UNIFYING GROUP RATIONALITY

MATTHEW KOPEC

Australian National University

ORCID: 0000-0002-6475-4368

Various social epistemologists employ what seem to be rather distinct notions of group rationality. In this essay, I offer an account of group rationality that is able to unify the dominant notions present in the literature under a single framework. I argue that if we employ a teleological account of epistemic rationality, and then allow that there are many different epistemic goals that are worth pursuing for various groups and individuals, we can then see how those seemingly divergent understandings of group rationality are all intimately related. I close by showing how the view has the additional benefit of allowing us to generate practical, normative suggestions for groups in the real world.

Keywords: Social Epistemology; Group Rationality; Epistemic Teleology; Pluralism; Epistemic Normativity.

1. Introduction

Undoubtedly, social epistemology has quickly become one of the more vibrant subfields in philosophy, and its relatively recent emergence offers a glance into some exciting new directions for our discipline. Social epistemologists have been examining a wide array of topics, from the more traditional, like testimony, expertise, and rational deliberation, to various newcomers such as epistemic injustice (e.g., Burroughs & Tollefsen 2016; Congdon 2015; Fricker 2007; McKinnon 2016), the value (or disvalue) of anonymity (e.g., Frost-Arnold 2014; Goldberg 2013), the epistemic consequences of implicit bias (e.g., Gendler 2011; Jönsson & Sjö Dahl 2017; Lassiter & Ballantyne 2017; Toribio in press), and conspiracy theories (e.g., Clarke 2007; Coady 2006; Dentith 2014; in press; Levy 2007; Pigden 2007), just to name a few. These inquiries employ an equally wide array of methods. Some social epistemologists prefer to present fictional scenarios to elicit our

---

Contact: Matthew Kopec <matthewckopec@gmail.com>

intuitions, while others prefer building computer simulations to examine how idealized agents would interact. Some prefer formal proofs, while others prefer to borrow empirical data from the behavioral, social, and organizational sciences. And, very often, these examinations seem to be driven by the desire to offer practically employable, normative guidance to real-world individuals, groups, and institutions.

But this kind of topical diversity and methodological pluralism can also come at some cost. In particular, it can be hard to tell how the various normative suggestions coming out of social epistemology relate to one another. It can also be hard to tell whether such suggestions are genuinely *epistemic*, or even whether they are genuinely *normative*. Take Goldman's 'systems-oriented' social epistemology as an example (Goldman 2010). For Goldman, an 'epistemic system' is "a social system that houses social practices, procedures, institutions, and/or patterns of interpersonal influence that affect the epistemic outcomes of its members" (2010: 197). On this definition, social epistemology includes the study of a very diverse class of entities, including such things as Twitter journalism, the education system, Wikipedia, and various funding agencies. While I agree that these examinations belong within social epistemology, I also feel there are important open questions whenever these examinations yield (supposedly) normative conclusions. For example, if an author were to argue that it would be irrational to have local education boards composed of elected officials, is this claim really expressing an epistemic judgment? And, if so, how does this epistemic judgment relate to the other epistemic judgments we make concerning various other individuals and groups? It won't do to simply insist that these assessments have something to do with truth.

In this article, I take up the challenge of explaining how the various epistemic requirements discussed in social epistemology could hang together, focusing here on the various notions of 'group rationality' that regularly appear in the literature. I will show that all of the various pieces fall into place when we think of epistemic rationality as an inherently teleological notion and maintain a kind of pluralism about what kinds of characteristics can be epistemically valuable to groups and individuals.¹ In addition to showing how the various notions of group rationality could all be intimately related, the account carries with it some other very desirable features. For example, it explains how the kind of normativity that binds the epistemic behavior of individuals could be the same kind of thing as both the normativity that binds the epistemic behavior of groups and that which ought to guide us in structuring our epistemic institutions. In short, it

---

1. Those familiar with Foley (1993) and Alston (2005) will see various parallels between my account and each of theirs. In the interest of space, I won't sketch the contours of how I disagree with them.

explains why we should see social epistemology as a genuinely normative subfield continuous with individual epistemology. The account also makes it clear that social epistemology is well situated to play the role of an *applied* normative subfield, akin to the role that applied ethics plays within value theory. Although my principal focus here will be on the unification task, these other benefits will emerge along the way.

The plan for the essay is as follows: In Section 2, I will present a number of notions of group rationality that appear in the social epistemology literature. In Section 3, I will give a quick sketch of my preferred account of epistemic rationality and explain how this account can unify the aforementioned notions of group rationality. In Section 4, I sketch some norms that follow from such a view for actual groups of inquirers. I conclude in Section 5.

2. Seemingly Disparate Notions of Group Rationality

In this section, I will discuss four main areas of inquiry where the notion of group rationality has either been, or could rightly be, employed in the social epistemology literature, namely, the areas concerning epistemic institutions, group agents, joint epistemic actions, and individuals with group directed epistemic goals.

Let us start with epistemic institutions. There are a number of social institutions whose central characteristics are epistemic in nature. Take the education system for example. Our society has a strong interest in properly educating its younger members so that they end up as rational as possible when they become adults. It's safe to assume that any society is more likely to flourish if its members are responsive to whatever evidence they have, seek evidence out when making important judgments, have a reasonable level of confidence in their beliefs, strive to correct inconsistencies in their belief systems, and so on. Therefore, we all have an interest in ensuring that our educational system is capable of increasing these characteristics in students over the duration of their education. If the majority of students make it through secondary education while remaining unresponsive to their evidence, overconfident in their beliefs, and generally not caring all that much about whether their beliefs are true (or even consistent), then that educational system is surely flawed and ought to be changed. There is a sense in which a failing epistemic institution like this can be thought to be irrationally organized. As another example, if the ways in which the European Research Council funds projects has the effect of actually inhibiting the growth of knowledge—for example by wasting funds on obviously dead-end projects for politically motivated or overtly corrupt reasons—this would also be irrational. I would argue that any institution that has an epistemic charter could be criticized

in a similar way when it fails to live up to that charter. If that's correct, then there are normative requirements for any institution covered under Goldman's (2010) aforementioned systems-oriented social epistemology, such as legal systems, journalism, Wikipedia, and much more besides.²

A second area of inquiry where we can rightly judge an entity's level of epistemic rationality involves group agents.³ There are certain groups, like a board of trustees, a small governmental agency, or the panel of judges that composes an appeals court, which can have the kind of unity necessary to judge and act much like an everyday singular agent (e.g., you or me). Rationality considerations enter the picture in a couple of places. First, List and Pettit (2011) argue that for a group of individuals to rightly count as an agent in its own right, the group needs to exhibit at least a minimal level of rationality in its intentional attitudes. Take a university board of trustees that lacks consistency in its preferences and judgments. For example, the spokesperson for the board regularly reports divergent assessments to the Chancellor (or President) about what the trustees would like the Chancellor to do, and the board initially rejects various proposals only to later change its mind, but without any corresponding changes in the relevant facts. This board of trustees doesn't meet the minimal rationality condition for agency. If an individual human were to behave like this, we would similarly doubt her claim to agency.

But there is another way in which rationality assessments can play a legitimate role in our theorizing about group agents. For example, we can also assess group agents according to whether they are responding well to the evidence held by the group (Hedden in press) or to whether the group agent is making judgments in the most reliable ways available (Goldman 2014; List 2005), just

---

2. It has been suggested to me by an anonymous referee (from an ultimately unsuccessful submission to *Episteme*) that some readers, especially those coming from the subfield of social ontology, might find my suggestion that the form of rationality at issue in epistemic institutions is a variety of 'group' rationality to be a bit jarring. After all, there is a vast literature on what it takes for a set of agents to count as a bona fide group, and the dominant accounts in that literature will rule out many (perhaps most) institutions. Additionally, those dominant accounts tend to treat matters of context or material culture as something outside the group (as a bona fide group), even though context and material culture will be crucial to our understanding of an epistemic institution. But I think this is largely a verbal dispute. To be clear, in this article I intend to be using the term 'group' broadly, in such a way that it could refer to mere collectives (or 'mobs') as well as highly structured groups, i.e., the kind of things that some might think have a rather special ontological status. I hope it is clear to such readers, from what I say below, that the latter kinds of special groups can generate their own modes of rational assessment, and that these modes of assessment cannot subsume institutional rationality in all cases.

