published in *Episteme, A Journal of Social Epistemology* 1: 11-22 (2004)

GROUP KNOWLEDGE VERSUS GROUP RATIONALITY:

TWO APPROACHES TO SOCIAL EPISTEMOLOGY

Alvin I. Goldman

1. Alternative Agendas for Social Epistemology

Social epistemology is a many-splendored subject.[1] Different theorists adopt different approaches and the options are quite diverse, often orthogonal to one another. The approach I favor is to examine social practices in terms of their impact on knowledge acquisition (Goldman 1999). This has at least two virtues: it displays continuity with traditional epistemology, which historically focuses on knowledge, and it intersects with the concerns of practical life, which are pervasively affected by what people know or don't know. In making this choice, I am not blind to the allure of alternative approaches. In this paper I explain and motivate the knowledge-centered approach by contrasting it with a newly emerging alternative that has a definite appeal of its own. According to this alternative, the chief dimension of social epistemological interest would be rationality rather than knowledge.

What is knowledge? As construed here, knowledge is true belief arrived at via appropriate means, methods, or sources. In particular, appropriate methods or means are reliable methods or means, ones with the general property of generating true belief (Goldman 1986, 1992, 2002b).[2] Knowledge-producing methods or sources can be either internal or external. On the internal side, they include perception, memory, and reasoning. On the external side, they include the testimony or discourses of others. Social epistemology concentrates on the external side, on patterns of interaction among epistemic agents that influence the beliefs of each. How do these networks or patterns of communication influence the quality of beliefs thereby generated? To what extent do they promote true or accurate belief, and hence knowledge, and to what extent do they promote ignorance or error? This focus on truth-related dimensions of epistemic systems is a core part of my favored approach.

Other conceptions of social epistemology minimize or reject the role of truth. Why should it be given pride of place? First, knowledge is a headline topic in traditional epistemology; why shouldn't it be equally pivotal in social epistemology? Second, the verb 'know' in its everyday, propositional use is a factive verb: one knows that p only if p is true. Virtually all mainstream epistemologists and semanticists agree on this. Third, quite apart from the term 'knowledge', people are inveterate seekers of truth. They desire truth either from intrinsic curiosity or from the recognition that truth facilitates success in practical encounters with the world (Goldman 1999, 72-74). This link to practical success makes truth possession a general desideratum and the conditions of its attainment a natural topic of investigation.

Suppose, then, that social epistemology starts down the path just charted, dedicated to studying social practices that promote or inhibit truth acquisition. A fork in the road is soon encountered. Who are the prospective knowers that social epistemology should study? Are they individuals or groups? Some social epistemologists argue that the best units to study are communities, groups, or collectivities – social entities of one sort or another. This raises the question: what is the ontological status of these alleged entities? Are there such entities to which it is appropriate to ascribe intentional and epistemic states? In earlier work (Goldman 1999) the fork I chose was to study individual knowers. Without categorically denying the existence of social collectivities, I set them aside. Social epistemology could be pursued by focusing on social transactions between individuals, or networks of relationships among individuals. The time is ripe, however, for a closer look at group knowledge. Many contemporary philosophers argue the case for treating institutions, organizations, and associations of individuals as proper subjects of intentional and epistemic states (Gilbert 1989; Tuomela 1995; Searle 1995; Bratman 1999; Nelson 1993). In common parlance, certainly, organizations are treated as subjects for knowledge attribution. In the wake of 9/11, there has been much commentary on what the C.I.A. and the F.B.I. did or didn't know about terrorist plans before the event itself. The following remark by Sandy Berger, national security adviser under President Clinton, is illustrative of a readiness to apply the "knowledge" predicate to an organization: "We've learned since 9/11 that not only did we not know what we didn't know, but the F.B.I. didn't know what it did know."

All this relates to an agenda for social epistemology under the knowledge-oriented conception. What about the rationality-oriented conception? Rationality is also an intensively investigated epistemic concept and group rationality seems like an eminently appropriate subject for social epistemology. What kind of thing is group rationality, and how (if at all) is it realizable? The detailed discussion begins by examining some recent intriguing proposals about group rationality. These proposals are not explicitly articulated under the heading of social epistemology, but up to a point they sit fairly comfortably within this niche. Nonetheless, I shall argue that group rationality – at least as conceived in the manner to be discussed – would be a sub-optimal conceptualization of social epistemology. This conception will be used as a counterpoint to the knowledge-oriented conception, in order to motivate the latter.

