BerkeleyLaw
UNIVERSITY OF CALIFORNIA
Kadish Center for Morality, Law & Public Affairs

General Aspects of Law

GALA

Dean's Seminar Room (215B Boalt Hall)

Thursday

November 15, 2012

4:00 PM – 6:00 PM

"When the Experts are Uncertain: Scientific Knowledge and the Ethics of Democratic Judgment"

By **Melissa Lane**

Professor of Politics, Princeton University

2012-13 Fellow of the Center for Advanced Study in the Behavioral Sciences, Stanford

NOTE

The GALA speaker will offer some brief introductory background remarks about the paper and the remainder of the time will be devoted to discussion. Participants are expected to read the paper in advance. Papers are distributed upon request to those on our mailing list or electronically when available at the GALA website: <http://www.law.berkeley.edu/9264.htm>. Others may obtain copies from Mr. Daniel Zussman at dzussman@law.berkeley.edu or 510.666.3781. Copies can be made available in an alternate format upon request. GALA events are wheelchair accessible. For any disability-related accommodations advance notice is requested.

GALA Seminar on 15 November 2012
Kadish Center for Morality, Law & Public Affairs, UC Berkeley
 mslane@princeton.edu

When the experts are uncertain: scientific knowledge and the ethics of democratic judgment

Melissa Lane
Professor of Politics, Princeton University
2012-13 Fellow of the Center for Advanced Study in the Behavioral Sciences, Stanford

Introduction

How, and how well, can ordinary citizens in a democracy evaluate the claims of scientific experts?[1] This problem often arises in multiple contexts, as when we choose our doctors or hear their diagnoses, when we evaluate evidence on juries, or when we assess scientific claims that bear on our votes in elections or our replies to opinion polls. In this paper, I focus on the ways that citizens form judgments about questions to which scientific expert testimony is relevant, judgments on which they may then draw in a range of political roles. In fact, a detailed answer to my opening question will have to be given case by case, depending on specific institutions, historical circumstances, and political contexts. Yet some scholars have attempted to give general answers to the question, answers which are sharply opposed. Given this conflict, it is necessary to reconsider the broad contours of the problem and prospects for a general approach to answering it. By showing what good democratic judgment of expertise requires and involves in general terms, I argue, we can better understand the challenges presented by particular contexts and develop ways to enable good judgment in spite of them.

Before giving an overview of the paper, I should explain why I use the term 'judgment.' 'Judgment' is the epistemic concept that the ancient Greeks used to describe responses to communicative acts of persuasion. Aristotle cast all auditors as judges: 'we

may say, without qualification, that anyone is your judge whom you have to persuade' (Aristotle, *Rhetoric*, II 18, 1391b; cf. I 2, 1356a; in Barnes 1984). From the assembly to the law courts to the prize committees for dramatic performances, Athenian institutions were designed precisely to allow non-experts to make up their minds about claims or products offered by those claiming superior knowledge or insight in these domains. Non-experts were required to evaluate proposals to fight wars made by those who had been elected as generals, to assess magistrates whose performance in office was subject to scrutiny, and to appraise tragedies and comedies written by playwrights selected to present their work in civic and religious festivals. In using the concept of lay 'judgment' I draw on these features of its Athenian deployment to indicate that the standard we seek is not one of epistemic parity between speakers and auditors. Rather, it requires a distinct and appropriate level of epistemic attainment on the part of the auditors that enables them to assess and engage with, while not duplicating, the claimed expertise of the speakers.[2]

But how are we to determine what constitutes good judgment? In the first part of the paper, I argue that the general approach to answering this question has been unduly limited in the way it frames the question: namely, it focuses on how lay people can *identify* experts. I survey two opposed positions that have been staked out in this limited program. Both positions utilize what has been called a 'novice – 2 expert' frame (Goldman 2001) in which the key issue is choosing between two (or sometimes more) rival putative experts. One camp is *skeptical*, arguing that the choice reduces to the evaluation of credentials but that such evaluation cannot be done by lay people incapable of understanding expertise in the relevant domain. The other camp is *optimistic*, contending that citizens are capable both of evaluating the credentials of experts and

assessing experts' arguments without having full understanding of their expertise. Not only is each camp is internally unstable, an observation I develop below, but also, both consider the best prospects for citizen judgments to be second-order judgments only, restricted to observing features of the experts' claims and standing rather than to engaging directly with the substance of their arguments. Such unduly narrow limits exclude cases where citizens not only must choose among rival experts but must also evaluate and determine the implications of acting in light of the claims that the experts make. This broader problem of evaluating and determining implications for action generates an additional burden of judgment: citizens must also evaluate the multiple levels of uncertainty that attend most expert claims.