3. There are a number of ways that authors have referred to group agents in the literature, including 'collective agents' (Searle 1990), 'corporate agents' (List & Pettit 2011; Pettit 2007), 'group agents' (Tollefsen 2015), and 'plural agents' (Helm 2008). In the interest of space, I will assume the differences between specific accounts attached to these labels don't much matter for our current purposes.

like we do with everyday individual agents.⁴ Simply forming and maintaining consistent collective intentional attitudes isn't enough to make a group agent fully rational. Rather, the group would need to meet further conditions that parallel those meant to ensure effective reasoning in the individual case.

It's worth taking a moment to explain why the notion of rationality at work in the case of epistemic institutions is not the very same notion of rationality at work in the case of group agents. The primary reason that these two notions must be treated as distinct is that many of the epistemic institutions of interest simply won't meet the minimal criteria to count as group agents. To be sure, some cases might come close. For example, the internet encyclopedia Wikipedia is often presented as a paradigm case of an epistemic system in Goldman's sense, since it is a kind of institutional framework that collects and presents information from a wide range of individuals in order to advance our knowledge of the world around us (Fallis 2008). And there is a kind of policing mechanism at work in the structure of this epistemic institution that is meant to ensure that its entries don't contradict other entries,⁵ and this would bring it closer to meeting the minimal consistency requirements for group agency (cf. Tollefsen 2009). But this kind of policing is certainly not present in large and amorphous institutions like the "educational system." For example, the way in which the Atlantic slave trade is discussed in history textbooks used in the conservative bastion of Texas is rather different from how it is discussed in the textbooks used in the progressive bastion of Vermont.⁶ Similarly, the policy decisions of school boards will at times differ drastically from place to place, and year to year, due to an ever changing political landscape. This all suggests that something like an educational system simply won't meet the minimal consistency criteria for counting as a group agent. And if such institutions don't count as agents, there is simply nothing to count as *its* judgments, or *its* decisions, or *its* evidence, and so there is really no way to make sense of how well *it* is making judgments and decisions based on the evidence. So, these notions must be distinct in at least some cases.

But there is a slightly more nuanced reason why the two notions cannot be coextensive, namely, that the effects we look for in a rational institution are very different from the effects we look for in a rational group agent. When we are

---

4. List (2005) considers such questions part of what he calls the 'knowledge challenge', since he understands the 'rationality challenge' to be solely about consistency. Lackey (2016) offers a challenging critique of the group agent understanding of justified group belief suggested by List and explicitly defended by Goldman (2014). Lackey favors a view with a more evidentialist flavor, somewhat akin to Hedden (in press), but the former strips away the requirement that the group at issue count as an agent in its own right, while the latter retains it.

5. Although, see Sanger (2009) who sketches how these mechanisms might regularly break down when stubborn (and anonymous) editors scare more qualified experts away from contributing.

6. See Strauss (2014) for some particularly egregious examples from Texan history textbooks.

examining the rationality of an institution, we are interested in whether the institution operates in a way that has beneficial effects on the individuals affected by the institution. But when we are examining the rationality of a group agent we are, instead, interested in matters like whether the group agent is forming the best judgments it can based upon the evidence it has access to within the group. And these two kinds of examinations can point in opposite directions when we have a group agent that also counts as an epistemic system. Let's return to Wikipedia. It's likely that Wikipedia editors are, themselves, made epistemically worse off by spending their time fighting over minutia on the comments forums or by tirelessly searching for inconsistencies amongst the various pages. Nonetheless, these efforts greatly improve Wikipedia's outputs on the whole. So, in an imaginary world where the only Wikipedia consumers were also the Wikipedia editors, we'd have a case of a group agent that's exceptionally rational overlapping with an epistemic system that's failing its community (essentially a failure of effective division of cognitive labor). Failures in the other direction are also common. For example, it's now well known that unstructured discussion is an awful way for a group to make judgments and decisions (see, e.g., Sunstein & Hastie 2015). But during such discussions, the members of the group surely learn all kinds of useful things, like various facts they simply hadn't run across before and new insights into the background values and beliefs of fellow group members. So, an unstructured group discussion could be a case of a rationally functioning epistemic institution that makes the group agent who uses it much less rational.

A third area of inquiry where group rationality should play a role in our theorizing involves groups who perform joint epistemic actions. Here I will follow the analysis of joint epistemic action proposed by Miller (2018). On Miller's account, groups of agents can be said to act jointly in cases in which each agent acts in order to achieve a common goal, where this collective end is a matter of mutual true belief (Miller 2001).⁷ These conditions can be met both for more typical joint actions, such as when a team of sailors tacks their yacht around a mark during a race, as well as for epistemic actions, such as when a team of investigators works to determine whether Smith is the murderer. And while group rationality isn't the focus of Miller's examinations of joint epistemic action, this is surely a place where we can make epistemically normative assessments of the

---

7. Those familiar with the literature on joint action should see some obvious parallels between Miller's account and Bratman's (1992; 1993; 2014). The main difference between their accounts that is relevant to this particular discussion is that the mutual true belief condition in Miller's account is much weaker than Bratman's common knowledge requirement (which is very stringent); so it will be easier for a group to hold joint epistemic ends on the account I prefer. Both accounts drastically differ from Gilbert's 'joint commitment' account (1989; 2013), which, admittedly, is not as amenable to the account of epistemic norms offered here. For some independent reasons to prefer an account like Bratman's or Miller's over Gilbert's, see Kopec and Miller (2018).

epistemic actions involved. For example, if the members on the team of detectives mentioned above refuse to share their individual information about the case with the other detectives, and this causes the team to fail to uncover that Smith is indeed the murderer, then we might assess the group as being irrational from a genuinely epistemic perspective. Importantly, the way ‘rationality’ is being used here is also distinct from how it is used in the case of institutions and group agents. For example, our group of detectives might be much too transient to count as a group agent—their judgment about whether Smith is guilty may well be the only determination they ever make together. Recall that for List and Pettit, there must be a minimal level of rational unity over time for an entity to count as an agent, and such one-off groups won’t meet that condition. And while there are surely institutions involved in criminal investigations, the failure of rationality in this group isn’t necessarily the failing of an institution—it’s much more likely a failing on the part of the individuals involved. Similarly, it’s possible for epistemic institutions to form and function without the efforts of any joint epistemic actors, which will happen in cases where the individuals working within the institution never have the level of contact necessary to form joint epistemic goals.⁸

I would add a fourth area of inquiry where group rationality considerations regularly arise but which has received much less discussion in the literature. Often individuals within a group have group-directed epistemic goals, and yet the epistemic actions in question don’t fit the criteria for them to count as joint actions. For example, there can be cases where an individual member of a group has, as one of her goals, that her group accomplishes something she finds epistemically valuable, even though all of the other members of the group lack this goal. Such cases will fall short of joint actions on Miller’s account, since the epistemic end in question isn’t held jointly amongst the members. Take, for example, the case of a team of scientists, which Miller presents as a paradigm case of joint epistemic action (Miller 2018). I would argue that most real-world teams of scientists actually lack mutual true belief concerning some common epistemic end. In reality, there will be members of the team that are only participating for the salary, grant money, prestige, or publications.⁹ And given that scientists know such folks are ubiquitous, even the more epistemically virtuous

---

8. Anyone familiar with Miller’s accounts of group agency and social institutions (2001; 2010) will note my disagreement with those accounts. Much of Miller’s work aims to analyze notions like group agency or social institutions in terms of joint action. I, on the other hand, believe that these three notions are distinct, and facts about entities of each sort can be the source of their own form of normativity.

9. Kitcher (1990) has argued that these kinds of impure motivations might actually be *beneficial* to the progress of science overall.

among them likely won't believe themselves to be surrounded entirely by other epistemically virtuous folk. I would argue that group-level normative requirements also apply to these individuals even though the joint epistemic action conditions aren't met. For example, if a virtuous lead investigator knows that one of her graduate assistants lacks the group-directed epistemic goals relevant to the project, there are still rational ways for her to educate him, delegate duties to him, share data with him, etc. That is, there are better and worse ways to pursue her own epistemic goals for the group, even if she's the only member of the group who actually desires to figure out the truth about her object of inquiry.¹⁰

So there are at least four rather distinct notions of epistemic rationality that can play a role in our assessments of the inner workings of social groups. There seem to be epistemically normative requirements for structuring institutions that can't be analyzed simply in terms of norms applying to group agents or those applying to individuals with group directed goals. There seem to be normative requirements for group agents that don't count as epistemic institutions and that are composed of individuals who lack the relevant joint epistemic ends. Normative requirements seem to apply to joint epistemic actors in a group that doesn't count as a group agent or an epistemic institution. And normative requirements also seem to apply to individuals with group-directed epistemic goals but who happen to be members of groups that lack all three of the special characteristics above. How are we to reconcile all of these seemingly distinct forms of group epistemic normativity?

