

The epistemology of inclusiveness (or) Particular epistemic communities are always a mess^{*†}

P.D. Magnus

March 31, 2011

Comments are welcome.
e-mail: [pmagnus\(at\)fecundity.com](mailto:pmagnus@fecundity.com)
web: <http://www.fecundity.com/job>

Abstract

The epistemology of inclusiveness poses the problem of how the epistemic life of a community should be organized. In §1, I distinguish three approaches: what we might call the standpoints of first-person judgement, third-person judgement, and group procedure. I argue that each of these approaches is inevitably partial and incomplete. There are significant limits to what can be said about the epistemology of inclusiveness in the abstract. The details of particular communities ultimately make a big difference. I illustrate this point by considering some specific examples: on-line sources like blogs and Wikipedia (§2) and the elicitation of scientific opinion (§3).

1 Three ways to understand the topic

The *epistemology of inclusiveness*, so called, is not one of the perennial and central concerns of philosophers. One would be hard pressed to put together a collection that starts with ancient conceptions of the epistemology of inclusiveness and includes selections from throughout history. There are not clear battle lines between familiar camps. Indeed, the relationship between knowledge and inclusiveness is vexed. There are reasons to think that involving more people in knowledge making would contribute justification and reliability, with formal results like the Condorcet jury theorem and rhetoric about the *wisdom of crowds*. Yet there are also reasons to think the opposite, if additional people are biased or concerned with gaming the system. How should this tension be resolved? More generally, how should the epistemic life of communities be understood?

^{*}This paper was written for a special issue of *Synthese* on the epistemology of inclusiveness. It will not appear there.

[†]Thanks to John Milanese and Heather Douglas for comments on drafts of this paper.

The general problem can be understood in several different ways.

First, one might adopt a first-person point of view: How should an individual live in a community? As someone with my own beliefs and commitments, how ought I react to peers who have different beliefs and commitments? This approach has received philosophical attention in the last decade as the *epistemology of disagreement*, and there is now a considerable literature; see, e.g. [FW10].

Second, one might adopt a third-person point of view: Given what members of the community believe, how can the collective opinion be construed? This approach has received considerable development under the heading of *judgement aggregation*. Although some of this work has been done by philosophers (for a survey, see [Car11]), judgement aggregation has been developed in formal detail by economists.

Third, one might adopt a procedural point of view: Given the various opinions in a group, how should the group's official and announced opinion be decided? The literature on judgement aggregation is somewhat relevant here, too, but the problem is more concrete. Academic philosophers are certainly familiar with the problem in practice, even if they do not think of it as a philosophical problem. Philosophy departments and university committees find particular solutions as a matter of course, guided both by local university procedures and interpersonal negotiation.

Although each of these approaches poses and answers an interesting set of questions, each is only partial. There are aspects of the general problem — the epistemology of inclusiveness — which each ignores or distorts.

Consider first the epistemology of disagreement. Richard Feldman [Fel07] takes, as his primary example, questions of personal religious belief: when a person with one set of religious commitments gets to know an intelligent and reasonable person with different commitments, should this lead them to doubt the truth of their religion? Feldman suggests that such a situation should lead them to be agnostic. Since we take religious belief to be a matter of personal conscience, it makes sense to adopt the first-person standpoint for a question like this.

However, authors writing about disagreement do not confine themselves to such questions. For example, Adam Elga offers this example: “You and a friend are to judge the same contest, a race between Horse A and Horse B. Initially, you think that your friend is as good as you at judging such races. In other words, you think that in case of disagreement about the race, the two of you are equally likely to be mistaken” [Elg07, p. 486]. Importantly, the suggestion here is not just that you and a friend are forming beliefs about the race, as you might form beliefs about God or an afterlife, but instead that you are *judging* the contest. Approaching such a case from the first-person point of view would be irresponsible. In many sporting events, there is a hierarchy of judges or referees; the head judge asks particular judges what they saw and breaks ties to make a final call. If the separate judges revised their opinion of *what they saw* based on what other judges said, then the process would break down.

