that, along with the changing social relations that made it possible, has enabled feminist science criticism and feminist knowledge more broadly." (Nelson 1993, pp. 147–8). See also (Koons 2019, chapter 4); see, also, Wolf (2012).

cognitive and behavioral coordination.²⁰ Inculcation into a shared, intersubjective world—the prerequisite of rational community—requires that each of us be supplied with a stock of knowledge, not merely of institutional facts (as described in IV.2), but also basic knowledge of the world we inhabit, to allow for sufficient shared meanings and shared relevances for basic communication and shared understanding. It is in virtue of this inculcation that our commitments have the status of an entitlement—one that is of necessity a *shared entitlement* not merely because of how it produces, but because of how this entitlement is supposed to allow us to inhabit a shared reality.

What about our beliefs whose entitlement less plausibly relies on a social practice—for example, our belief that water is drinkable, or that rain is wet, or that birds fly, and so on? Now, it follows from what I have said above that entertaining or expressing such beliefs or commitments involves holding or expressing a commitment in the we-mode in the strong sense. But one might say that to say this is merely to make a claim about what is involved in having the concepts necessary to formulate beliefs or commitments using intersubjective language or linguistic concepts. A persistent criticism of the Sellarsian project is that it confuses conceptual priority—or what is necessary for the formulation of various beliefs or mental states—with epistemic priority. So we may grant that none of this yet demonstrates that such beliefs involve collective entitlement or collective knowledge.

I will begin, though, by reiterating the previously made claim: Mere commitment is not sufficiently stable to maintain cognitive and behavioral cooperation. We inhabit a shared world, and it would be odd to claim that one can coordinate cognitive and behavioral activity with someone who has *ungrounded commitments* about that world. It would seem pure chance that these commitments turned out to be accurate, or that they matched with mine, or that I could successfully coordinate my behavior with this person on the basis of their commitments. Thus, commitment by itself is not enough—the person must have *entitled* commitments. I must take them to *know* the world they inhabit—as they must take me also to know the world. There must be a mutual expectation of this knowledge—the element of mutual expectation will confer the status of (mutual) entitlement on these commitments. And so even in the case of these simple beliefs, commitment is not sufficient to sustain cooperation. There must be mutual (we-mode) knowledge.

And of course on the pragmatist order of explanation, overt attributions of normative statuses merely make explicit what is already implicit in practice. So if the knowledge implicit in practice is we-mode, then *a fortiori* so are explicit attributions of such knowledge.

4.4 | Intermediate conclusions/looking ahead

To argue that the maintenance of consensus is specifically a function of our *epistemic* practice, I needed to argue that it is a specifically *epistemic* status that underlies commitment coordination. Thus, I began by arguing that commitment coordination can only reliably be accomplished by shared entitlement, not by shared commitment. I then argued that common knowledge of institutional reality—that very reality which in a fundamental way constitutes human society—is underlain by shared

²⁰The skeptic will argue that shared commitment can produce reliable coordination—brains in vats (who lack justification for their beliefs) lack justification, but successfully coordinate. I would say that what examples like this show is the bankruptcy of conceptions of justification that make it radically external to all possible social practice. Within their vat-world, the brains *do* make attributions of entitlement, and this practice allows them to maintain consensus—something they could not do if they did not have their own epistemic practice which functioned more or less successfully, relative to the vat-world. The fact that we, from our God's-eye perspective, see that their world is ultimately an illusion, changes nothing about the fact that *their* epistemic practice serves the same function as does *ours*, and that *our* failure to attribute entitlement to their claims can have no bearing on their actual lived practice.

entitlement (and, indeed, shared knowledge). I then extended the argument to common knowledge of the non-institutional variety.

However, the claim I wish to defend in this paper is the radical one that *all* knowledge attributions attribute a collective status. This is easier to prove with regard to entitled beliefs that are actually communally held. But what about entitled beliefs that are not widely held—entitled beliefs that are only held by one, or a handful, of individuals? In other words, what about *uncommon knowledge*? I need to make the case that even when we attribute knowledge in these cases, we are still attributing a collective status—we are still attributing knowledge to the entire epistemic community. It is to this task that I now turn.

5 | UNCOMMON KNOWLEDGE

I conclude that in the great majority of cases where knowledge is properly attributable, it is attributable as a we-mode status in the strong sense. It certainly does not follow that knowledge is an *inherently* communal notion, or an inherently communal status. And indeed, common accounts of collective attitudes treat we-mode knowledge as a special case, only existing when various (often stringent) conditions are met. Tuomela, for example, holds that group knowledge implies that members of a group have we-commitment to the content of the relevant belief, and argues that “We say that a person is we-committed...to a content *p* if and only if he is committed to *p* and believes that the others are similarly committed to *p* and that this is mutually believed in the group.” (Tuomela, 2004, p. 113). Tuomela also emphasizes the need for a shared or joint *justification*.²¹ Thus, it seems that Tuomela (like other authors) makes little room for us to speak of knowledge in the we-mode where members of the group by and large do not share the commitment in question or do not share entitlement to that commitment.