2. Rationality and Judgment Aggregation

The approach to group rationality to be explored is due to Christian List and Philip Pettit (2002, forthcoming; Pettit, forthcoming). It builds on what others have called the "doctrinal" or "discursive" paradox. The paradox arises in frameworks in which sets of judgments are to be aggregated. Judgments are modeled on acts of assent or dissent, assertion or denial. They differ from credences (subjective probabilities) in being binary, on or off, affairs. The category of judgments, as treated by List and Pettit, includes both decisions about what to do as well as judgments of fact. To illustrate the factual category, consider judgments of a promotions committee of whether a candidate is formally qualified for the promotion, whether the candidate has sufficient experience, whether he or she is likely to establish good working relations with colleagues after the

promotion, and finally whether he or she merits promotion. Another example is where a three-judge court has to decide whether a defendant is liable under a charge of breach of contract. These factual or semi-factual judgments are what, in my opinion, bring the topic into the domain of social epistemology.

The doctrinal or discursive paradox arises as follows (Kornhauser and Sager, 1986; Kornhauser and Sager, 1993; Chapman, 1998; Brennan, 2001). Suppose that a committee or other formal body has to make a judgment based on certain related issues. In the breach of contract case, legal doctrine says that a court should find against a defendant if and only if it finds, first that a valid contract was in place, and second that the defendant's behavior was such as to breach a contract of that kind. In other words, the court should accept the "conclusion" that the defendant is liable if and only if it accepts the two "premises" that a valid contract was in place and the defendant breached the contract. Imagine that the three judges, 1, 2, and 3, vote as follows on the two premises and the conclusion of whether the defendant is liable. The 'yes' or 'no' in any row represents the relevant judge's judgment to accept or reject the indicated premise or conclusion.

	Valid contract?	Breach?	Liable?
1	yes	no	no
2	no	yes	no
3	yes	yes	yes

Given these individual opinions, there are two ways the court might aggregate them. It might let the judges do their individual reasoning from the premises to the conclusion (liability), and then aggregate their conclusion votes on, say, a majority basis. This may be called the "conclusion-driven" approach. Under this approach the liability proposition would be rejected, because there would be only one vote (from judge 3) in favor of liability. Alternatively, the judges' votes on each individual premise could be aggregated, and those aggregate premise judgments would determine the aggregate conclusion. Since there is a two-to-one majority for each premise, both premises would be accepted and hence the conclusion would be accepted as well. This may be called the "premise-driven" approach. The doctrinal, or discursive, paradox consists in the fact that the two procedures yield different outcomes, and it isn't clear which is right. The paradox generalizes in many different ways.

What worries List and Pettit is that if we go the normal, majoritarian way in aggregating a collective set of judgments from individual sets of judgments, we may end up with a collective set that is irrational. While every individual is perfectly rational in their pattern of assent to premises and conclusion, the collectivity would not be rational in its pattern of assent if it followed the procedure of majority voting on each premise and on the conclusion. It would fail to be rational because it would endorse the premises but

reject the conclusion they deductively support. List and Pettit point out that there is a tension between two plausible demands that might be made on the aggregation of judgment. The first demand is that aggregation should be responsive to individual members' views on each of the judgments involved. The second is that in aggregating judgment a group should reach a set of collective judgments that is rational at the collective level. These two demands sometimes conflict.

These considerations suggest that if we take the sets of views held among a group of people on a range of propositions, then even if those sets of views each satisfies rationality constraints, the set of collective views derived by majority voting on each issue may not satisfy the same rationality constraints. Furthermore, a danger looms that any procedure that is akin to majority voting in certain ways will be vulnerable to this same threat of collective irrationality. Indeed, List and Pettit go on to formulate and prove the impossibility of finding a procedure for aggregating judgment that satisfies rather undemanding (by their lights) rationality conditions.

Three minimal conditions are proposed for generating rational collective judgments. Consider a procedure or function, F, to be called a judgment aggregation function, whose input is a profile of personal sets of judgments and whose output is a corresponding collective set of judgments by the group as a whole. The three minimal conditions are as follows:

Universal Domain. A judgment aggregation function, F, should accept as admissible input any logically possible profile of personal sets of judgments, provided that each person's set of judgments (on the issues in question) is complete, consistent, and deductively closed.

Anonymity. The collective set of judgments that is yielded by F should be invariant under any permutation of individuals. In other words, no one individual's judgment should be given special weight in determining the collective judgments; all persons should be treated even-handedly by the aggregation function.