One might worry that I am putting too much weight on the details on Elga's

	P	Q	$(P\&Q)$
Alice	yeh	nay	nay
Bob	nay	yeh	nay
Charles	yeh	yeh	yeh
majority	yeh	yeh	nay

Table 1: There are three judges: Alice, Bob, and Charles. Each has consistent beliefs on the matters P , Q , and $(P\&Q)$. The *majority* rule yields the inconsistent combination of affirming P and Q but denying $(P\&Q)$.

example. One might revise the case, as Thomas Kelly [Kel10] does, so that you and your friend see the horse race but are not judges. The point, however, is not just that Elga’s example is poorly framed. Rather, the point is that the first-person point of view inevitably leaves out a great deal. What it means for an epistemic community to be functioning well cannot be teased out just by considering how members of the community ought to behave if they are just looking out for themselves.

Consider next the literature on judgement aggregation. Take a basic and typical problem, such as the *discursive dilemma*. Specify a community of judges; for any proposition, the *majority* rule yields the judgement endorsed by a majority of the judges. The rule has the nice features that it treats every judge equally and that it does not bias the conclusion toward one judgement or another. Yet it can lead to inconsistent collective judgements, even if all the judges considered individually have consistent beliefs. (See table 1.)

The discursive dilemma is just this problem: The *majority* rule can yield inconsistent results even if every individual judge has consistent beliefs.

A good deal of ink has been spilled precisely specifying various nice features that we might want a judgement aggregation procedure to have, proving that some desiderata are inconsistent, and proving that others are consistent. More sophisticated rules than *majority* have been developed, advocated, and disputed (see *infra*, §3). Although important, the insights are partial in two respects.

First, even if it can be proven that a set of desiderata cannot be satisfied in *all* cases, they may still be jointly satisfied in some instances. Although the *majority* rule can lead to contradiction, it will not do so in every case. As a practical matter, we might begin by trying out a simple rule (like *majority*) and add sophistication only if the actual community has judgements like those in table 1.¹ I argue below (in §3) that presumptive reliance on *majority* would be a mistake, but not because of some formal failing of the rule.

Second, judgement aggregation rules take the judges themselves as inputs. A particular judge may well, from the first-person perspective, reconsider their opinion in the face of a concrete instance of the discursive dilemma. As a pro-

¹The strategy of adding complications only as necessary can be applied generally to decision problems. For example, intransitive preferences wreck dominance reasoning. Yet one might presumptively employ dominance reasoning until one actually faces a case where there are intransitive preferences.

cedural matter, it might be valuable to have judges compare notes and debate one another; alternately, this may produce premature convergence and the appearance of confident agreement in cases where caution is justified.

The epistemic life of any community involves a variety of activities: Individuals are deciding for themselves what to think about particular questions. Individuals (who might or might not themselves be members of the community) are trying to figure out what the community knows. Small groups are attempting to settle collective opinion in a form that can be announced publicly.

The three approaches which I identified above each concentrate on one of these activities. Yet the activities, in an actual community, occur concurrently and with complicated interactions. So it is inevitable that each of the approaches should be abstract and partial.

Reckoning with the epistemology of inclusiveness in actual instances thus requires more than any of these approaches can provide. I want to suggest, moreover, that it depends on the nature of the questions being asked and the nature of the community doing the asking. The epistemology of inclusiveness, for real communities, will depend on their membership, norms, and resources. A general account is impossible.

I do not have an *a priori* argument for the kind of particularism that I am advocating. It is hard to know how an *a priori* argument could be mounted without the kind of general account which I insist is impossible. In the remainder of the paper, however, I illustrate the claim with a number of examples. I apologize that the connection between the examples and the general issue is more suggestive than decisive.

2 Blogs and Wikipedia

On-line information sources differ in many respects from their counterparts, and one major difference is in respect to their inclusiveness: Anyone with an internet connection can set up a blog or edit Wikipedia. Blogging citizen journalists are not directly responsible to editors and advertisers, in the way that traditional journalists are. Wikipedia, the encyclopedia that anyone can edit, is the product of a many anonymous incremental contributions. These are inclusive epistemology made real.

So it seems reasonable to ask whether news blogs are more or less reliable than traditional media, and whether Wikipedia articles are trustworthy sources of information. Yet these questions, although of *prima facie* interest, are ultimately not the right ones to ask.