An advantage of the present account is that it treats knowledge not as a collective attitude, but as a collective *status*. This difference might make it easier to explain how we could attribute knowledge to a community in the absence of a widely shared commitment to a particular proposition.

5.1 | Uncommon knowledge attributions in the we-mode

Consider the point of attributing epistemic entitlement to others. The point can be seen by again contrasting such attributions with attributions of commitment. Doxastic commitments can be incompatible with each other. But to know that someone has a commitment incompatible with one of your own commitments does not *in itself* place any sort of *normative* constraint on you.

Attributions of entitlement, however, have an essentially social normative element. To attribute entitlement to someone with respect to a commitment, or a range of commitments, is to afford a particular conversational status to some subset of this person's utterances. To attribute entitlement to Smith with respect to *p* is to treat her utterances that *p* as authoritative—not just for Smith, but for me, and for other members of my epistemic community. It is thus to recognize *our* entitlement to *p*. For example, if I say that Smith is an expert entomologist—and notice that there are many ways to attribute

²¹Tuomela has revised his account somewhat in recent years, but he still seems to make little room for we-mode knowledge that is not by and large shared by the group. For example, one element of his 2011 account of group knowledge is, “A group's knowing that *p* qua a group entails that the group must have accepted that *p* as true and that the group is justified in accepting that *p*. Group acceptance entails that the group is committed to *p* as a group. Group justification involves here that the group must have a good epistemic reason for its acceptance of *p* as true” (Tuomela 2011, p. 89).

epistemic status to individuals; the use of the words “knows,” “justified,” and other epistemologists’ favorites are probably not the most common—then I convey that non-experts such as myself ought to defer to Smith when she, say, identifies a particular specimen as a rhinoceros beetle, or tells me that velvet ants are not really ants. Furthermore, this allocation of epistemic authority further serves the point of consensus maintenance in society. Attributing a status like this to Smith involves undertaking a commitment to defer to Smith on this issue: We need to make *our* beliefs conform to *hers*. Or if I claim entitlement to a particular commitment—by saying, “I know that Smith was here this morning,” or “I saw her with my own eyes,” or something similar—then I am saying, in essence, that you ought epistemically to defer to me. The community ought to make *its* beliefs conform to *mine*. An entitlement attribution involves, therefore, a claim about whose commitments should be regarded as controlling with respect to a certain set of questions. Indeed, the epistemic division of labor (and the corresponding inegalitarian distribution of epistemic authority) only works if we defer to certain individuals, groups, or organizations—if we treat their entitlements as *prima facie* governing for our commitments regarding certain issues. This involves, then, treating these entitlements as issuing licenses (and obligations) for the entire discursive community. They are, as it were, *communal entitlements*.

This point about deferral to expert authority allows me to tie our present discussion back to our earlier discussion. As John Hardwig (1985) points out, much—perhaps most—of what we know, we know on the authority of others, whether these others are scientists, economists, car mechanics, trusted news sources, physicians, laboratory technicians, teachers, textbook authors, or whatever.²² Recall, again, our pragmatist order of explanation: Explicit normative claims merely say what one is committed to doing by acting in this way or that way. Thus, by deferring to a scientist, or a textbook author, or a car mechanic, one is implicitly acknowledging this person's entitlement to a certain claim or set of claims. Thus, inhabiting a common world, structured by sufficient common beliefs to allow anything like a cohesive community, requires broadly similar epistemic practices—for example, similar practices of deferring to similar authorities on similar questions. Epistemic practices maintain this “common world” of commitment by recognizing certain sources as authoritative for commitment formation and maintenance, and others as not. Thus, the fragmentation of an epistemic practice (e.g., when scientists are no longer widely or universally recognized as authoritative on matters of climate change or vaccine safety) represents a failure of the smooth functioning of epistemic practice to do what it normally does, quietly in the background—namely, maintain consensus and coordinate doxastic behavior. And while the community can survive a certain level of epistemic breakdown, it cannot survive extensive breakdown (much less total breakdown), as then members of the community would cease to inhabit a common world.

Here, then, is the failure of individualistic epistemic programs laid bare. To treat entitlement or knowledge as a property of *individuals* is to act as though the coordination of beliefs in society is something that happens fortuitously, or as a side effect of good epistemic practice by individuals. Coherentists, for example, face the traditional challenge of internally consistent belief systems that do not match the world. But perhaps the deeper problem is of different individuals having internally consistent belief systems that do not match *each other's* systems. What mechanism guarantees coordination of belief systems across members of society? There is nothing intrinsic to coherentism or foundationalism, or other individualistic epistemological theories that would vouch for such coordination.

²²Hardwig, of course, has his critics—for example, Fuller (2002, pp. 277–80). But the point that a layman might have a good reason not to defer to a particular expert (or a particular set of experts) in no way undermines the basic fact of our radical epistemic dependence upon a huge range of people who—for this or that reason (including expertise)—are more epistemically privileged than we with respect to this or that subject matter. For a set of interesting empirical arguments that strategic copying and learning from others is key not only to the evolutionary success of individuals, but also to the success of the species as a whole, see Laland (2017, especially Chapter 3).