Systematicity. The same pattern of dependence of collective judgments on individual ones should hold for all propositions and profiles in the domain of F. For example, if majority rule or some other procedure of aggregation is used to decide the collective view on some propositions, then that procedure should be used to decide the collective view on all.

List and Pettit prove that it is impossible to satisfy all of these conditions. There is no judgment aggregation function F generating complete, consistent and deductively closed collective sets of judgments that satisfies the requirements of universal domain, anonymity and systematicity. Many readers will recognize strong parallels between this theorem and Kenneth Arrow's (1963) famous proof of a similar impossibility result in the domain of social choice. The current theorem, however, is a wholly distinct one.

Given this impossibility result, how do groups manage routinely to aggregate judgments? How do they avoid the problem identified by the impossibility result? The answer is that, strictly speaking, they don't avoid it. Either they don't quite live up to the rule of rationality (the rule of completeness, consistency and deductive closure) or they sacrifice responsiveness to their members' views. List and Pettit identify several types of strategies for "collectivizing reason", as they put it. These include relaxing universal domain, relaxing anonymity, relaxing systematicity, etc. One might say – this is my own formulation – that the study of social epistemology under the aggregation-rationality conceptualization becomes the study of how, and to what extent, the demands of group rationality can be accommodated, though partly compromised, in collective judgmental settings.

3. Groups as Intentional and Personal Subjects

In another paper Pettit (forthcoming) builds on the impossibility result to argue for the existence of a certain type of group that he calls an "integrated collectivity" or a "social integrate." Such groups, he says, qualify as intentional subjects, as the bearers of intentional attitudes like belief and intention, and they qualify as such by the usual criteria employed for such subjects.

In working up to this idea, Pettit begins by pointing out that some groups are not mere unorganized collocations but are united by a common purpose. Purposive groups will almost inevitably confront discursive dilemmas and will be under pressure to collectivize reason. Such groups will be faced across time with sets of rationally connected issues forcing them to choose between maximizing responsiveness to the views of their members, on the one hand, and ensuring rationality at the collective level, on the other. Groups in this situation are what Pettit calls "social integrates". He then proceeds to argue that these groups are distinct intentional subjects that should be counted over and above their members in any inventory of intentional subjects.

In defending this thesis, Pettit appeals to familiar, rationality-based conceptions of intentional subjects. A system is an intentional subject just in case it preserves intentional attitudes over time, and forms, unforms, and acts on those attitudes in a rationally permissible manner. In short, an intentional system is a system that displays a certain rational unity (Pettit 1993). Purposive groups, says Pettit, will generally be constrained to be such systems. Moreover, they will be forced, at least some of the time, to make judgments at odds with many of their members, so their attitudes will not match those of their "natural" members. Hence, we are warranted in regarding such a group as a distinct intentional system, over and above its natural members. True, such groups do not have their own faculties of perception or memory, and they will not be capable of existing in the absence of their natural members. But Pettit does not regard these as disqualifying characteristics. So here we have a rationale for treating a certain type of group as an ontologically distinct entity and the bearer of doxastic or epistemic states. This seems like a serious candidate for the sort of social knower that social epistemology should take an interest in.

4. Judgment Aggregation and the Truth Desideratum

It bears emphasis that Pettit and List do not themselves advertise their proposals as specimens of social epistemology. The term 'epistemology' nowhere appears in these writings, and whether they would be happy with the label remains to be seen. But others might regard this as an appropriate classification. Certain styles of epistemology, especially probability-based approaches like Bayesianism, characteristically take rationality as their guiding beacon. Similarly, philosophers of science commonly view rationality, rather than knowledge or truth, as the paramount norm or goal of science. Although collective rationality of the type explored by Pettit and List is not the image of rationality customarily employed in these epistemological enterprises, one might be attracted to their form of rationality as a prototype for social epistemology, especially the sector of social epistemology concerned with group epistemic agents. How should this idea be received?

If we examine List and Pettit's conception of rational collective judgment, we find an important assumption of equality of weight among group members. This assumption is embedded in the anonymity condition. Anonymity says that collective judgments should be invariant under any permutation of individuals. No one individual's judgment should be given special weight over others. Now this is an eminently plausible condition in any democratic framework, and both Pettit and List are, primarily or substantially, democratic theorists. It is hardly surprising, then, that they and other political theorists should find an assumption of equal voice, or non-discrimination, appealing. The anonymity assumption is a staple of social choice theory.[3] It is much less, clear, however, that such an assumption is appropriate for epistemology, even social epistemology.