Regarding blogs, Alvin Goldman [Gol08] argues that there is no question of whether traditional news or blogs are more reliable *tout court*. Rather, the epistemic value of any medium — the amount of knowledge it produces — depends on what kind of audience it has. “For the acquisition of knowledge to occur,” he argues, “it isn’t sufficient that there be a free press that publishes or broadcasts relevant truths. It is equally critical that members of the public receive and believe those truths” [Gol08, p. 113]. Only relative to particular participants

and psychological details about them is there any determinate answer to how reliable a social arrangement is.²

Regarding Wikipedia, I have argued elsewhere [Mag09] that the tremendous variability between topics and articles makes it impossible to say whether Wikipedia is trustworthy *tout court*. For a topic with an active communities of editors, Wikipedia articles can be remarkable good. For topics without a committed group of knowledgeable contributors or with so many vandals and cranks that editors are overwhelmed, the articles can be abysmal. Although there are some rules of thumb about which topics will be well behaved and which vicious, it ultimately depends on contingent facts about the particular people who contribute to a given entry.

3 What scientists know

When describing scientific knowledge, it is natural to talk about ‘what we know’. This is clearly not a claim made from the first-person point of view; it is not just I report that I know such-and-so or an exhortation that you should believe it, but instead a claim about the state of our science. Although the state-of-the-art may sometimes be reflected in committee reports or agency findings, it is not primarily a procedural matter either. So it may be tempting to understand this kind of talk as reflecting the third-person approach of judgement aggregation.

I will consider a specific approach and problems which arise for it, but I will argue that the problems are ultimately more general.

Stephan Hartmann, Gabriella Pigozzi, and Jan Sprenger [HPS10][HSng] develop a judgement aggregation rule specifically to escape discursive dilemma. Their procedure involves polling judges only regarding matters of independent evidence. For matters which are consequences of the evidence, the procedure derives consequences from the aggregated judgements. In the simple case given in table 1, for example, the procedure would affirm P and Q (because each is affirmed by a majority) and also $(P \& Q)$ (because it is a consequence of P and Q). Call this the *premise-majority* rule. Presuming that claims can all be divided into independent evidence claims and dependent consequences, then *premise-majority* will generate a consistent set of judgements.

There are several difficulties with *premise-majority*, as a way of aggregating expert scientific opinion.³

First, *premise-majority* inevitably produces some determinate answer. As

²Goldman’s approach represents a general outlook does not follow any of the three approaches that I identified in §1. He explains, “A major ambition of social epistemology (in the guise I have presented it) is to compare the knowledge consequences of alternative social practices, institutions, or mechanisms” [Gol08, p. 120]. Yet this *veristic social epistemology* is a high-level approach and so is not a competitor to any of the three approaches with which I began. Instead, it provides criteria for judging proposals within the third-person or procedural approaches.

³Since Hartmann et al. are thinking about the general problem of judgement aggregation, rather than the problem of expert elicitation, these are objections to the application of the rule rather than to the rule *as such*.

Brams et al. [BKZ98] show, it is possible for a combination of separate elections to result in an overall outcome that would not be affirmed by any of the voters. Moreover, a judge’s inconsistency will necessarily be between some belief about evidence and some belief about the consequences of the evidence — since the evidence claims are stipulated to be independent — but *premise-majority* does not query their beliefs about consequences at all. So it will generate a consistent set of judgements even if many judges are inconsistent. As such, *premise-majority* will generate determinate results even when the community is confused or fractured into competing camps. Yet in considering scientific opinion, we certainly only want to say that there is something ‘we know’ when it is a consensus or at least majority view.