But a practice of belief regulation is useless if it does not produce belief coordination—if it fails to maintain consensus, to make sure members of a community are on the same page, doxastically speaking. This is the basic prerequisite for the existence of any rational community. So we have to treat entitlements as communal rather than individual.

Support for this conception of epistemic practice can be found in Kukla and Lance (2009). An entitled declaration has an output that is, in their terms, “agent neutral”—that is, it “issues reassertion and inference licenses that are not indexed to any specific agent or kind of agents” (2009, p. 26). Naturally, this does not mean that everyone will now possess the relevant reassertion license—due to ignorance or other causes, many will not. But the license becomes an epistemic resource now *available* to the community. Thus, “the actual agent-neutrality of the output goes along with its universality as a *regulative ideal*. It is, as it were, a claim *for everyone*, which strives to contribute to the bank of public knowledge shared by the discursive community” (2009, p. 26).

Thus, to speak of *what is known* is to make a claim about what epistemic resources are *available to the community*. Thus, “There is an important sense in which once Daniel Mazia discovered that mitosis is a form of reproduction...it became true that *we know* [this fact], even though not everyone in the community knows [it]. As a textbook might put it, ‘We have known since 1951 that mitosis is a form of reproduction’ (a statement most assuredly not true of the authors of this book)” (Kukla & Lance, 2009, p. 27). Such knowledge claims are of course false if read distributively—i.e., as attributing a bit of knowledge to *each* member of the community—but true if read as a claim about a piece of knowledge that is possessed by the community at large and to which all members of the community are in some sense accountable.

The thesis that knowledge claims can plausibly be seen as community-wide, and hence making claims about what *we* know, is supported by reflection on the normative consequences of such a claim. As Kukla and Lance note, “there is a trivial sense in which not knowing something that is true is a defect” (2009, p. 28). We are not omniscient, after all. But if someone makes an entitled declaration that-*p*, then there is now a much stronger sense in which failure to know that-*p* is a defect: “In uttering a justified declarative, a speaker...adds it to the public bank of knowledge—it is now part of what we know... An individual’s failure to know what her discursive community knows puts her in a position of discursive deficiency—susceptibility to legitimate correction by others—that is concretely different from a mere failure of omniscience” (2009, p. 28). Just as importantly, if Smith knows that-*p*, then this precludes anyone else in the community from knowing that not-*p*. This is another strong sense in which the knowledge is communal: “The achievement of an agent-neutral entitlement always precludes entitlement by anyone else to any claim incompatible with the claim in question: Once we discover planets orbiting other stars, it can never be the case that anyone can be properly entitled to the belief that there are no such planets (although, given incomplete knowledge, someone may still have good reasons for such belief)” (Kukla & Lance, 2009, p. 29).

The present account makes sense of examples like Bird’s Case of Dr. N. (Bird, 2010). In this case, Dr. N. carries out research and makes a discovery, publishing it in the *Journal of X-ology*. The research garners little attention; Dr. N. dies and the result is forgotten. Later, Professor O. discovers this publication and cites it in a widely read article. Bird asks, “Was Dr. N.’s discovery part of scientific knowledge? I argue that it was so *throughout* the period in question” (2010, p. 32). If knowledge is a *mental state* or *attitude*, this cannot be; but if knowledge is a *status* (as I have argued), then Bird’s conclusion can be sustained.

That knowledge attributions attribute a collective status is demonstrated by reflecting on some of the specific normative features of epistemic entitlement. As Gilbert (2014) has famously argued, we do not ordinarily have the standing to rebuke people for violating personal commitments. If I make a commitment to go to Café Caffè for lunch, but fail to follow through, then you do not have the standing

to rebuke me for failure to follow through on my commitment. But if *we* make a commitment to meet at Café Caffè for lunch, and I fail to show up, then you do have the standing to rebuke me: Joint commitments give one the standing to rebuke others for violation.

But notice the following fact about making assertions. (Here, I am assuming that in the standard case, the norm of assertion is that the speaker at least implicitly takes herself to be entitled to the content of her assertion; an argument for this can be found in [Koons & Wolf, 2018].)²³ Making a non-entitled assertion gives most anyone in the discursive community the standing to rebuke you.²⁴ It may be unwise or socially inappropriate to rebuke you. But as a specifically *epistemological* point, anyone (or *at least* anyone who possess the relevant epistemic authority) has the standing to rebuke you. Thus, if the actor Jim Carrey were to post something on social media claiming that vaccines are the cause of autism, literally anyone has the standing to rebuke him in that forum, or any other appropriate forum. By engaging in the act of assertion, you make yourself part of a discursive community with the rest of us, and your commitments are thereby we-mode commitments. More importantly, the entitlement you claim for your commitments thereby aspires to be authoritative for this community. Making a claim involves implicitly claiming a we-mode entitlement—one that makes a claim on other members of the discursive community—a claim that others might reject. Thus, those who reject your entitlement, *and the claim that it makes upon them*, have the standing to rebuke you—to dispute your entitlement, and the authority you are claiming for your commitments.