3. How to Unify Group Rationality

I would argue that all of the abovementioned uses of group rationality aren't really as distinct as they seem. On a certain account of epistemic rationality, all of these notions of group rationality are intimately related to each other, and, additionally, they all stem from the same kind of normativity as the kind involved in individual epistemic rationality. In this section I will quickly sketch this ac-

---

10. The existence of norms that apply to individuals with group directed epistemic goals will, in turn, explain why much of the recent work in so-called 'network epistemology' (e.g., Mayo-Wilson, Zollman, & Danks 2011; Zollman 2013) should be seen as examining genuinely epistemic norms, as opposed to merely instrumental norms. If some agents in these networks care about their groups converging upon the theory with the highest expected epistemic utility, then they ought to do what they can to form networks in ways that allow such convergence. In a similar vein, genuinely epistemic norms will automatically apply to groups that contain members who desire for their group to hold justified group beliefs in the sense discussed by Lackey (2016). Whether such groups will also generate epistemic norms of the other kinds discussed above would depend on the further details of the particular group.

count of epistemic rationality, and then point out how this unifies the seemingly distinct notions of group rationality into a coherent whole.¹¹

On my account, rationality is an essentially goal-oriented, that is, teleological, notion, where the rational attitudes to hold are those that are conducive to attaining the relevant goals.¹² I believe that this is just as true in the case of practical rationality, where teleological accounts are rather popular, as it is in the case of epistemic rationality, where such accounts are not quite as common. In fact, I argue that epistemic rationality is simply a special form of practical rationality, where the latter is understood in the teleological sense. The main difference is that epistemic rationality assessments examine how well an agent is doing at pursuing certain cognitive or epistemic goals, whereas practical rationality assessments look at how well an agent is doing at pursuing various goals more generally (which can include both epistemic and non-epistemic goals). Importantly, an account of this sort need not collapse into a form of pragmatism (like, e.g., James 1907/1995; Rinard 2017; Stich 1990). For example, we can assess an agent as being practically rational for holding the belief that she'll recover from pancreatic cancer, while at the same time being epistemically irrational for doing so. Her belief might best serve her non-epistemic goals, since her wishful thinking might boost her chances of survival, even though such a belief might not serve her epistemic goals of believing truthfully or believing in accord with the evidence. (That is, if she's aware of the very low survival rates.)¹³

To make things a bit more precise, we could state the view as follows:

---

11. For a more in depth look at the kind of pluralistic view I spell out here, but solely at the individual level, see Kopec (2018).

12. An anonymous referee has rightly encouraged me to flag the fact that the term 'teleological' is sometimes used to refer to a broader range of normative accounts that have goal-oriented aspects, including some views that don't define normativity so as to sanction the attitudes or acts that are conducive to attaining the relevant goal. For example, in epistemology, Littlejohn (2018) and Mehta (2016) both propose knowledge-oriented teleological accounts that are not conducivist in the way that my view is. (Some, like Littlejohn himself, refer to these views as non-consequentialist teleological accounts.) Also, one could see fitting-attitude accounts in value theory, like the one developed specifically in epistemology by Sylvan (2012; 2018; in press), as teleological in a similar way, since they hold that agents ought to have the attitudes that are the fitting ones to have given some specified end that is of final value (such as having a respect for the truth, as in Sylvan's account). Sometimes, the kind of teleological views I have in mind are referred to as varieties of 'instrumentalism', i.e., those with a commitment to the instrumentality of derivative value (see, e.g., Sylvan 2018). I shy away from using the qualifier 'instrumental' here because there is also a substantial literature on the topic of 'epistemic instrumentalism', spawned largely by an attack on the view by Kelly (2003). Although that topic is related to this discussion, as discussed in Kopec (2018), it is ultimately about a distinct set of views. Unfortunately there seems to be no completely settled and unambiguous terminology to use in this context.

13. Admittedly, the notion of having practical reasons to hold beliefs is controversial. For example, Kelly (2002), Shah (2006), and Parfit (2011) dispute the possibility of practical reasons for belief. See Leary (2017) and Rinard (in press) for some recent defenses of their possibility.

A particular attitude (behavior, social structure, etc.) is maximally epistemically rational for an agent (or entity) if and only if that attitude (etc.) is amongst the most effective means toward pursuing some specified set of epistemic goals.

So, in the simple individual case, an agent's belief in some proposition will only be maximally epistemically rational if forming that belief is the most effective means for her to pursue some specified epistemic ends. (I'll discuss the more complicated group cases in a moment.) But notice that even this more precise formulation contains a number of aspects still in need of disambiguation. For one, we would need to specify which epistemic goals are the relevant ones for such assessments. That is, do we assess the agent according to *her own* epistemic goals or, instead, according to some goal or goals we believe to be of independent epistemic value (regardless of whether she herself values it or them)? Additionally, we would need to specify whose perspective is the relevant one for the assessment. That is, do we assess the agent according to how the world looks from her perspective, or instead from the way the world actually is?

I'll start with the first question. In the case of practical rationality, authors have been split on how to answer this question (see, e.g., Kolodny & Brunero 2018). So-called 'desire-based' theorists believe that the agent (practically speaking) ought to do what best promotes the satisfaction of her own desires, whatever those might be. So-called 'value-based' theorists believe that the agent's own desires are largely irrelevant, since what she (practically speaking) ought to do is whatever best promotes the attainment of things of independent value, even if she lacks any relevant desires to attain them. The latter views are then further split into monists, who believe there is only one kind of state that's of independent value, and pluralists, who believe there are a number of them (Heathwood 2015). For example, take an agent who most desires to smoke cigarettes and engage in generally risky behavior but who genuinely lacks any desire for long-term wellbeing. A desire-based view would have to admit that whichever course of action best secures this individual the capacity to smoke and engage in risky behavior is the course she ought to take. Most value-based views will disagree with the desire-based verdict here, since that kind of behavior doesn't best promote the kinds of things that are really of value, be that pleasure, sustained friendship, overall objective wellbeing, etc.

Although most authors working on practical rationality have seen the desire-based/value-based distinction as a dilemma necessitating a theoretical choice, I actually believe that both views are partly correct, since they both tap into a legitimate source of practical normativity. It's true that the agent ought to do whatever is best in order to pursue her own goals, and it's also true that she ought to do whatever's best to pursue the things of independent value, even if

these two oughts pull her in different directions. We can think of it like this: the ‘ought’ involved in each claim is indexed to a distinct realm of practical normativity.¹⁴ In other words, the question “What ought she (practically speaking) to do?” is actually ambiguous. It doesn’t always have an unambiguous answer until we specify which mode of assessment is under examination. And since I hold that epistemic rationality is a subset of practical rationality, I believe this kind of pluralism similarly applies in the epistemic case. We can examine whether an agent’s epistemic behavior best promotes her own epistemic goals, which I call a ‘liberalistic’ assessment. Or we could instead examine whether it best promotes something of independent epistemic value, which I call an ‘idealistic’ assessment. And both of these examinations can yield a legitimate assessment of epistemic rationality. In short, questions like “What ought she (epistemically speaking) believe?” are similarly ambiguous. In many cases we would need to specify whether we are making a liberalistic or idealistic assessment before we could answer this question.

I must address one other question before moving on to the matter of perspectives. There is a vast literature addressing the question of what sorts of states or qualities have independent (or final) epistemic value. Some have suggested that the coherence of one’s beliefs or credences is what’s valuable (e.g., Quine & Ullian 1970). Others think what’s valuable is having beliefs that accord with the evidence (e.g., Conee & Feldman 2004). Still others think it’s holding beliefs that are true (e.g., Goldman 1986) or credences that are accurate (e.g., Pettigrew 2016). And finally, some think only knowledge, in the now standard sense of justified true belief plus some anti-Gettier condition, is the only thing of independent epistemic value (e.g., Williamson 2002). While I don’t have space to fully defend my view here, I believe that we can legitimately take all of these, and various others besides (e.g., wisdom, understanding, etc.), to be of independent epistemic value. In short, there are various contexts in which we rightly care about these different qualities of beliefs (behaviors, systems, etc.), and in varying degrees, in which case these values must be balanced against each other. Thus my view, on the idealistic side of the divide, is somewhat parallel to so-called ‘objective list’ theories in the practical rationality literature (cf. Nussbaum & Sen 1993).