Second, there may not be any clear division between which judgements are evidence and which are conclusions. As Fabrizio Cariani notes, *premise-majority* “requires us to isolate, for each issue, a distinguished set of logically independent premises” [Car11, p. 28]. He constructs a case involving three separate, contentious claims and an agreed upon constraint, such that any two of the three claims logically determines the third. It would be arbitrary to treat two of the claims as evidence (and so suitable for polling) and the third as a consequence (and so fixed by inference). The *premise-majority* simply is not applicable in cases where the line between premises and conclusions is so fluid. Cariani concludes, “Different specific aggregation problems may call for different aggregation rules” [Car11, p. 29]. The problem is especially acute for scientific judgement, because inference can occur at different levels. Individual measurements like ‘35° at 1:07 AM’ are not the sort of thing that would appear in a scientific publication; individual data points are unrepeatable and not something about which you would query the whole community. And scientists may take things like the constancy of the speed of light to be evidence for a theory; the evidence here is itself an inference from experiments and observations.⁴ So it is unclear what we would poll scientists about if we applied *premise-majority*.

Third, *premise-majority* is constructed for cases where the conclusion is a deductive consequence of the premises. In science, this is almost never the case.⁵ Scientific inference is ampliative, and there is uncertainty not only about which evidence statements to accept but also about which inferences ought to be made on their basis. One might avoid this difficulty by including inferential relations among the evidential judgements. To take a schematic case, judges could be asked about R and $(R \rightarrow S)$; if the majority affirms both, then *premise-majority* yields an affirmative judgement for S . The problem with this suggestion is that it presumes that scientists can say, independently of everything else, whether the inference from R to S is appropriate. There is a deep reason to deny that they can: Scientific inference must be informed by *more* than just the particular

⁴There are different labels for these different levels. Trevor Pinch [Pin85] calls them observations of differing *externality*. James Bogen and James Woodward [BW88] distinguish *data* from *phenomena*.

⁵I say ‘almost’ because sufficiently strong background commitments can transform an ampliative inference into a deduction from phenomena. Of course, we accept inductive risk when we adopt the background commitments; cf. [Mag08].

evidence — the appropriate scientific conclusion depends (in some cases) on the risks and values involved. I argue that this difficulty is a problem for more than just Hartmann et al.’s specific proposal.

First, we need to be clearer on the way in which inference can be entangled with values and risk; I argue for this connection in the next section. In the subsequent section, I return to it as a problem for *premise-majority* and judgement aggregation.

3.1 The James-Rudner-Douglas thesis

Here is a quick argument for the entanglement of judgement and values: There is a tension between different epistemic duties. The appropriate balance between these duties is a matter of value commitments rather than a matter of transcendent rationality. So making a judgement of fact necessarily depends on value commitments.

The argument goes back at least to William James, who puts the point this way: “*We must know the truth; and we must avoid error* — these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws” [Jam48, p. 99]. Although James has in mind personal matters of conscience (such as religious belief), Richard Rudner makes a similar argument for scientific judgement. Rudner argues that

the scientist must make the decision that the evidence is *sufficiently* strong... to warrant the acceptance of the hypothesis. Obviously our decision regarding the evidence and respecting how strong is “strong enough”, is going to be a function of the *importance*, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis. [Rud53, p. 2]

There is not only a tension between finding truth and avoiding error, but also between making risking one kind of error and risking another. Any particular test involves a trade-off between making the standards too permissive (and so mistakenly giving a positive answer) or making them too strict (and so mistakenly giving a negative answer). The former mistake is a *false positive* or *type I* error; the latter a *false negative* or *type II* error. There is an inevitable tradeoff between the risk of each mistake, and so there is a point at which the only way to reduce the risk of *both* is to collect more evidence and perform more tests. Yet the decision to do so is obviously a practical as well as an epistemic decision. In any case, it leaves the realm of expert elicitation — it would mean having different science, rather than discerning the best answer our present science has to the question. In the end, values come into play. Heather Douglas concludes, “Within the parameters of available resources and methods, some choices must be made, and that choice should weigh the costs of false positives versus false negatives. Weighing these costs legitimately involves social, ethical, and cognitive values” [Dou09, p. 104].

Plotting a curve through these 19th, 20th, and 21st-century formulations, call this the *James-Rudner-Douglas* or *JRD thesis*: Anytime a scientist announces a judgement of fact, they are making a tradeoff between the risk of different kinds of error. This balancing act depends on the costs of each kind of error, so scientific judgement involves assessments of the value of different outcomes.