Notice carefully what we are rebuking in this case. We are in a sense rebuking Carrey's commitment—he has the wrong commitment. But examining more closely the actual structure of rebuke, we can see that at the same time, we are rebuking his *entitlement* to this commitment. Criticism of Carrey will point out (for example) the fraudulent nature of the original study purportedly linking the MMR vaccine to autism, the *Lancet's* retraction of that study, the exhaustive subsequent research that failed to establish any such link, and so on. The fact that we have the standing to rebuke discursive partners for the entitlement they implicitly claim in making assertions is strongly suggestive that this entitlement has we-mode status—that it purports to bind us all.²⁵

We can see, then, a crucial respect in which my account of knowledge attribution differs from Kusch's—and a sense in which Kusch's communitarianism is not well-served by his account. A common feature of theoretical accounts of group attitudes and statuses is that such accounts presuppose that a group attitude and status requires that some corresponding belief or intention be held by a majority of the members of the relevant community. And this seems to be Kusch's attitude about the status of knowledge. He writes, for example, “As far as the social status of knowledge is concerned, to acquire this status for one of my beliefs is for this belief to be shared by others” (2002, p. 147). This view has important consequences for his account of knowledge attribution: If I attribute knowledge to myself before my belief is widely accepted, “In thinking of beliefs as knowledge I am making a prediction as to how they will fare” (2002, p. 148).

²³A number of authors have argued that there is an epistemic norm governing assertion; see, for example, Goldberg (2015).

²⁴I am leaving open the possibility that this standing might reflect some discursive-community relativity. Thus, I will seldom (if ever) have the standing to rebuke a particle physicist who makes a non-entitled assertion about particle physics. But anyone in *that* discursive community has, in principle, the epistemic standing to do so. I am obviously, for simplicity's sake, setting aside important issues of epistemic injustice.

²⁵Bird (2010) argues that members of “organic” groups—like a community of scientists—do not have the standing merely to rebuke contradictory commitments *qua* commitments. When rebukes to other scientists are issued, these “criticisms...will be primarily epistemological and are distinct from the rebuke that Gilbert refers to, which arises from the breaking of a *commitment* the individual makes to the group” (2010, p. 38). I am largely in agreement with Bird; except I think the fact that we have standing to rebuke others' assertions indicates that the entitlement they (implicitly) claim for such assertions is we-mode entitlement.

This cannot be correct. As I noted previously, a common claim among standpoint epistemologists is that marginalized groups often have certain kinds of epistemic privilege—particularly regarding the types of oppression and injustice they face at the hands of the dominant group. A member of an oppressed group could self-attribute knowledge about such injustice while at the same time pessimistically predicting that this knowledge will never be accepted by the larger community. (One could try to relativize knowledge to the smaller community of marginalized people, but Kusch seems eager to reject such relativizing moves; see particularly the start of Chapter 11 of his 2002). Also, Kusch's particular account of knowledge attributions would render nonsense ordinary phrases such as, “When I find out, you'll be the first to know.” On Kusch's account, one in principle cannot be the first to know anything.

More promising, I think, is to focus on the semantics of knowledge attributions. As I noted in the introduction, this account of knowledge takes its inspiration from Sellars's account of moral judgments. On such an account, moral judgments express *we-intentions*. Crucially, a moral judgment takes this form *even if you know that nobody else shares or agrees with your judgment*. Thus, it is a theory about *the form of such judgments*; it is not a sociological theory about whether such judgments are or are not shared. The point is that such judgments purport to bind the community of rational beings—and purport to do so whether or not these judgments are acknowledged by other rational beings.

So, too, with knowledge attributions. Because of their role in maintaining consensus, and because we must see epistemic communities as central to a social epistemology, we must see knowledge attributions not as *predicting* adoption of a belief, but as *legislating* adoption of this belief. Thus, such attributions are always prescriptive in character, and always prescriptive toward the larger epistemic community.

All of this naturally raises the problem of how knowledge attributions are supposed to be true if they are not shared by others. It is to this problem that we now turn.

5.2 | Attributing statuses to groups (even when they reject them)

We have seen that cases of common knowledge involve a mutual attribution of entitlement that elevates this to entitlement in the *we-mode*, in the strong sense. But if a knowledge- or entitlement claim is not widely recognized or reciprocated, how can the attribution of such count as a strong *we-mode* attribution?

The problem is general: How can a status properly be attributed to a group if the members of the group generally would reject attribution of this status? Writers such as Tuomela generally focus on cases where group members not only hold the commitment in question, but mutually believe that others hold this commitment; and entitlement to the commitment is similarly shared. This focus is explained primarily by Tuomela's goal to explain the role group knowledge plays in group action—which is, as already noted, best-facilitated by widely shared commitments and entitlements.

However, less widely shared commitments and entitlements can also count as group knowledge. Consider an institutional example.²⁶ Tuomela notes that even in the case of “strongly social concepts,” like that of being a leader (i.e., in an organization), there are “constraints for correct collective acceptance” (2002, p. 147). Thus, if Smith is elected leader of Organization, but it turns out that a quorum was not present at the meeting, then the election was not valid, and he is not in fact the valid leader, *even if people take him to be so*. Therefore, someone could challenge his status as leader, demand a re-vote, etc. Thus, I may be the only person in Organization to hold the attitude.