Suppose that a local weather bureau hires five meteorologists and wishes to amalgamate their weather judgments daily to make the best prediction possible of the following day's weather. What amalgamation practice should it adopt? More specifically, suppose that it will make a binary prediction choice: rain or non-rain. Further suppose that each of two of the meteorologists has a 90 percent success rate in making such judgments and each of the other three has a 60 percent success rate. Should the bureau determine its collective prediction by letting the five make their individual predictions (independently) and adopting the majority opinion? That would not be optimal. Given their several competences at forecasting, this would yield a probability of correctness for the bureau of .877. That success rate could be trumped by simply letting one of the two more expert meteorologists dictate the bureau's prediction. This would yield a correctness probability of .900. Even this practice can be trumped, however, by another amalgamation scheme, a weighted voting scheme that accords greater weight to the more expert weather forecasters than to the others. An optimal weighted voting scheme would yield a correctness probability of .927 for the bureau.[4] So, clearly, the most sensible amalgamation, the most epistemically appropriate amalgamation, would not be equal weighting. This suggests that when it comes to epistemic matters, truth is of paramount importance. Rational aggregation principles of the sort List and Pettit present

have insufficient bearing on the project of truth determination. It is precisely this feature that disables the conception as a suitable prototype for social epistemology.

Similar remarks apply to purposive organizations like intelligence agencies. No reasonable intelligence agency would restrict itself to assigning equal weights to different agents in all matters. Suppose two agents gather intelligence on the same topic from different sources. A higher-level analyst who receives these reports need not assign them equal weights. The analyst may have good reason to think that one of the sources is more reliable than the other, or that one agent is more skilled and hence reliable than the other. In either case, equality of weighting is by no means imperative.

Turning to the question of collective agents, Pettit makes an interesting case for what he calls "integrated collectivities", and this is undoubtedly one kind of collective epistemic subject that social epistemology should study. But it isn't the only kind. I shall discuss two other kinds of collective entities that are worth investigating from a social epistemological standpoint.

The first additional kind of collective entity may be called a summative collectivity (Quinton 1975). This is the kind of collectivity whose intentional states are little more than summations of the intentional states of their diverse members. An instructive illustration of such collectivities are honey bee colonies.

The social epistemology of honey bees has been intensively studied by Thomas Seeley (1995), with striking results.

> [A] honey bee colony operates as a thoroughly integrated unit in gathering its food. It monitors the flower patches in the countryside surrounding its hive; it distributes its foraging activity among these patches so that nectar and pollen are collected efficiently, in sufficient quantity, and in the nutritionally correct mix; and it properly apportions the food it gathers between present consumption and shortage for future needs. In addition, a colony precisely controls its building of beeswax combs for honey storage, strictly limiting this costly process to times of clear need. (1995, p. 6)

Seeley sometimes goes so far as to describe the colony as a "superorganism" (Seely 1989). However, I would not say that honey bee colonies are "integrated" units in exactly Pettit's strong sense. Nonetheless, there is enough integration within the colony to induce this expert in the field to speak of it as a kind of collective entity. Let us examine one aspect of honey bee social epistemology: how they acquire information about food sources from other members of the hive.

At any given time, forager bees can be classified as either employed or unemployed. Employed foragers are busy carrying food (either nectar or pollen) from flowers to the hive. Unemployed foragers do not have a current flower source that is profitable. They therefore go to a particular area of the hive – the dance floor – and watch employed foragers do a waggle dance, by means of which they "advertise" their

profitable food sources to hive-mates who observe them. Each dance communicates the distance, orientation, and profitability of a food source that the dancer recently visited and exploited. The longer the dance – i.e., the more waggles it features – the greater is the profitability of the source. Because there are many dancers on the floor, an unemployed forager has an opportunity to become well informed about the various food sources exploited by her colony. In principle, an observing honey bee could compare the various food sources and choose the most profitable one as her next work site. Do unemployed foragers do so, or do they adopt a new forage site based on more limited information?

Seeley discovered that the latter is the case. Unemployed foragers do not acquire information from a large number of dancers before setting out to search for a new food source. Indeed, each search is preceded by just a single bout of following one dancer for numerous waggle runs. It is rare for an unemployed forager to follow a dance from start to finish. For one thing, few bees will be near enough to a dancer when she begins dancing -- observation occurs by listening to the dancer's wing vibrations with antennae extended in close proximity. In addition, once an observer follows enough of a dance to learn that the source is substantially profitable, she breaks away from the dancer, scrambles out of the hive, and searches for the food source indicated by the dance just observed. So the observer seems to form a belief in the proposition, "At such-and-such a location is an adequately profitable food source."