Concerning the second question, about which perspective is relevant for the assessment of epistemic rationality, I take a similarly pluralistic line. This question parallels the classic debate between the internalists and externalists in traditional epistemology. I believe that we can legitimately assess an agent’s epistemic behavior according to whether it best promotes the relevant epistemic

---

14. This kind of view, as it relates to practical rationality, has close relatives in those who hold ‘dualist’ views concerning practical and moral rationality (e.g., Copp 1997), and those who deny that there are ‘all-things-considered’ practical oughts, or an ‘ought simpliciter’ (e.g., Tiffany 2007 and Baker 2018).

goals given how the world seems to her, whether she is a brain in a vat or a more typical individual perceiving the real world along with the rest of us. I also believe we can legitimately assess an agent's epistemic behavior according to whether it best promotes those goals according to the way the world actually is. There are various contexts in which these different assessments make perfect sense, and I don't feel there is any need to preemptively choose between them. We simply need to be clear about what kind of assessment we are making. This kind of pluralism about internal and external reasons is now a well-established position in the practical rationality literature (cf. Schroeder 2007), and I think it's natural to extend the position into the epistemic realm.

To sum up my account, we have four different ways of assessing an attitude (behavior, social structure, etc.) for epistemic rationality:

- (1) A **liberalistic internalist** assessment judges whether the attitude (etc.) is among the best means for the entity to promote its own epistemic goals from its own perspective on the way the world is.
- (2) A **liberalistic externalist** assessment judges whether the attitude (etc.) is among the best means for the entity to promote its own epistemic goals according to the way the world actually is.
- (3) An **idealistic internalist** assessment judges whether the attitude (etc.) is among the best means for the entity to promote some set of laudable epistemic goals from its own perspective on the way the world is.
- (4) An **idealistic externalist** assessment judges whether the attitude (etc.) is among the best means for the entity to promote some set of laudable epistemic goals according to the way the world actually is.

There will be many cases where all four assessments will yield identical verdicts. For example, the verdicts would be guaranteed to agree in any case where an agent desires to attain the laudable goals at issue and has an accurate perspective of the way the world is. And they will agree in many other cases outside this limited class. But, importantly, they will not always agree. Therefore, whenever asking questions about epistemic rationality, it is important to be clear about which mode of assessment we intend to discuss.

Once we move to this kind of pluralistic, epistemic-goal-oriented account of epistemic rationality, we are able to explain the close relationships between the seemingly distinct forms of group rationality discussed in the previous section. Let's start with epistemic institutions. Epistemic institutions are endowed with epistemic goals by the very role that they play in our social system. This is not to say that an epistemic institution is some kind of agent with intentional attitudes, and among these attitudes are certain epistemic goals. Rather, the in-

stitutions have the goals they do because we construct them to fulfil a certain social epistemic purpose. Therefore, it is fully legitimate to claim that an institution is operating irrationally when it is doing worse than it could at promoting the relevant goals given the resources available. I would argue that the answer to the question of which goals are the most relevant to our assessment will ultimately depend upon which institution we're examining. For example, we might assess the National Science Foundation (NSF) according to how well its funding procedures have advanced our understanding of nature. And we might assess our educational system according to how critically minded and responsive to evidence our children become on average. If these institutions are operating rationally, they'll be promoting these values. And if not, we have epistemic reasons to make changes.¹⁵

My view offers a similarly satisfying account of how epistemic rationality operates with group agents. Some groups that count as agents in their own right, in the sense sketched by List and Pettit (2011), have collective epistemic goals. Likely examples are juries, governmental intelligence teams, and closely-knit groups of scientists (i.e., some 'labs'). Since these group agents have their own epistemic goals, we can start to assess them in the liberalistic mode, that is, according to how well they are pursuing those goals, whatever those goals happen to be. For example, I would argue that the goal of a jury, if they have a proper understanding of their job, is to come to the correct judgment concerning the weight of (what is typically) a proper subset of the evidence, since they are supposed to set many of their prior beliefs, and often times various things said in court, to the side when making their final determinations. So, if a jury is operating in such a way that makes it unlikely to judge the proper weight of the admissible evidence, for example by discussing facts they privately know but that were not presented in court, then this jury isn't acting rationally. But this isn't the only mode of assessment available to us; we might also assess this jury according to whether it is promoting some other goal it doesn't necessarily hold, like whether its actions tend to track the truth. And to the extent that we can make sense of a group agent's perspective on the world—that is, to the extent that it has the relevant beliefs—we can also assess the group agent either according to its own perspective or, instead, according to the way the world actually is. Importantly, all of these assessments can diverge in various ways.

---

15. Note that the fact that these are *epistemic* reasons, as opposed to merely practical reasons, is important, since there might be very strong practical reasons for an institution like the NSF to operate in an epistemically sub-optimal way. Understanding the aspects of biology that lead to advances in biological weapons or treating disease might not add as much to our overall understanding of the biological world as understanding other aspects—such as more foundational issues in evolutionary theory—but the former are probably, practically-speaking, much more important for the NSF, given the non-epistemic goals the US government has endowed upon it.

In the case of groups whose members perform joint epistemic actions but fall short of full group agency, we can also make perfect sense of epistemic rationality assessments within an account like the one I propose. It follows from the very definition of ‘joint epistemic action’ that the individuals involved have certain joint epistemic goals, and we can use these in a liberalistic assessment that judges how well these individuals are promoting the various goals they hold in common. And since the individuals will have their own perspectives on how the world happens to be, we can either use these perspectives in internalist assessments of the agents, or we could instead judge their behavior externalistically by looking at what best promotes their goal given how the world actually is.

Finally, my account explains how group-level epistemic rationality can put normative pressure on a member of a group even when that group lacks collective agency, contains no members with joint epistemic goals, and doesn’t itself count as an epistemic institution. On my account, as soon as an individual member of a group develops a group-directed epistemic goal, this group-directed goal can be used in a liberalistic assessment of that agent’s attitudes or behavior. Take, for example, a group of federal investigators held together simply through their common (non-epistemic) purpose of breaking up a certain organized crime ring. Let’s assume they never take votes or hold collective discussions to settle upon collective judgments, they lack a spokesperson to make any group proclamations, and they never coordinate their actions in any meaningful way. Further assume that even though the members all have various epistemic goals that are somewhat related to their common purpose of breaking up the crime ring, these epistemic goals are completely disjoint, that is, no two members have exactly the same group-directed epistemic goals. For example, Detective Grey aims for someone in the group (whether that be herself or someone else) to discover which corrupt bank the crime ring has been using to launder its dirty money, while Detective Green aims for someone in the group to discover whether the crime ring is connected to a particular string of execution-style murders, and so on. Given this setup, this group would fail the conditions both of group agency and of joint epistemic action, and therefore previous accounts would be unable to diagnose any group-directed epistemic behaviors as irrational from anything more than the perspective of individual rationality.¹⁶ On my account, if one of these agents behaves in a way that is detrimental to the group attaining the epistemic goal that she personally has for the group, she can be assessed as behaving irrationally. For example, if Grey fails to share pertinent banking information

---

16. Admittedly, given that they possibly count as having a joint practical goal of breaking up the organized crime ring, we may still be able to say they act irrationally relative to that practical goal. But this, importantly, won’t be an epistemic failure. And, anyway, if this group lacks mutual true belief about their joint goal, it wouldn’t count as having a joint practical goal under Miller’s framework (or Bratman’s, for that matter).

with Green, and Green just happens to be adept at following the money trail, then Grey can be criticized as acting irrationally according to a liberalistic assessment.¹⁷ And, just like in the case of joint epistemic actions, we can assess her either from her own perspective, or according to how the world actually is. For example, we might be able to claim Grey was behaving irrationally even if she couldn't have known Green was the financial specialist in the group.

One crucial thing to notice is that, since the relevant social entities can overlap and their respective group directed goals can conflict, the various modes of assessment might come into conflict. For example, it's possible to have a group of individuals who all have group directed epistemic goals, while maintaining other goals that rise to the level of joint epistemic goals, and this group could have the kind of coherent unity in judgment and behavior to make it count as an agent in its own right, possibly with a completely different set of epistemic goals.¹⁸ Even though the ultimate source of epistemic normativity is the same kind of thing at each level of assessment, the assessments might push the agents in drastically different directions. For example, it might best promote the group agent's epistemic goal if the individuals shared information in one way, best promote the members' joint epistemic goals to share it in a different way, and best promote each individual's group directed epistemic goals to not share information at all. Much will depend on the details. Things will get even more complicated once we recall that the individuals in these groups will also have their own epistemic goals regarding their solitary epistemic lives, and they can be judged even with reference to various laudable epistemic goals that they don't necessarily hold for themselves. The picture of rationality that emerges is one where agents will often be pulled in different directions by the various sources of epistemic normativity.