A standard objection to the thesis is that responsible scientists should not be making categorical judgements. They should never simply announce “ P ” (the objection says) but instead should say things like “The available evidence justifies $x\%$ confidence in P .” This response fails to undercut the thesis, because procedures for assigning confidence levels also involve a balance between different kinds of risk. This is clearest if the confidence is given as an interval, like $x \pm e\%$. Error can be avoided, at the cost of precision, by making e very large. Yet a tremendous interval, although safe, is tantamount to no answer at all.⁶

Another step is required to connect the JRD thesis to inclusiveness. After all, James does not apply it to empirical scientific matters. He is concerned with religious and personal matters, and he concludes merely that we should “respect one another’s mental freedom” [Jam48, p. 109]. Rudner, who does apply the thesis to empirical judgements, nevertheless hopes that the requisite values might themselves be objective. What we need, he concludes, is “a science of ethics” [Rud53, p. 6]. Yet such a science is not, as Rudner calls it, a “task of stupendous magnitude” [Rud53, p. 6]. Instead, it is a fool’s errand.

A great virtue of Douglas’ discussion is that she considers the concrete question of how to determine the importance of the relevant dangers. She argues for an *analytic-deliberative process* which includes both scientists and stakeholders [Dou09, ch. 8]. Such a process is only required when the scientific question has a bearing on public policy. There are further conditions which must obtain in order for such processes to be successful. For one, “policymakers [must be] fully committed to taking seriously the public input and advice they receive and to be guided by the results of such deliberation” [Dou09, p. 166]. For another, the public must be “engaged and manageable in size, so that stakeholders can be identified and involved” [Dou09, p. 166]. Where there are too many stakeholders and scientists for direct interaction, there can still be vigorous public examination of the values involved. Rather than pretending that there is any all-purpose procedure, Douglas calls for “experiment with social mechanisms to achieve a robust dialog and potential consensus about values” [Dou09, p. 169]. Where consensus is impossible, we can still try to elucidate and narrow the range of options. Douglas’ approach is both a matter of policy (trying to increase trust in science, rather than alienating policymakers and stakeholders) and a matter

⁶A substantially more subtle reply to the standard objection is given by Eric Winsberg and Justin Biddle [BW10]. In the case of climate modeling, they argue that scientists’ estimates both of particular quantities and of confidence intervals depend on the histories of their models — e.g., if they model ocean dynamics and then add a module for ice formation rather than vice-versa. The history of a model reflects decisions about what was considered to be important enough to model first, and so it incorporates prior value judgements.

of normative ethics (claiming that stakeholders' values are ones that scientists should take into consideration).

In cases where Douglas' concerns are salient, procedures for inclusiveness must depend on facts about the actual communities of scientists, policymakers, and stakeholders. Nevertheless, her concerns will not be salient in all cases, because some science is far removed from questions of policy. So the significance of the JRD thesis depends on the question being asked and on broader questions of political organization.

3.2 JRD and *premise-majority*

I argued above that the *premise-majority* rule was inapplicable in many scientific contexts because it only worked for cases of deductive consequence. Formally, this worry could be resolved by asking scientists about which inferences would be justified; we poll them about claims like $(R \rightarrow S)$, where we are polling them about R but not about S .

However, the problem raised by the JRD thesis is not solved by this formal trick. Where the judgement has consequences, the inference itself is an action under uncertainty. So the appropriate inference depends on the values at stake. For example, questions of climate science may not be entirely separable from questions of the value of the environment. If we merely poll scientists, then we will be accepting whatever judgements accord with their unstated values.

We instead want the procedure to reflect the *right* values, which in a democratic society means including communities effected by the science. Importantly, this does not mean that stakeholders get to decide matters of fact themselves. Rather, they help determine how the risks involved in reaching a judgement should be weighed. Nor does it mean that politicized scientific questions should be answered by political means. Climate scientists can confidently identify general trends and connections, even allowing for disagreement about the values involved. What they cannot do is provide an account that is value-neutral in all its precise details.⁷

This is not merely a problem for *premise-majority*, but for any judgement aggregation rule that treats judges merely as separate inputs to an algorithm. Where there are important values at stake that scientists are not taking into account, or where the value commitments of scientists are different than those of stakeholders, the present judgements of individual scientists cannot just be taken as given. What Douglas calls an analytic-deliberative process is needed, and the appropriate analytic-deliberative process will depend on facts about the communities involved. In terms of the approaches that I identified in §1, this means that the third-person approach shades into the procedural approach for cases like these. And the procedural approach offers no uniform answers.