Interestingly, the limited information-seeking behavior of the bee is actually better for the colony. If all unemployed foragers compared all dances and chose the single best source, the colony would be overinvested in that source, which might be highly profitable but small or short-lived. This would also leave the colony underinformed should foraging conditions change. As things stand, the bees distribute themselves among all reasonably profitable forage sites. Indeed, the colony achieves high foraging success by allocating foragers in proportion to the profitability of sources, which is reflected in the duration of the dance for each source. If we are willing to treat the colony as a summative collectivity that is the subject of intentional states, we might regard the colony as "believing" the proposition: "This is a profitable way for <u>us</u> to distribute our foraging activity". This might be regarded as a summative belief, not in the sense that the same aggregate proposition is believed by an "integrated" collectivity (in Pettit's sense), nor in the sense that each of the foraging bees believes this proposition, but in the sense that the sum-total of "local" beliefs of individual foragers "adds up" to a (distributed) belief by a summative collectivity in the aggregate proposition.

Here is an analogous example from an entirely different domain, the domain of human politics. According to the epistemic theory of democracy, at least in its traditional form, a vote in an election expresses the judgment, "Candidate X will best serve the interests of the electorate as a whole." Contemporary political theory views it as unlikely, however, that voters generally cast their votes on the basis of what serves the common interest. They are more likely to be cast on the basis of what voters believe to be in their <u>own</u> interest, at least if "interest" is allowed to subsume anything the voter approves of for whatever reason (including altruistic reasons). So beliefs of individual

citizens in propositions of the form "Candidate X will best serve my interest" are the beliefs that guide those citizens' votes. This type of proposition is what I have called a "core voter proposition" (Goldman 1999, chap. 10). I also show that if all voters in a majoritarian election have core voter knowledge -- i.e., true belief in their core voter proposition -- then their votes will ensure a certain positive outcome from the vantage point of democracy. Namely, it will ensure that a majority of voters will have their interests better served by the winning candidate than by the loser. That outcome isn't guaranteed without full core knowledge. This shows the importance of knowledge in a democratic setting.

If we wish, in this case, we can also "sum" the knowledge of the individual voters into an aggregate item of knowledge on the part of the electorate construed as a summative collectivity. Under full core knowledge, each individual voter knows that candidate X (or Y) is best for him or her. If we then select the candidate so chosen by a majority of voters -- call that candidate X -- then we can say that the electorate knows the "aggregate" proposition that X will best serve the interests of a majority of voters. I concede that this talk of aggregate knowledge is something of a creative fiction. But it seems to be an appealing and instructive fiction.

Recall the second half of Sandy Berger's dictum: "the F.B.I. didn't know what it did know." Let us analyze this carefully. Berger is saying that under some conception of "the F.B.I.", the F.B.I. did have knowledge highly pertinent to 9/11 and under another conception of "the F.B.I." it lacked this knowledge. What is the conception under which it had this knowledge? What he presumably means is that knowledge was possessed in a distributed fashion by the collection of agents in the field (e.g. in Minneapolis and Phoenix), in particular, by those agents who were each aware that this or that future hijacker was engaged in flight training. In this distributed or summative way, the F.B.I. had knowledge of the flight-training pattern of the several hijackers. Thus, under one conceptualization of the F.B.I. employed by Berger, it is a collectivity of members whose distributed knowledge can be thought of as the aggregate knowledge of a single entity. What is the conception under which it lacked such knowledge, under which it was ignorant? I turn to this next.

5. Knowledge in Hierarchical Groups

If Berger's quip is true – and it has a strong ring of truth – the entity that lacked a certain piece of knowledge cannot be the same entity that also possessed the very same knowledge. So what is the F.B.I. entity that was ignorant of what the summative collectivity knew?