While I admit that possible conflicts between these different modes of assessment might arise, I still maintain that the resulting way of understanding epistemic rationality serves to unify group rationality within a single normative realm. In particular, the ways we could adjudicate whether a certain attitude or behavior is irrational at the various levels are all similar in a very important way. We just need to be told what the relevant goals are, and what the relevant perspective is, and we are able to embark on a largely naturalistic examination of what behaviors best promote those goals when the world matches the relevant perspective. This allows us to use the resources of the social sciences, for example, social psychology, organizational behavior, rational choice and decision

---

17. The aforementioned work in network epistemology (see Footnote 10) will have fairly obvious normative applications in cases like these.

18. The possibility that individual and group goals can come apart like this is one of the motivations behind List and Pettit's (2011) account. Their often discussed discursive dilemma is one such case.

theory, and agent-based modelling, to determine what the agents in question ought to do, or how the institutions in question ought to be structured. In the next section, I'll walk through some examples of how the framework can be applied to real-world cases, so as to use the empirical sciences to generate conclusions concerning individuals, groups, and institutions.

But before moving on to that task, it is worth taking a moment to address some commonly posed questions about this kind of pluralistic framework. First, one might wonder whether one of the four modes of assessment should be seen as privileged over all others. For example, when we encounter an entity that is in the grips of a conflict between one mode of assessment and another, does one of those modes trump the others in order to yield a final answer as to how that entity ought to believe, act, or be structured? A view on which one mode of assessment trumps all others might seem especially promising within the context of social epistemology, since there is only one mode of assessment that is open to all four of the types of groups I've discussed, that is, the idealistic externalist assessment.¹⁹ So, does the idealistic externalist mode generate epistemic norms that are guaranteed to take precedence over any conflicting norms generated via other sources?

I believe that a pluralist like myself ought to deny that one of the modes of assessment takes precedence over the others, because any such privileging carries what I take to be serious costs.²⁰ For example, if we were to privilege idealistic assessments over the others, we would jeopardize our ability to offer motivating rational criticism to agents who happen to lack the laudable goals we're tempted to base our assessment on. Just as a toy example, if an agent couldn't care less about whether her beliefs are true, but cares a great deal about having a consistent belief system, then our pointing out that her beliefs aren't great from a veritistic perspective surely won't do much to motivate the agent to do any better. In short, if we allow liberalistic assessments, we always have a retort to the "who cares" kind of response to our criticism. The agent herself must care, given the goals that ground our liberalistic assessment. And if we were to ultimately privilege externalist assessments, we would in essence lose a mode of assessment that promises to guide the agent "from the inside," as it were. To borrow an example that Jackson and Smith (2016) use in a different context, a normative theory that tells you how to make your way through a labyrinth by instructing

---

19. Recall that some epistemic institutions won't count as group agents, in which case the institution is guaranteed not to have its own personal perspective on the world or its own personal epistemic goals from which to assess the institution. But we can always assess an entity, be it agential or otherwise, according to how well it's doing at pursuing our specified laudable epistemic goals given how the world really is.

20. Each of the losses discussed below is basically the inverse of one of the benefits discussed in more detail in Kopec (2018).

you to “find a path to the exit” is obviously unhelpful when you’re in the middle of the labyrinth and just can’t find your way out. It is better to have a source of authoritative norms that works with what the agent has at her disposal, and, in our current context, that’s a body of beliefs that simply doesn’t come with a roadmap to tell the true ones apart from those that are false. Lastly, we would need to accept that internalist assessments come with their own authority in order to capture the robust intuitions that internalist accounts in epistemology were designed to capture.

But I hope these points in favor of allowing that liberalistic and internalist assessments have their own independent authority won’t leave the reader with the impression that liberalistic or internalist assessments must trump the others. First of all, as mentioned earlier, some entities we want to assess won’t have the requisite goals or perspectives needed as inputs to liberalistic or internalist modes of assessment. And, even when they do have the requisite goals or perspectives, idealistic and externalist assessments play vital roles of their own. For example, we need to be able to assess agents when their epistemic goals are deficient, and one way of doing this is to note that the entity’s attitudes or behaviors would be irrational when assessed according to the laudable set of epistemic goals which that agent, in an idealistic sense, ought to have (whether she cares to or not). Second, we will need to be able to anchor our assessments into how the world really is whenever our aim is to revise our epistemic practices to deal with problematic cognitive or social biases. In short, empirical facts about how individual or group thought and inquiry could be improved have obvious normative relevance. Allowing that these facts can generate norms with their own epistemic authority, in turn, explains why we (epistemically speaking) ought to strive to educate groups and individuals on how to recraft their epistemic lives. And, lastly, externalist positions in epistemology have their own intuitive appeal, and I feel the most promising pluralistic position will also respect those robust intuitions.

This kind of pluralistic position does admittedly come with some costs. In particular, once we accept that each mode of assessment has its own independent authority, with no assessment able to trump the others, we can no longer comfortably assert that there is something an individual or group epistemically ought to do, *all-things-considered*. As alluded to earlier, there will certainly be cases where all of our assessments line up, in which case it becomes clear what the correct epistemic course must be. But in cases of conflict, it becomes much less clear how the normative push of each kind of assessment is to be weighed. Although I am admittedly less confident in these waters, I expect there won’t be any satisfactory process for weighing the normative push of the various assessments. In other words, we ought to prepare ourselves for the possibility that we’ll encounter incommensurable epistemic requirements, each pushing agents

and groups in different directions. Although this might seem troubling, this kind of possibility is starting to find proponents working on normativity more generally (see, e.g., Baker 2018). I, as a personal matter, have found this possibility less troubling as I have considered the nature of group-level normativity more generally. Groups of all sorts can overlap in myriad ways, and any agent caught in such an overlap will occasionally find herself pulled in different directions. She could be a member of one group with a certain joint epistemic aim, a member of a distinct agential group with a completely different epistemic aim, a member of a third, less-structured group, for which she holds group-level epistemic goals, and also a member of a steering committee with some power to change an epistemic institution. For example, she might be a scientist who is, with a close colleague, jointly trying to develop a certain cure for cancer; who is also a member of a well-structured lab studying the effects of a particular gene; who is also a member of a less-structured team studying a completely different gene; and who is also a member of the board of research directors for the National Institute of Health (NIH). I feel it's obvious that this scientist's membership in each of these groups will often push her in different directions—sometimes irresolvable ones. Given that I feel these kinds of conflicts are simply a deep and ubiquitous matter of our social epistemic lives, I see little reason to assume our individual epistemic lives must be so different. I know that even when I sit in solitary contemplation over the nature of my own rationality, I feel much more like this deeply conflicted scientist than I do some caricature of Descartes.

Although I think normative conflicts are simply a fact of our epistemic lives, it's worth pointing out that there will also be many cases of conflict where the context of assessment will rightly drive us to privilege certain modes of assessment over others. For example, say we are trying to devise a method for aggregating probability assessments within small groups so that we can take the final assessment and employ it in policy making. In a case like this, we don't so much care whether any particular small group finds our method to be justified, or whether they are motivated to use it. And we might not care whether they have any idea about how the method works. All we care about in a case like this is how to make the final assessment as accurate as possible, so that we can force the group to use the very best method. So even if internalist and liberalistic assessments were to point away from some method that proves most accurate, we're unlikely to care. But in other contexts, we may care a great deal about what epistemic goods a group aims to attain, and there are very many cases where the way the world looks to a group matters to our assessments. For example, if we are trying to determine the most efficient way to run a 'consensus conference' (see Solomon 2011), where the express goal is to settle on a unified group opinion to present to the public, then any idealistic assessment that suggests a course of action that's very unlikely to lead to a group consensus attitude won't

ultimately have any pull for us. This suggests that we will favor liberalistic assessments in some contexts. And we will also favor internalist assessments in some contexts. For example, say we are trying to determine whether a tobacco industry player, like the company Phillip Morris, should be held accountable for promoting smoking during the 1970s. In a case like this, what's relevant to us is whether it would have been rational for the relevant group of corporate actors to have believed that smoking caused cancer given the information they had access to at the time. We now know that smoking causes cancer, but the way the world happens to have turned out isn't what's most relevant in this context of assessment. We're interested in how the world looked *to them*. In short, the possibility of conflicts should seem less troubling once we notice that context will often determine which kind of assessment we should focus our attention on.²¹

4. Some Normative Applications

In their book *Wiser: Getting Beyond Groupthink to Make Groups Smarter* (2015), Sunstein and Hastie document a number of empirically confirmed pitfalls that afflict groups attempting to make judgments. In particular, open and unstructured group deliberation tends to lead to rather unreliable judgments under certain circumstances. For example, unstructured group deliberation tends to elicit only a subset of the total relevant information held by the group's members. Information that is held in common (i.e., that which multiple members would antecedently agree on) tends to be ruminated upon, while information held solely by a single member is often not even mentioned, and, if mentioned, such information is less likely to affect the overall decision or judgment of the group.²²

This general phenomenon has been reproduced many times in laboratory studies. To get a handle on the phenomenon, take the classic study by Stasser and Titus (1985). In this study, university students were split into four-person groups and tasked with deciding, as a group, which of three candidates was the most qualified to be student council president. The study's designers constructed profiles for each candidate that contained a total of sixteen pieces of information, some of which were positive, some neutral, and some negative. For example, a positive piece of information might be that the candidate has a very

---

21. I thank an anonymous referee for pressing me to clarify the cases in which the context of inquiry will cause us to privilege one kind of assessment over another.