⁷Douglas [Dou09, esp. ch. 6] provides an excellent discussion of how (what I have called) the JRD thesis is compatible with objectivity.

4 Conclusion

I began by noting that the *epistemology of inclusiveness* is not a coherent sub-discipline within philosophy. I surveyed three approaches — the first-person approach of deciding individual belief, the third-person approach of discerning the community’s opinion, and the procedural approach of settling on an explicit collective judgement. Each of these approaches is legitimate; they ask and answer important, interesting questions. Nonetheless, each is inescapably partial.

The epistemology of inclusiveness involves disparate issues which cannot be addressed in the abstract. Answers depend on the particularities of the questions under consideration and the communities doing the knowing. Perhaps this means that the epistemology of inclusiveness could never be a coherent sub-discipline. This should not turn us back, however. There are important questions here — even if they are a motley — and we should not be afraid of digging into the details.

References

- [BKZ98] Steven J. Brams, D. Marc Kilgour, and William S. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.
- [BW88] James Bogen and James Woodward. Saving the phenomena. *Philosophy of Science*, 97(3):303–352, July 1988.
- [BW10] Justin Biddle and Eric Winsberg. Value judgements and the estimation of uncertainty in climate modelling. In P.D. Magnus and Jacob Busch, editors, *New Waves in Philosophy of Science*, pages 172–197. Palgrave MacMillan, Basingstoke, Hampshire, 2010.
- [Car11] Fabrizio Cariani. Judgment aggregation. *Philosophy Compass*, 6(1):22–32, 2011. doi:10.1111/j.1747-9991.2010.00366.x.
- [Dou09] Heather E. Douglas. *Science, Policy, and the Value-free Ideal*. University of Pittsburgh Press, 2009.
- [Elg07] Adam Elga. Reflection and disagreement. *Noûs*, 41(3):478–502, 2007.
- [Fel07] Richard Feldman. Reasonable religious disagreements. In Louise M. Antony, editor, *Philosophers without Gods*, pages 194–214. Oxford University Press, 2007.
- [FW10] Richard Feldman and Ted A. Warfield, editors. *Disagreement*. Cambridge University Press, 2010.
- [Gol08] Alvin I. Goldman. The social epistemology of blogging. In Jeroen van den Hoven and John Weckert, editors, *Information Technology*

- and Moral Philosophy*, pages 111–122. Cambridge University Press, 2008.
- [HPS10] Stephan Hartmann, Gabriella Pigozzi, and Jan Sprenger. Reliable methods of judgment aggregation. *Journal for Logic and Computation*, 20:603–617, 2010.
- [HSng] Stephan Hartmann and Jan Sprenger. Judgment aggregation and the problem of tracking the truth. *Synthese*, forthcoming.
- [Jam48] William James. The will to believe. In Alburey Castell, editor, *Essays in Pragmatism*, pages 88–109. Hafner Publishing Co., New York, 1948. Originally published June, 1896.
- [Kel10] Thomas Kelly. Peer disagreement and higher-order evidence. In Feldman and Warfield [FW10], pages 111–174.
- [Mag08] P.D. Magnus. Demonstrative induction and the skeleton of inference. *International Studies in the Philosophy of Science*, 22(3):303–315, October 2008. doi:10.1080/02698590802567373.
- [Mag09] P.D. Magnus. On trusting Wikipedia. *Episteme*, 6(1):74–90, 2009. doi:10.3366/E1742360008000555.
- [Pin85] Trevor Pinch. Towards an analysis of scientific observation: The externality and evidential significance of observational reports in physics. *Social Studies of Science*, 15:3–36, 1985.
- [Rud53] Richard Rudner. The scientist qua scientist makes value judgments. *Philosophy of Science*, 20(1):1–6, January 1953.