²⁶Drawn from (Koons 2019, chapter 4).

B: I Believe_{Org} [Smith is not the valid leader]

and yet my group attitude is correct—I am *entitled* to this attitude—though this attitude is not shared by the organization as a whole.

Notice a few things about this example. First, my possession of belief B, which I hold *qua* member of Organization, presupposes *massive agreement* with other members of organization—agreement about the structure of Organization (e.g., that there is a leader, about the relation of the leader to ordinary members), about the bylaws (voting procedures, procedures for conducting meetings, rules of quorum, rules for moving and seconding motions, etc.), and so on. Thus, I am only able to *dissent* *qua* member of organization *because we already share so many attitudes characterizable as “Attitude_{Org}.”* My belief B—which counts as a Belief_{Org}—is a group attitude *in virtue of the rules constitutive of Organization*. And these are rules to which both I and other members of the organization are at least in principle committed—and in virtue of which we can even entertain beliefs characterizable as Belief_{Org}.

This is true not only in institutional cases, but wherever we dispute over norms. A dispute over a semantic norm (e.g., the proper meaning of “beg the question”) presupposes massive agreement, without which we could not even be characterized as having a *semantic dispute*. Given the constant evolution of epistemic standards, and their significant contextuality, when Jones makes a knowledge claim, Jones is making one that is justified in light of a certain socio-historical set of epistemological standards. That is, Jones is judging that she knows by the standards of her discursive community. She might recognize that members of her community do not in fact share the commitment in question; she might recognize that they also will not recognize her entitlement to this commitment. But in saying she knows that-*p*, she is posing this claim as one her community is obligated to recognize, based on shared commitments. So when Galileo first observed the phases of Venus, which he knew to be incompatible with the Ptolemaic model of the universe, he knew that Venus orbited the sun. Further, his knowledge relied on a broadly shared set of theoretical tools along with shared epistemic methods and standards—for example, centuries of observations of the motions of heavenly bodies (along with sophisticated models of same) and standards involving observation of the motion of heavenly bodies and methods of relating them to standard models of the universe. And if Galileo had no understanding of the optics underlying telescopes, nevertheless it was widely understood at the time that such devices could be constructed “by putting a weak convex and strong concave spectacle lens into a tube and stopping down [i.e., reducing the size of] the aperture” (Van Helden, 2010, p. 187). Thus, even if others did not at that moment accept the claim that Venus orbited the sun—and even if they denied it, on being presented with the evidence—Galileo had the (epistemic) standing to rebuke them, as they were in violation of communal standards of evidence.²⁷ (I should note that epistemic standing does not imply political or any other kind of standing. Nor does it imply that one is afforded epistemic authority by one’s discursive partners. For example, standpoint epistemologists often point out that certain types of epistemic privilege accrue to oppressed and marginalized people precisely in virtue of their oppression. See, for example, Elliott [1994].)

To return to our Organization example, of course, those who disagree with me also hold attitudes properly characterized as Belief_{Org}. Thus, while I hold the following belief:

B: I Believe_{Org} [Smith is not the valid leader]

²⁷Thus, this account should not preclude a knowledge attribution to an individual who makes a discovery through her own methodological innovation; for the innovative method will in all cases be grounded in a mountain of shared theory and standards. The innovator will have *reasons* for thinking this innovative method is accurate; and these reasons will essentially be shareable.

others in the organization hold the contrary belief:

B': I Believe_{Org} [Smith is the valid leader]

Their beliefs are properly characterized as Belief_{Org} because even though they are mistaken (lacking the same factual foundation as my belief), they are still grounded in the rules constitutive of Organization. The same is true with knowledge claims: The mere fact of disagreement is compatible with our making collective claims from the standpoint of the group. Indeed, the very *possibility* of constructive disagreement—and subsequent resolution of this disagreement—presupposes the possibility of making claims with this structure.

A further, crucial point to be gleaned here is that I am *entitled* to B. And this fact has important consequences. I am entitled to B *qua* member of organization, specifically because of the way in which my entitlement arises. I am entitled to B because of:

- (i) The rules that are constitutive of Organization, and
- (ii) Various facts which I have observed—such as the number of voting members of Organization present at the meeting where Smith's leadership was voted upon—facts which have no specific relevance apart from the rules constitutive of Organization. Indeed, many of these are facts that *cannot be stated* (e.g., “voting members of Organization,” “meeting,” “quorum,” “Smith's leadership was voted upon”) without the conceptual and institutional resources provided by the rules constitutive of Organization.