22. There are certain types of problems where these phenomena won't tend to occur. In particular, they won't occur in so-called 'demonstrable' problems, where once the answer is mentioned it becomes immediately obvious to everyone in the group that that answer is correct. In what follows, I'm assuming that the kind of problem the group is working on isn't demonstrable. The majority of group problems probably fit within this category, since most predictive problems and most value-laden problems do.

high overall GPA, while a negative piece might be that the candidate holds a very unpopular stance regarding campus alcohol regulations. In all experimental conditions, the complete profiles were designed to favor Candidate A over Candidates B and C. The total profile for A held eight positive, four neutral, and four negative characteristics, while the total profiles for B and C contained four positive, eight neutral, and four negative characteristics.

The profiles were then selectively given to the group members depending upon which experimental condition that group was selected for. In one condition, all four group members were given the full profiles for each candidate. In another condition, A's eight positive characteristics, and B's four negative characteristics, were divided up among the group, so that, for example, only one member was told about A's high GPA and B's unpopular position regarding campus drinking. In both conditions, the group was told to discuss the relevant facts with the rest of the group and to collectively come to a decision about which candidate was most qualified for the position. The results were dramatic. Groups under the condition where every piece of information was given to each member of the group ended up deciding A was the best candidate 85 percent of the time after discussion. But groups under the condition where positive pieces of information about A and negative pieces about B were each selectively given to just one member of the group decided B was the best candidate 71 percent of the time after discussion. So, even though both kinds of groups were given the same information about the candidates in total, their discussions, and corresponding judgments, were obviously very different. Studies like these strongly suggest that unstructured discussions tend not to elicit, or to allow for properly weighing, all of the relevant information held within the group.

Sunstein and Hastie (2015) catalogue a few strategies that social psychologists have devised to fight these kinds of information loss tendencies. One strategy is to instill a kind of division of cognitive labor, where each agent is told that she is charged with representing the information she has been given during the group's discussion. Another is to ask each agent to reflect on what she knows that might be relevant to the discussion before any information is exchanged within the group. And, finally, group judgments made simply through anonymous voting can't exhibit this kind of tendency. (Although anonymous voting can carry its own problems.) In many cases, groups will make more effective judgments if they take some of these steps instead of holding an open discussion.

One mechanism thought to be at work in groups that exhibit the tendency mentioned above is an effect called an 'information cascade'. This effect occurs when an individual receives information about a certain state of affairs through the actions or reports of others in her social environment, and this social information causes her to place less weight on her own private informa-

tion when making a judgment. The resulting judgment correspondingly affects that agent's utterances or behaviors, which then add to the social information received by other agents yet to form judgments. And the cycle continues down the cascade. It is much easier for cascading groups to form a kind of false consensus, since the early information has a disproportionate effect upon the attitudes of the group.

One particularly striking example of an information cascade was elicited in a laboratory experiment conducted by Anderson and Holt (1997). In this experiment, groups of six subjects were presented with an urn that they were told either contained two dark marbles and one light marble, or vice versa. They were then instructed to draw from the urn, observe the marble privately, replace the marble in the urn before the next subject's draw, and then publicly report their best guess about which type of urn the group was drawing from. One trial outcome is represented below in Table 1. Notice that in this example, respondent #3 reports a guess contrary to her private evidence, and this starts off a cascade of these contrary responses within this group. In fact, the members of the group unanimously guess the urn to be the one with two light and one dark marble, even though the majority of the group drew dark marbles.

Table 1. Observed Reverse Cascade (Anderson & Holt 1997)

	#1	#2	#3	#4	#5	#6	URN USED
Observation	Light	Light	Dark	Dark	Dark	Dark	Dark
Report	Light	Light	Light	Light	Light	Light	

On a view like mine, the empirical studies on these effects, as well as those that catalogue the ways of avoiding these effects, can be used to infer various epistemically normative constraints for groups. I'll start by examining what normative requirements we might infer would apply to a group that has the kind of collective unity in judgment and action that would allow it to count as an agent in its own right. For a concrete example, say that within the US Environmental Protection Agency there is a close-knit working group on carbon pricing tasked with figuring out the most economically efficient way to price CO₂ emissions. Further suppose that the members don't themselves care much about efficient carbon pricing, that this group makes all of their determinations through unstructured, open, collective deliberation, and that none of the members are aware of the pitfalls of making collective judgments in that way. We can use the empirical work sketched above to criticize this group. Since the working group behaves in a way that fails to best promote the epistemic task the group is charged with, we can use an idealistic externalist assessment to charge the group with irrationality. The group would do better in pursuing their stated task by either taking steps to structure their discussions, or perhaps even by dropping open deliberation

altogether and simply voting anonymously. Thus, there is a legitimate sense in which, epistemically speaking, they ought to take one of the latter paths.

Now I'll examine how we might use some of the empirical work above to infer various norms for groups containing individuals with joint epistemic commitments. Take a group facing a task like the one in Anderson and Holt's (1997) study, and, for whatever reason, they can only exchange information in a way parallel to the study, that is, each must privately observe the draw and publicly report her guess. But, unlike the study, say that each member gets to make a final vote after hearing all the reports, and that each member is rewarded if and only if the majority vote is correct. In this case, the practical reward would likely lead to the joint goal of getting the correct majority vote. Now return to the particular trial run discussed above. In this run, the third respondent faces a bit of a dilemma. If she's a rational individual, she would form the belief that the urn they are drawing from is the one with more light marbles. Although she drew a dark marble, the previous two individuals almost certainly pulled light marbles (that is, unless they were purposefully misrepresenting their draw, which we could suppose is unlikely). From her perspective, chances are it's the light urn, and so individual rationality dictates that she delivers a report that conflicts with the shade of marble that she drew. But obviously the group as a whole is more likely to vote correctly if each individual simply reports the color of her marble. So, the third respondent would do better at pursuing her joint epistemic goal if she reports a guess that the urn is the dark one, as opposed to the light one that individual rationality would dictate she report.

Just so that the reliabilist flavor of the previous example won't leave the impression that I think such a framework is generally superior, here is another example. Say we have a twelve-person jury tasked with determining whether Smith should be found guilty of killing Jones. During the lengthy trial, a wide range of evidence is presented. Some pieces of the evidence are rather striking, which leads every member of the jury to vividly remember those details. But, for many of the more mundane pieces of evidence, only a single member will recall the details. Let's also suppose that the prosecution presents an especially shocking piece of evidence, something that absolutely proves Smith's guilt, but this conclusive evidence was determined to be inadmissible by the judge.²³ Before the jury is sent off to do their deliberations, the judge reminds the jury to

---

23. Just to flesh out the example a bit, say that some blood from a fresh wound on Smith's right hand ended up on the murder weapon, but the sample for the DNA test that confirmed the blood belonged to Smith was illegally obtained, i.e., without consent or a warrant. During discovery, the judge rules this piece of evidence inadmissible since it was illegally obtained. Nonetheless, the prosecuting attorney, upon questioning the forensic scientist, lets it slip that blood matching Smith's DNA was found on the murder weapon, something the scientist inadvertently confirms.

completely disregard that piece of inadmissible evidence. How should the jury go about their deliberations?