A very important kind of group or organization is what I'll call a hierarchical group. This type of group does not make decisions by consulting its membership, in a majoritarian or semi-majoritarian manner. Decision-making power is vested in a single individual or directorship, and, to a first approximation, what this individual or directorship decides is the decision of the organization. Of course, such organizations often confer central decision-making power on an individual by some sort of vote. A

CEO is appointed by vote of the board of directors. A President is elected by vote of the electorate. But once this appointment is made, and for a certain period, the chosen official acts for or on behalf of the organization. It is therefore unsurprising that we find it natural, up to a point, to identify the intentional states of the organization with the (organizationally relevant) intentional states of the decision-maker. The organization's plans and intentions are her plans and intentions – at least her plans and intentions vis-à-vis the organization's actions (not her own private actions). If America's aims, intentions, or attitudes are in question on a foreign policy issue, those plans, intentions, or attitudes can be identified with the President's plans, intentions, or attitudes. One won't take into account the spectrum of attitudes on that foreign policy issue by the electorate at large. The President's stance on the issue may be wildly unpopular, but his view constitutes America's view, at least the government's view.

This simple formulation needs qualification, of course. A President does not have an opinion or stance on every issue that concerns some arm of the government's executive branch. Certain foreign policy issues, for example, may not come to the President's attention. The government's attitude may be defined by that of the Secretary of State, or of some mid-level official in the region. By delegation or chain of authority, it is often a lower-level official whose attitude fixes or determines the government's attitude. Thus, a government's attitude should not invariably be identified with the attitude of the highest-ranking official. The important point, however, is that a group's intentional attitude is not some aggregate function of attitudes of all members of the group, e.g., not all citizens nor all executive branch employees. It is the relevantly authorized decision-makers whose attitudes most naturally represent, or stand in for, the corporate entity's attitudes.[5]

Similarly, consider the criminal law concerning corporations. In criminal matters, guilt typically requires the presence of a mens rea (mental state) element. What is the mens rea element in the case of a corporation accused of a crime? Which individual or individuals must be shown to have had the pertinent knowledge or intent in order for the corporation to be guilty? Certainly it isn't all employees of the corporation. Nor is it, necessarily, the CEO. In general, it is the officer or officers of the corporation that comprised the "control" group with respect to the act in question. It might be, for example, the chief financial officer.[6]

All this applies quite naturally to Berger's attribution of ignorance to the F.B.I. The F.B.I. that lacked knowledge was constituted by the Bureau's decision-makers in Washington. It was the ignorance of the Director, and the ignorance of lower-level officials in Washington, that comprised the ignorance of the Bureau; they were the "seat" or "substrate" of the corporate ignorance. Those relatively high-level officials were in a "control" position vis-à-vis the potential actions in question, so it was their epistemic state that was critical to the epistemic state of the Bureau construed as a corporate entity. This construal of the corporate epistemic state completes our rendition of Berger's quip so that it comes out true. It shows that there was one entity, a summative collectivity, that (plausibly) had a certain body of knowledge, while another entity, a hierarchical group, lacked that knowledge – and both of them can be called "the F.B.I." More

felicitously, perhaps, there are two ways of conceptualizing the referent of "the F.B.I.", one in which knowledge was present and the other in which knowledge was absent.

6. Testing the Two Approaches

The F.B.I. was guilty of a gargantuan failure, a social-epistemic failure. It would be an impoverished approach to social epistemology that could not identify the pre-9/11 ignorance of the F.B.I. as some sort of epistemic debacle. Since it was substantially a failure to "put the pieces together", or "connect the dots," where individual pieces or dots were at least within the ken of separate agents, it was not just an epistemic failure but a social-epistemic failure. A minimal test of adequacy for a social epistemology is to identify and explain the type of misfire or misexecution that occurred. How do the two approaches under discussion fare on this test?

The problem should be divided into two parts: identifying the flaw in the outcome and identifying flaws in the process responsible for the outcome. The first part is straightforward under the knowledge approach. The F.B.I. (at least in the hierarchical conception) was ignorant of an important fact, a fact that it was a crucial part of their mission to know. This was the fact that several terrorist suspects were engaged in similar training involving airplanes. How does the rationality approach handle the outcome in question? The rationality approach does not have a natural treatment of this. It does not feature knowledge as a basic desideratum or ignorance as a basic defect. This is a straightforward deficiency in the account when viewed as a comprehensive approach to social epistemology (which is not how its proponents present it, of course). But let us set this aside; perhaps it can say useful things about procedural flaws.