Thus, my entitlement has not merely an epistemic, but also an *institutional* character. It can be characterized as Entitlement_{Org}. It is such *because* of (i) and (ii). Because of (i) and (ii), my entitlement is binding on other members of Organization; and this also demonstrates the collective nature of this entitlement. My entitled declaration of B obligates others to recognize the invalidity of Smith's election, and to take the appropriate steps (as designated by the rules constitutive of Organization) to rectify the situation. Again, to say this is not to say that others will recognize this obligation, or even that it is wise for me to go around declaring B, or that Smith will not go on being recognized as the leader of Organization. I am merely making a point about the normative consequences of an entitled declaration, and how even non-shared entitled declarations have collective consequences, and therefore a collective character. Thus, we can further say of B:

I am Entitled_{Org} to Belief_{Org} [Smith is not the valid leader]

Or even:

I Know_{Org} [Smith is not the valid leader]

Thus, in making a knowledge claim, one makes a claim that purports to bind one's discursive community; the structure of such claims precludes a strictly individualistic reading of them. In claiming that she knows that-*p*, Jones is also committed to the claim that failure to acknowledge her entitlement to *p* constitutes an epistemic defect or failure on the part of other members of her epistemic community. And in recognizing their epistemic defect, Jones is (a) implicitly recognizing that they are members of the same discursive community, and therefore members of an epistemic “we”; (b) implicitly recognizing that as such, they are bound by and beholden to the same set of epistemic norms; and (c) implicitly making the

“knows” claim binding on all other members of her discursive community, making it therefore not a mere description of Jones herself, but instead a *communal statement of entitlement and obligation*.

5.3 | Giving a more formal characterization of communal statures

To offer a more formal account of an entitlement that is (in Tuomela's strong sense) a we-mode entitlement, even if it is not widely shared by the relevant group, we will begin with something like Tuomela's notion of an *ethos*. Tuomela defines an ethos as follows: “The *ethos* of a group *g* in its strict sense is defined as the set of *constitutive* goals, values, beliefs, standards, norms, practices, and/or traditions that give the group motivating reasons for action” (Tuomela, 2007, p. 16). So the ethos of Organization will be—at least in part—the rules constitutive of Organization, such as the charter, bylaws, and other rules that make organization exist, and in virtue of which we can be said to possess anything characterizable as $Attitude_{Org}$. Similarly, the ethos of a rational community will be the set of norms, practices, and standards (including semantic and epistemic norms) that bind it together as a rational community. Group cohesion—indeed, group constitution—requires “the collective acceptance of E [the ethos] (involving collective commitment)” (Tuomela, 2007, p. 32). To put this in the pragmatist terms of this essay, we might say that I have am committed to an ethos *qua* member of a group iff:

- a. I am committed to the ethos;
- b. I have an expectation that other members of the group are committed to the ethos; and
- c. I have an expectation that other members of the group have an equivalent expectation; i.e., an expectation that I am committed to the ethos.

By “I am committed to the ethos” I mean simply “I am committed to the rules, procedures, practices, etc., constitutive of the ethos.” I am eliding a number of no doubt thorny difficulties, as practically speaking, no two agents ever have *exactly* the same set of commitments regarding any set of norms, be they semantic, moral, epistemic, or whatever. But these incommensurability worries are not unique to this type of account; and in any case, what I am trying to offer here is an account of how an attitude can still count as a group attitude even if it is *not* shared—simply on the basis of the overwhelming amount of the ethos that *is* shared. So this account relies merely on the fact of rational agents sharing a “form of life,” even if precisely what is shared will vary from individual to individual.

Given that agents share an ethos, what is it for a non-shared attitude based on this ethos to be a we-attitude in the strong sense? For example, a non-shared commitment that-P (say, that Smith is not the valid leader of Organization) will be a strong we-mode commitment iff, in addition to (a)-(c) holding:

- (d) I have a commitment to P, which commitment is epistemically based on the ethos (plus perhaps on certain other factual claims to which I am entitled).

Clause (d) requires some commentary. I think the above phrasing is more felicitous than the following candidate:

- (d') A commitment on my part that the ethos (plus perhaps certain other factual claims to which I am entitled) entitles me to P.

The alternative clause (d') is too meta-epistemological, in that it requires not merely a commitment to P (and one based on the ethos), but a commitment about what entitles me to P. It is enough, say, that I believe B (that Smith is not the valid leader of Organization) based on the ethos, plus factual information such as the number of people attending the meeting where the election was held; I should not be required to believe that I believe B *because* of the ethos and this further factual information.

The resulting notion of an attitude qualifies as a we-attitude in the strong sense.²⁸ It is certainly not a strictly *individual* commitment; it is one I hold *qua* member of Organization. But it is stronger than what Tuomela would call a “weak we-mode” attitude. Consider Tuomela's discussion of weak we-mode action. Tuomela writes that acting as a group member in the weak sense “is based on the group members’ sharing the ethos but being only privately (and not collectively) committed to it” (2007, p. 29), and gives the following example: “Thus there may be a group of persons who grow flowers in the village commons and intend to make their small village look beautiful in this and perhaps other ways. Each of them is only privately committed to making the village beautiful, and they mutually know or believe that the others are similarly committed” (2007, p. 29). But my belief B is not like the villagers’ commitment regarding the village commons. For example, the key difference between weak and strong we-mode action for Tuomela is that strong we-mode action involves acting for a group reason. The individual villagers are not beautifying the village commons *because* of a joint commitment to do so, even if in a sense they constitute a group and as such have a joint commitment to do so. They act on *private* commitments. But my belief B is not held for private reasons—it is held for reasons essentially tied to the ethos of Organization (and indeed for reasons that cannot even be *articulated* without reference to the ethos, or constitutive rules, of Organization). It is a belief I hold *qua* member of Organization. It is essentially a group belief. It is much more plausibly characterized as an attitude in the we-mode in the strong sense, rather than in the weak sense.