As I suggested earlier, this jury's epistemic goal certainly doesn't *seem* to be the goal of determining the truth in the matter of whether Smith killed Jones. If that were the case, then they already have their answer, since the inadmissible evidence proves that Smith is the murderer. Their goal actually seems to be an evidentialist goal of sorts. The jury must determine whether the *admissible* evidence, on balance, speaks in favor of Smith's guilt. Even with these sorts of evidentialist goals, we can use results from the social sciences to make normative suggestions for this group. Notice how the structure of the case puts our jury into a somewhat similar situation to the university students in the Stasser and Titus study discussed above. Every member of the jury will recall the more sensational pieces of evidence, but each member will also have her own unique recall of various pieces of more mundane evidence. Under such circumstances, simply holding an open and unstructured discussion will predictably be a sub-optimal way for this jury to pursue its goal. If they really want to determine whether *all* of the admissible evidence weighs in favor of Smith's guilt, they would need to take some steps to ensure that the shocking evidence that they all share isn't overvalued and, similarly, that the mundane evidence that they privately hold isn't undervalued. And, perhaps more importantly, they need to take steps to ensure that the one crucial piece of inadmissible evidence doesn't sneak back in. This is one of those cases where the group may well be rational in judging that Smith is innocent, even though every member of the jury knows otherwise.²⁴

These sorts of cases expose some of the perplexing consequences of living in a social world. When we inquire in groups, the various modes of epistemic rationality can pull us in very different directions. In essence, we regularly must make value judgments concerning which epistemic goals we hold most dear for ourselves, and which we feel groups and institutions ought to do their best

---

24. Since I focus on different goals in these two examples, one might wonder how we are to determine the relevant epistemic goals when we are making an idealistic assessment. It depends. In some cases, relevance will be determined by our own concerns in doing the assessment, e.g., when we judge an educational system. In other cases, the relevant goals will be determined by a group's own charter, e.g., when we judge a funding agency like the Australian Research Council. In still others, the specific social context will determine the goal. (Arguably, scientific communities fall in the latter category.) I feel social epistemologists ought to be open to a range of factors determining which goals make for the most sensible idealistic assessments. And, just to emphasize a point made earlier, I hope the fact that I focus here on externalist assessments won't leave the impression that these kinds of assessments should be privileged in all contexts. Although I don't sketch any examples here, there are certainly cases where we will tend to privilege internalist assessments as well, especially in cases where we are trying to assess whether a group is epistemically culpable (as in the Phillip Morris example discussed earlier). I thank an anonymous referee again on this last point.

at pursuing. But once we specify these details, we can begin a largely empirical investigation of how individuals and groups ought, epistemically speaking, to behave.

5. Conclusion

Even though social epistemology is a comparatively young field of research relative to other more established areas within philosophy, it has already become seemingly disjointed. In particular, different groups of scholars have been examining what seem at first to be totally distinct aspects of group rationality, many of which seem to clash with more traditional enquiries in epistemology. In this essay, I have shown that social epistemology isn't really as disjointed as it seems. There is a core notion of epistemic rationality that relates to each of the seemingly distinct notions and that allows for a continuum between traditional epistemology and social epistemology. A nice consequence of the goal-oriented, pluralistic account that I propose is that we can make full use of a wealth of empirical research in the social and behavioral sciences in order to generate normative guidance for real-world groups. This allows social epistemology to play a similar applied role in epistemology to the role played by applied ethics in value theory. And the account does justice to the complexities and dilemmas that invariably arise when we enquire with others.

Acknowledgements

For all of their many comments, challenges, questions, suggestions, and otherwise helpful discussion, I would like to thank Kylie Bourne, David Coady, Al Hájek, Brian Hedden, Sandy Goldberg, Seumas Miller, and Mary Walker; the audiences at the Australian National University, the University of Tasmania, the (now shuttered) Centre for Applied Philosophy and Practical Ethics's workshop on "Rationality, Responsibility, and Agency in Social Groups", the 2015 meeting of the Australasian Association of Philosophy, and the 2016 meeting of the New Zealand Association of Philosophers; and a number of anonymous referees from various other journals. I especially would like to thank the two anonymous referees from this journal for all the work they put into refereeing the paper and for comments and questions that helped me greatly improve the paper. Lastly, I'd like to thank Toby Solomon for his help with editing the paper. (I obviously retain responsibility for all errors that remain.) My sincere apologies to anyone I've accidentally omitted. My work on this project was supported by the Andrew W. Mellon Foundation's Sawyer Seminar at Northwestern University entitled "The-

oretical Issues in Social Epistemology” (Grant #21300628) and the Australian Research Council’s Discovery Early Career Researcher Award (DE180101119) for my project entitled “Making More Effective Groups.”

References

- Alston, William (2005). *Beyond Justification*. Cornell University Press.
- Anderson, Lisa and Charles Holt (1997). Information Cascades in the Laboratory. *The American Economic Review*, 87(5), 847–862.
- Baker, Derek (2018). Skepticism about Ought Simpliciter. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 13, 230–252). Oxford University Press. <https://doi.org/10.1093/oso/9780198823841.003.0011>.
- Bratman, Michael (1992). Shared Cooperative Activity. *Philosophical Review*, 101(2), 327–341. <https://doi.org/10.2307/2185537>.
- Bratman, Michael (1993). Shared Intention. *Ethics*, 104(1), 97–113. <https://doi.org/10.1086/293577>.
- Bratman, Michael (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199897933.001.0001>.
- Burroughs, Michael D. and Deborah Tollefsen (2016). Learning to Listen: Epistemic Injustice and the Child. *Episteme*, 13(3), 359–377. <https://doi.org/10.1017/epi.2015.64>.
- Clarke, Steve (2007). Conspiracy Theories and the Internet: Controlled Demolition and Arrested Development. *Episteme*, 4(2), 167–180. <https://doi.org/10.3366/epi.2007.4.2.167>.
- Coady, David (Ed.) (2006). *Conspiracy Theories: The Philosophical Debate*. Routledge.
- Conee, Earl and Richard Feldman (2004). *Evidentialism: Essays in Epistemology*. Oxford University Press. <https://doi.org/10.1093/0199253722.001.0001>.
- Congdon, Matthew (2015). Epistemic Injustice in the Space of Reasons. *Episteme*, 12(1), 75–93. <https://doi.org/10.1017/epi.2014.34>.
- Copp, David (1997). The Ring of Gyges: Overridingness and the Unity of Reason. *Social Philosophy and Policy*, 14(1), 86–106. <https://doi.org/10.1017/S0265052500001680>.
- Dentith, Matthew (2014). *The Philosophy of Conspiracy Theories*. Springer. <https://doi.org/10.1057/9781137363169>.
- Dentith, Matthew (in press). Conspiracy Theories on the Basis of the Evidence. *Synthese*. Advance online publication. <https://doi.org/10.1007/s11229-017-1532-7>.
- Fallis, Don (2008). Toward an Epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology*, 59(10), 1662–1674. <https://doi.org/10.1002/asi.20870>.
- Foley, Richard (1993). *Working Without a Net: A Study of Egocentric Epistemology*. Oxford University Press.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>.
- Frost-Arnold, Karen (2014). Trustworthiness and Truth: The Epistemic Pitfalls of Internet Accountability. *Episteme*, 11(1), 63–81. <https://doi.org/10.1017/epi.2013.43>.
- Gendler, Tamar (2011). On the Epistemic Costs of Implicit Bias. *Philosophical Studies*, 156(1), 33–63. <https://doi.org/10.1007/s11098-011-9801-7>.
- Gilbert, Margaret (1989). *On Social Facts*. Princeton University Press.