Turn, then, to the question of process. Under the knowledge (or "veritistic") approach to social epistemology, practices are evaluated by their effectiveness in promoting the attainment of knowledge, especially knowledge in the heads of those who seek it, or have an interest in it. What went wrong in the pre-9/11 situation? How were the practices deficient? The precise details, of course, need not concern us (nor are they publicly known). What was dramatically deficient, however, was the F.B.I.'s communication system, the network linking reports from agents in the field to higher-level analysts in Washington. Also deficient, apparently, was the Bureau's system for coding and pooling information from multiple sources. Now, these kinds of communication and information-processing structures are among the topics that the veritistic conception of social epistemology places squarely and prominently on its agenda (see Goldman 1999, chap. 6). Both information technology and the organization of communication networks are core arenas to be studied under the umbrella of veritistic social epistemology. In the present case, at least one Bureau official with appropriate decision-making authority had to receive messages from the various field agents, had to believe those messages, and had to pool or amalgamate them into a larger pattern. Where, exactly, the system broke down is unclear, but the deficiency or combination of deficiencies is readily subsumed into the veritistic approach.

What can the rationality approach say? Under this approach, the failure must be explained in terms of its key ingredients: rationality and/or responsiveness. How would such an explanation go? Rationality might have failed at either the individual level or the collective level. Was there a rationality failure at the individual level, according to the List-Pettit rationality constraints? For any pre-selected set of interconnected propositions, they require each individual's judgments to be complete, consistent, and deductively closed. Completeness here means that the individual either assents to a proposition or assents to its negation. Consistency and closure have their standard definitions. Is there any reason to suppose that individual agents or officials were irrational in the List-Pettit sense? That depends on what is chosen as the pre-selected set of propositions for each individual. Here is the problem. Field agent A_1 knew about the activities of suspected terrorist T_1 but not about terrorist T_2; the converse was true of field agent A_2. Agent A_1 was presumably agnostic (withheld judgment) about the proposition "T_2 is taking flight training lessons." Similarly, agent A_2 was agnostic (withheld judgment) about proposition "T_1 is taking flight training lessons". Now withholding judgment vis-à-vis P is neither assenting to P nor assenting to its negation. So if we put all the indicated propositions in a field agent's pre-selected set of judgments, he/she will have violated the List-Pettit rationality requirement, by failure of completeness. But why should we make this pre-selection? If anything, the requirements might be faulted for imposing too stringent a constraint on rationality. So there is no clear case for finding a compelling rationality flaw at the individual level.

Was there a rationality violation at the collective level? This is very hard to determine, because it isn't clear what defines the collective-level judgments for the Bureau. As we have explained, the Bureau isn't an integrated collectivity in Pettit's sense, and there is no aggregation procedure of the kind List and Pettit discuss. So it's a bit difficult to fit this into their scheme. Suppose, however, that we use principles appropriate to the hierarchical collectivity that the Bureau consists in. Let's identify the Bureau's judgments with those of the Director, or those of any appropriate mid-level official that may have been responsible for assimilating information about domestic terrorism threats and decision-making relevant to them. It is not clear in the public record (as far as I know) that any official of this sort had an irrational set of judgments in the List-Pettit sense. Conceivably some such official was apprised of the assorted "dots" and failed to "connect" them; that would be a failure of closure. But it's not clear that this transpired.

Perhaps the Bureau's failure consisted not in irrationality but in unresponsiveness. The whole idea of rational aggregation is that group judgments should be responsive to members' judgments. Unresponsiveness sounds like an apt description of the F.B.I. situation, but does unresponsiveness in the List-Pettit sense really apply? Responsiveness is making the group's judgment conform to the members' judgments, but what does conformity consist in? That has a definite meaning only if a specific aggregation function is operative, but the List-Pettit approach does not choose such a function. (In fact, it says that there can be no function that meets the stipulated conditions.) Furthermore, in a case where one member (a field agent) assents to a certain proposition

and all (or many) other members are agnostic, it's just not clear what a reasonable aggregation function would dictate.

Finally, I return to a point mentioned earlier. The anonymity constraint, or perhaps the systematicity constraint, implies that a unique group judgment should be determined when a specified distribution of individual judgments obtains, no matter what else is true. This may sound appealing, but does it make good epistemic sense? Not really. All sorts of varying reasons, evidence, or causes may underlie a given individual's judgment. His assent to a specified proposition may result from direct observation, from solid evidence, or from sheer fantasy. If another person or a group is deciding how much weight to assign to an individual's judgment, these factors ought to be taken into account (to the extent that they are known). A principle that abstracts from such variations in the causes or grounds of an individual's judgment cannot be a sound epistemological principle. Thus, even if there is a way of explaining the F.B.I.'s deficiency in terms of List-Pettit unresponsiveness, I have to question whether this approach to the subject is epistemologically sound. Again, it must be recognized that Pettit and List themselves have not explicitly set about to resolve all epistemological issues in terms of their framework. Their project is much more modest. But if anyone hoped to erect a comprehensive approach to social epistemology on the basis of the framework they have articulated, the foregoing problems would have to be addressed.