Next—and perhaps most importantly—to the extent that commitment P is epistemically based on the ethos (plus perhaps on certain other factual claims to which I am entitled), commitment P represents an *entitlement* that I have achieved. And just like Belief_{Org}, such entitlement is a *we-mode entitlement*, as its epistemic status derives from the ethos (plus other facts whose relevance and status is determined by the ethos—*no* facts have an interpretation *except* in light of an ethos); and my entitlement to P places corresponding obligations on other members of the community, as discussed in section V.1 above. Thus, I can conclude with the following definition: I will have strong we-mode entitlement to a non-shared commitment that-P (say, that Smith is not the valid leader of Organization) iff, in addition to (a)-(d) holding:

- (e) The ethos (plus perhaps on certain other factual claims to which I am entitled) *entitle* me to commitment P.

But because we, *qua* members of a group, are mutually committed to the ethos; and because my commitment P counts as a we-mode commitment; and—most importantly—because per clause (e) I possess entitlement_{group} to P, this entitlement extends to other members of the group. It is binding on them, whether they recognize it or not.

²⁸I earlier denied this (Koons 2019, chapter 4) and was willing to concede that such non-shared we-attitudes could be we-attitudes in the weak sense. Ronald Loeffler convinced me that this concession was a mistake.

5.4 | Final conclusions

We ended section IV with the conclusion that common knowledge is underlain by shared entitlement. In section V, I extended my argument to the more radical conclusion that *every* knowledge claim attributes a communal status—every attribution of knowledge or entitlement implicitly has the form, “We know that-p” or “We are entitled to believe that-p.” This follows from the way in which such claims purport to bind all members of the epistemic community (and from other features of such attributions, such as how non-entitled assertions give anyone in the epistemic community the right to rebuke the asserter). In sections V.2 and V.3, I argued (based on an institutional analysis) that such attributions can be true even when the commitment in question is not widely shared in the community—even when most members of the community would explicitly reject the commitment in question.

Thus, I have reached the conclusion of the argument: All attributions of knowledge or entitlement attribute a we-mode status, not an individual status; and as such, there can be no research program in epistemology that is not radically social.

6 | CONCLUSION

Traditional theories of justification have often focused on necessary and sufficient conditions for an *individual's* belief to count as knowledge. Even contemporary social epistemology often focuses on how the knowledge of *individual agents* depends on various social factors. To the extent that these projects focus on epistemology primarily as a tool for epistemically assessing individuals, they miss out on the chief function of our actual, lived epistemic practice. This practice exists largely to parcel out epistemic statuses—such as entitlement, epistemic authority, and the like—in a way that maintains a common doxastic world, thereby allowing for coordinated behavior, the accumulation of knowledge, and the construction of our social world. As such, epistemic statuses are in the first instance social statuses; our attributions of such statuses to this or that individual are parasitic on these wider social statuses and communal goals.

Finally, as noted in the introduction, the present account sidesteps debates in social epistemology about the ontology of group minds. Knowledge claims attribute a normative *status*; such claims are not descriptive, and we thus do not need to fit the “normative facts” they allegedly describe into our ontology. This is an advantage of expressivist accounts generally and was a chief motivation for Sellars's expressivism. It is true that the account of collective statuses on offer here relied heavily on a pragmatist account of collective attitudes; there can be no attribution of collective statuses without an ethos consisting of shared attitudes. However, as I noted at the end of II.1, an advantage of the present account is that while acting on norms (such as linguistic norms) presupposes—as a *conceptual* matter—the existence of group attitudes (such as collective commitments), purely as an *ontological* matter such collective commitments presupposes nothing beyond individually held attitudes of commitment and expectation. Thus, we are able to offer a pragmatist account of collective attitudes that does not commit us to anything ontologically suspect, such as group minds.²⁹

ACKNOWLEDGMENTS

I am grateful to the National Endowment for the Humanities for funding my participation in the workshop “Philosophical Responses to Empiricism in Kant, Hegel, and Sellars,” at the University of New Hampshire (19 June – 19 July 2019), during which this project was initially developed and

²⁹As noted earlier, I argue for this at greater length in (Koons 2019, chapter 4).

presented. I am particularly grateful to Bill deVries and Jim O'Shea for directing the seminar, and to the participants of the seminar (as well as to Bill and Jim) for offering much helpful feedback on this project in its earlier stages. Michael P. Wolf also read an early draft and offered helpful feedback and criticism. Finally, a referee for Analytic Philosophy provided a good deal of very helpful constructive criticism leading to the current version of the paper. Open access funding enabled and organized by Projekt DEAL. Open access funding provided by the Qatar National Library.