- Gilbert, Margaret (2013). *Joint Commitment: How We Make the Social World*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199970148.001.0001>.
- Goldberg, Sanford C. (2013). Anonymous Assertions. *Episteme*, 10(2), 135–151. <https://doi.org/10.1017/epi.2013.14>.
- Goldman, Alvin (1986). *Epistemology and Cognition*. Harvard University Press.
- Goldman, Alvin (1999). *Knowledge in a Social World*. Oxford University Press. <https://doi.org/10.1093/0198238207.001.0001>.
- Goldman, Alvin (2010). Systems-Oriented Social Epistemology. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 3, 189–214). Oxford University Press.
- Goldman, Alvin (2014). Social Process Reliabilism: Solving Justification Problems in Collective Epistemology. In Jennifer Lackey (Ed.), *Essays in Collective Epistemology* (11–41). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199665792.003.0002>.
- Heathwood, Chris (2015). Monism and Pluralism about Value. In Iwao Hirose and Jonas Olson (Eds.), *The Oxford Handbook of Value Theory* (136–157). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199959303.013.0009>.
- Hedden, Brian (in press). Reasons, Coherence, and Group Rationality. *Philosophy and Phenomenological Research*. Advance online publication. <https://doi.org/10.1111/phpr.12486>
- Helm, Bennett (2008). Plural Agents. *Noûs*, 42(1), 17–49. <https://doi.org/10.1111/j.1468-0068.2007.00672.x>.
- James, William (1995). *Pragmatism*. Dover Publishers. (Originally printed 1907)
- Jackson, Frank and Michael Smith (2016). The Implementation Problem for Deontology. In Errol Lord and Barry Maguire (Eds.), *Weighing Reasons* (279–291). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199315192.003.0014>.
- Jönsson, Martin L. and Julia Sjö Dahl (2017). Increasing the Veracity of Implicitly Biased Rankings. *Episteme*, 14(4), 499–517. <https://doi.org/10.1017/epi.2016.34>.
- Kelly, Thomas (2002). The Rationality of Belief and Some Other Propositional Attitudes. *Philosophical Studies*, 110(2), 163–196. <https://doi.org/10.1023/A:1020212716425>.
- Kelly, Thomas (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*, 66(3), 612–640. <https://doi.org/10.1111/j.1933-1592.2003.tb00281.x>.
- Kitcher, Philip (1990). The Division of Cognitive Labor. *The Journal of Philosophy*, 87(1), 5–22. <https://doi.org/10.2307/2026796>.
- Kolodny, Niko, and John Brunero (2018). Instrumental Rationality. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2018 ed.). Retrieved from: <https://plato.stanford.edu/archives/win2018/entries/rationality-instrumental/>
- Kopec, Matthew (2018). A Pluralistic Account of Epistemic Rationality. *Synthese*, 195(8), 3571–3596. <https://doi.org/10.1007/s11229-017-1388-x>.
- Kopec, Matthew and Seumas Miller (2018). Shared Intention Is Not Joint Commitment. *Journal of Ethics and Social Philosophy*, 13(2), 179–189. <https://doi.org/10.26556/jesp.v13i2.250>.
- Lackey, Jennifer (2016). What Is Justified Group Belief? *Philosophical Review*, 125(3), 341–396. <https://doi.org/10.1215/00318108-3516946>.
- Lassiter, Charles and Nathan Ballantyne (2017). Implicit Racial Bias and Epistemic Pessimism. *Philosophical Psychology*, 30(1–2), 79–101. <https://doi.org/10.1080/09515089.2016.1265103>.

- Leary, Stephanie (2017). In Defense of Practical Reasons for Belief. *Australasian Journal of Philosophy*, 95(3), 529–542. <https://doi.org/10.1080/00048402.2016.1237532>.
- Levy, Neil (2007). Radically Socialized Knowledge and Conspiracy Theories. *Episteme*, 4(2), 181–192. <https://doi.org/10.3366/epi.2007.4.2.181>.
- List, Christian (2005). Group Knowledge and Group Rationality: A Judgment Aggregation Perspective. *Episteme*, 2(1), 25–38. <https://doi.org/10.3366/epi.2005.2.1.25>.
- List, Christian and Philip Pettit (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- Littlejohn, Clayton (2012). *Justification and the Truth Connection*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139060097>.
- Littlejohn, Clayton (2018). The Right in the Good: A Defense of Teleological Non-Consequentialism. In H. Kristoffer Ahlstrom-Vij and Jeffrey Dunn (Eds.), *Epistemic Consequentialism* (23–47). Oxford University Press.
- Mayo-Wilson, Conor, Kevin Zollman, and David Danks (2011). The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science*, 78(4), 653–677. <https://doi.org/10.1086/661777>.
- McKinnon, Rachel (2016). Epistemic Injustice. *Philosophy Compass*, 11(8), 437–446. <https://doi.org/10.1111/phc3.12336>.
- Mehta, Neil (2016). Knowledge and Other Norms for Assertion, Action, and Belief: A Teleological Account. *Philosophy and Phenomenological Research*, 93(3), 681–705. <https://doi.org/10.1111/phpr.12222>.
- Miller, Seumas (2001). *Social Action: A Teleological Account*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511612954>.
- Miller, Seumas (2010). *The Moral Foundations of Social Institutions: A Philosophical Study*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511818622>.
- Miller, Seumas (2018). Joint Epistemic Action: Some Applications. *Journal of Applied Philosophy*, 35(2), 300–318. <https://doi.org/10.1111/japp.12197>.
- Nussbaum, Martha and Amartya Sen (1993). *The Quality of Life*. Oxford University Press. <https://doi.org/10.1093/0198287976.001.0001>.
- Parfit, Derek (2011). *On What Matters*. Oxford University Press. <https://doi.org/10.1093/acprof:osobl/9780199572816.001.0001>.
- Pettigrew, Richard (2016). *Accuracy and the Laws of Credence*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198732716.001.0001>.
- Pettit, Philip (2007). Responsibility Incorporated. *Ethics*, 117(2), 171–201. <https://doi.org/10.1086/510695>.
- Pigden, Charles (2007). Conspiracy Theories and the Conventional Wisdom. *Episteme*, 4(2), 219–232. <https://doi.org/10.3366/epi.2007.4.2.219>.
- Quine, W.V.O., and J.S. Ullian (1970). *The Web of Belief*. Random House.
- Rinard, Susanna (2017). No Exception for Belief. *Philosophy and Phenomenological Research*, 94(1), 121–143. <https://doi.org/10.1111/phpr.12229>.
- Rinard, Susanna (in press). Believing for Practical Reasons. *Noûs*. Advance online publication. <https://doi.org/10.1111/nous.12253>.
- Sanger, Lawrence M. (2009). The Fate of Expertise after Wikipedia. *Episteme*, 6(1), 52–73. <https://doi.org/10.3366/E1742360008000543>.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199299508.001.0001>.
- Searle, John (1990). Collective Intentions and Actions. In Philip R. Cohen, Jerry Morgan, and Martha Pollack (Eds.), *Intentions in Communication* (401–415). MIT Press.

- Shah, Nishi (2006). A New Argument for Evidentialism. *The Philosophical Quarterly*, 56(225), 481–498. <https://doi.org/10.1111/j.1467-9213.2006.454.x>.
- Stasser, Garold and William Titus (1985). Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion. *Journal of Personality and Social Psychology*, 48(6), 1467–1478. <https://doi.org/10.1037/0022-3514.48.6.1467>.
- Stich, Stephen (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Bradford Books.
- Strauss, Valerie (2014, September 12). Proposed Texas Textbooks Are Inaccurate, Biased and Politicized, New Report Finds. *The Washington Post*. Retrieved from: <https://www.washingtonpost.com/news/answer-sheet/wp/2014/09/12/proposed-texas-textbooks-are-inaccurate-biased-and-politicized-new-report-finds/>
- Sylvan, Kurt (2012). Truth Monism Without Teleology. *Thought: A Journal of Philosophy*, 1(3), 161–169. <https://doi.org/10.1002/tht3.26>.
- Sylvan, Kurt (2018). Veritism Unswamped. *Mind*, 127(506), 381–435. <https://doi.org/10.1093/mind/fzw070>.
- Sylvan, Kurt (in press). Reliabilism Without Epistemic Consequentialism. *Philosophy and Phenomenological Research*. Advance online publication. <https://doi.org/10.1111/phpr.12560>
- Sunstein, Cass and Reid Hastie (2015). *Wiser: Getting Beyond Group Think to Make Groups Smarter*. Harvard Business Review Press.
- Tiffany, Evan (2007). Deflationary Normative Pluralism. *Canadian Journal of Philosophy*, 37(5), 231–262. <https://doi.org/10.1353/cjp.0.0076>.
- Tollefsen, Deborah (2009). Wikipedia and the Epistemology of Testimony. *Episteme*, 6(1), 8–24. <https://doi.org/10.3366/E1742360008000518>.
- Tollefsen, Deborah (2015). *Groups as Agents*. John Wiley & Sons.
- Toribio, Josefa (in press). Accessibility, Implicit Bias, and Epistemic Justification. *Synthese*. Advance online publication. <https://doi.org/10.1007/s11229-018-1795-7>
- Williamson, Timothy (2002). *Knowledge and Its Limits*. Oxford University Press. <https://doi.org/10.1093/019925656X.001.0001>.
- Zollman, Kevin (2013). Network Epistemology: Communication in Epistemic Groups. *Philosophy Compass*, 8(1), 15–27. <https://doi.org/10.1111/j.1747-9991.2012.00534.x>.