In summary, the knowledge-centered approach passes the F.B.I. adequacy test quite easily, whereas the rationality-centered approach struggles. Of course, this is just one case, by no means drawn at random. And we should keep in mind that the rationality-aggregation approach was not advanced as a general program for social epistemology. Nonetheless, the present comparison instructively highlights the strengths of the knowledge-based, or "veritistic," approach to social epistemology.[7]

References

Arrow, Kenneth (1963). Social Choice and Individual Values. New York: Wiley.

Bratman, Michael (1999). Faces of Intention: Selected Essays on Intention and Agency. Cambridge: Cambridge University Press.

Brennan, G. (2001). Collective coherence? International Review of Law and Economics 21: 197-211.

Chapman, B. (1998). More easily done than said: Rules, reason and rational social choice. Oxford Journal of Legal Studies 18: 293-329.

Gilbert, Margaret (1989). On Social Facts. Princeton, NJ: Princeton University Press.

Goldman, Alvin (1986). Epistemology and Cognition. Cambridge, MA: Harvard University Press.

Goldman, Alvin (1992). Liaisons: Philosophy Meets the Cognitive and Social Sciences. Cambridge, MA: MIT Press.

Goldman, Alvin (1999). Knowledge in a Social World. Oxford: Oxford University Press.

Goldman, Alvin (2002a). What is social epistemology? A smorgasbord of projects. In Goldman (2002b).

Goldman, Alvin (2002b). Pathways to Knowledge: Private and Public. New York: Oxford University Press.

Goldman, Alvin (2004). The need for social epistemology. In B. Leiter, ed., The Future of Philosophy. New York: Oxford University Press (pp. 182-207).

Kornhauser, L.A. and Sager, L.G. (1986). Unpacking the court. Yale Law Journal 82.

Kornhauser, L.A. and Sager, L.G. (1993). The one and the many: Adjudication in collegial courts. Califronia Law Review 81: 1-59.

List, Christian and Pettit, Philip (2002). Aggregating sets of judgments: An impossibility result. Economics and Philosophy 18: 89-110.

List, Christian and Pettit, Philip (forthcoming). Aggregating sets of judgments: Two impossibility results compared. Synthese.

Nelson, Lynn Hankinson (1993). Epistemological communities. In L. Alcoff and E. Potter, eds., Feminist Epistemologies. New York: Routledge.

Pettit, Philip (1993). The Common Mind: An Essay on Psychology, Society and Politics. New York: Oxford University Press.

Pettit, Philip (2003). Groups with minds of their own. In F. Schmitt, ed., Socializing Metaphysics. Rowman and Littlefield.

Quinton, Anthony (1975). Social objects. Proceedings of the Aristotelian Society 75.

Searle, John (1995). The Construction of Social Reality. New York: Free Press.

Seeley, Thomas (1989). The honey bee colony as a superorganism. American Scientist 77: 546-553.

Seeley, Thomas (1995). The Wisdom of the Hive: The Social Physiology of Honey Bee Colonies. Cambridge, MA: Harvard University Press.

Sen, Amartya (1970). Collective Choice and Social Welfare. London: Holden Day.

Shapley, Lloyd and Grofman, Bernard (1984). Optimizing group judgmental accuracy in the presence of interdependence. Public Choice 43: 329-343.

Tuomela, Raimo (1995). The Importance of Us. Stanford, CA: Stanford University Press.

---

[1] For surveys of different styles and rationales of social epistemology, see Goldman 2002a, 2004.

[2] A more canonical approach to knowledge views it as justified true belief (where the justification isn't defeated). There is no essential conflict between the justification account and the reliability account, because justification may consist in the use of reliable methods (Goldman 1986, chaps. 4-5; 1992).

[3] For example, see Sen 1970, p. 68.

[4] Mathematical details of this case are elaborated in (Goldman 1999, pp. 81-82). The example and the analysis behind it originate with Shapley and Grofman (1984).

[5] The present account bears some similarity to Raimo Tuomela's (1995) "positional" analysis of group belief. Tuomela assigns a belief to a group G as a function of the beliefs of certain "operative" members of G, who hold relevant positions in the group (1995, pp. 312-316).

[6] I am indebted here to Raphael Goldman.

[7] I am grateful to Don Fallis for valuable comments on an earlier draft of the paper.