ORCID

Jeremy Randel Koons  <https://orcid.org/0000-0002-5492-348X>

REFERENCES

- Berger, P. L. & Luckmann, T. (1966). *The social construction of reality: A treatise in the sociology of knowledge*. Penguin Books.
- Bird, A. (2010). Social knowing: The social sense of 'scientific knowledge. In J. Hawthorne & J. Turner (Eds.), *Philosophical perspectives, Vol. 24: Epistemology* (pp. 23–56). : Ridgeview Publishing Company.
- Brandom, R. B. (1994). *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard University Press.
- Brandom, R. B. (2015). *From empiricism to expressivism: Brandom reads Sellars*. Harvard University Press.
- Bratman, M. E. (1999). *Intentions, plans, and practical reason*. CSLI Publications.
- Clifford, W. K. (1876–1877). The ethics of belief. *Contemporary Review*, 29 (December 1876–May 1877), 289–309.
- Dang, H. (2019). Do collaborators in science need to agree? *Philosophy of Science*, 86(5), 1029–1040.
- Elliott, T. (1994). Making strange what had appeared familiar. *The Monist*, 77(4), Feminist Epistemology (October), 424–433.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Fuller, S. (2002). *Social epistemology*. Indiana University Press.
- Gilbert, M. (2014). *Joint commitment: How we make the social world*. Oxford University Press.
- Goldberg, S. C. (2015). *Assertion: On the philosophical significance of assertoric speech*. Oxford University Press.
- Goldberg, S. C. (2018). *To the best of our knowledge: Social expectations and epistemic normativity*. Oxford University Press.
- Green, A. (2016). *The social contexts of intellectual virtue: Knowledge as a team achievement*. Routledge.
- Hardwig, J. (1985). Epistemic dependence. *The Journal of Philosophy*, 82(7) (July), 335–349.
- Heal, J. (1978). Common knowledge. *The Philosophical Quarterly*, 28(111) (April), 116–131.
- Koons, J. R. (2019). *The ethics of Wilfrid Sellars*. Routledge.
- Koons, J. R. (2021). Sellars on rational agency as presupposing collective Attitudes. In L. Koreň, H. B. Schmid, P. Stovall, & L. Townsend (Eds.), *Groups, norms and practices: Essays on inferentialism and collective intentionality* (pp. 189–213). Springer.
- Koons, J. R. & Wolf, M. P. (2018). The ordinary language case for contextualism and the relevance of radical doubt. *Contemporary Pragmatism*, 15(1), 66–94.
- Kuhn, T. S. (1977). The essential tension. In T. S. Kuhn (Ed.), *The essential tension: selected studies in scientific tradition and change* (pp. 225–239). : The University of Chicago Press.
- Kukla, R. & Lance, M. (2009). 'Yo!' and 'Lo!': *The pragmatic topography of the space of reasons*. Harvard University Press.
- Kusch, M. (2002). *Knowledge by agreement*. Oxford University Press.
- Laland, K. N. (2017). *Darwin's unfinished symphony: How culture made the human mind*. Princeton University Press.
- Longino, H. (2002). *The fate of knowledge*. Oxford University Press.
- Miller, B. (2013). When is consensus knowledge-based? Distinguishing shared knowledge from mere agreement. *Synthese*, 190(7), 1293–1316.
- Miller, B. (2015). Why (Some) knowledge is the property of a community and possibly none of its members. *The Philosophical Quarterly*, 260(65) (July), 417–441.
- Nelson, L. H. (1993). Epistemological communities. In L. Alcoff & E. Potter (Eds.), *Feminist epistemologies* (pp. 121–159). Routledge.

- Schutz, A. (1962). In M. Natanson (Ed.), *Collected papers, Volume I: The problem of social reality*. Martinus Nijhoff.
- Searle, J. (2009). *Making the social world: The structure of human civilization*. Oxford University Press.
- Sellars, W. (1963). Empiricism and the philosophy of mind. *Wilfrid sellars, science, perception and reality*. (pp. 127–196). Ridgeview Publishing Company.
- Tollefsen, D. P. (2015). *Groups as agents*. Polity Press.
- Tuomela, R. (2002). *The philosophy of social practices: A Collective acceptance view*. Cambridge University Press.
- Tuomela, R. (2004). Group knowledge analyzed. *Episteme*, 1(2) (October), 109–127.
- Tuomela, R. (2007). *The philosophy of sociality*. Oxford University Press.
- Tuomela, R. (2011). An account of group knowledge. In H. B. Schmid, D. Sirtes & M. Weber (Eds.), *Collective epistemology*. (pp. 75–117). : Ontos Verlag.
- Van Helden, A.(2010). Galileo and the telescope. In A. Van Helden, S. Dupré, R. van Gent, & H. Zuidervaart (Eds.), *The origins of the telescope* (pp. 183–201). KNAW Press.
- Wolf, M. P. (2012). Boundaries, reasons, and relativism. *Journal of Philosophical Research*, 37, 205–220.

How to cite this article: Koons JR. Knowledge as a collective status. *Analytic Philos.* 2021;00:1–28. <https://doi.org/10.1111/phib.12224>