To address this broader and more fundamental conception of the problem of lay judgment of expertise, I develop my alternative account in three steps (Parts II to IV). Part II argues that lay judgment is best understood (and best enabled) when it is conceived not as the result of experts dumping information on the laity, but rather as an embedded practice in an iterative and normative relationship of communication. Part III develops an account of the ethical norms attaching to both experts and laity in such an ongoing communicative relationship. Here I focus on the problem of communicating uncertainty in particular, showing how this demands, ethically, the cultivation of epistemic norms, habits and strategies of relating to self and other, again by both parties. To illustrate each of these arguments, I conclude in Part IV with a brief discussion of the dramatic conviction of the members of an Italian scientific risk commission in L'Aquila as a relevant case. Collectively, Parts II – IV construct and apply the ethics of democratic judgment to which my title refers.

I. Opposite Camps: Optimistic and Skeptical Approaches to Non-Expert Judgment

Why should lay people defer to experts at all? A helpful framework is often drawn from the approach to authority developed by Joseph Raz, most fully in *The Morality of Freedom* (1986). Raz begins by treating theoretical and practical authorities as sharing a certain structure in which one's reason for deferring to an authority is that such deference better enables one to comply with applicable duties (and other reasons for action). Raz focuses on developing the implications of this approach for practical authorities. Others, including Linda Zagzebski (2012), have sought to develop a parallel approach for theoretical authorities ('scientific' in my parlance in this paper), on which one's reason to defer to an authority is that it better enables one to pursue and attain knowledge, something which one has a natural concern to do.[3]

One strength of a Razian approach is that, by developing an account of authority that is consonant with a particular understanding of autonomy, it sidesteps a particular skeptical challenge, the worry about epistemic dependence that motivates a classic paper by John Hardwig (1985). Hardwig had argued that we exhibit an objectionable lack of 'epistemic autonomy' when we rely on theoretical or practical experts. By constructing autonomy differently, namely as the capacity to be part-author of our lives, Raz and his followers effectively argue that this is a role we play best precisely when we rely on others where their knowledge is likely to be superior to ours.

While Raz's rejoinder avoids Hardwig's skeptical challenge, a more troubling difficulty arises for the Razian approach. It fails to attend to how experts – either theoretical or practical – can actually be identified in practice. Raz suggests that, in normal circumstances, so long as I can judge an authority to be a genuine authority, I

have reason to obey it, and it is normally justified in exerting its authority over me. But how can I make that identificatory judgment? As I have argued previously (Lane 1998-99), Raz's major work on the subject, *The Morality of Freedom* (1986), does not acknowledge this as a problem: there, his focus is always on some one authority and a layperson, rather than on competition between would-be authorities. In a later summary article in the *Minnesota Law Review*, Raz briefly remarks that 'to fulfill its function, the legitimacy of an authority must be knowable to its subjects,' '[s]ince the point is to improve conformity with reason' (2005-06: 1025). Even here, however, Raz assumes that recognizing this legitimacy will be unproblematic in the case of practical authorities. He treats it as a common lay task: 'We engage in such assessments every day of the week' (p.1025).

This insouciance may often be justified in the practical case and, in particular, the political case, in which I am normally confronted with a pre-existing and massively superior authority such as the state. But even in political contexts, two would-be authorities may compete for my allegiance, as after a coup or an attempted secession. In such cases, it is not clear that Raz has much guidance to give. That lack of guidance is even more acute in epistemic contexts of *theoretical* authority, in which citizens may be confronted with rival groups of contending experts. The recent extension of a Razian approach to the epistemic domain by Linda Zagzebski (2012), impressive as it is, has little to say about the problem of identifying experts in such cases.

There is, however, an extensive literature on how lay people can identify experts, one which has come to dominate the discussion of the relationship between lay people and expert knowledge altogether. Alvin Goldman (1999 and 2001), for example, has

developed a robust approach to offer identifying marks of expertise. Interestingly, Goldman's account has striking affinities with the approach taken by Socrates in the Platonic dialogues, texts that are deeply preoccupied with how ordinary people can possibly identify genuine experts and forms of expertise, and how such specialized expertise relates to politics.[4] Consider the following list of marks of expertise that have been offered by Socrates, Goldman, or both. (I draw on my own analysis of Platonic texts but owe the specific wording and enumeration of items (i) – (v) in the Socratic list to LaBarge 1997; see also Gentzler 1995.) I have ordered the criteria to bring out the parallels that exist in four out of the six cases: *explanation offered; agreement among experts; mutual recognition of experts*; and *success or track record*. Socrates adds one criterion (listed here first) -- *the expert's ability to teach the expertise* – that has no parallel in Goldman, while Goldman adds one criterion (listed here sixth) -- *evidence of biases* -- which has no parallel in Socrates.

Table 1

Socrates (drawing in part on LaBarge 1997)	**Goldman** (drawn from 2001: 93)
(i) Teachability: expert can teach student the expertise	[no analogue]
(ii) Explanation offered by the expert and **Elenctic testing**	**Argument offered by the expert**

(iii) **Agreement among experts**	**Agreement with additional putative experts**
(iv) **Mutual recognition of experts by one another**	**Appraisal by meta-experts** [or other experts] **of their expertise or credentials**
(v) **Success (track record)**	**Track record**
(vi) [no analogue]	**Evidence of biases –** **tested by aspects of research procedure** (funding source, etc)

To the Socratic eye, the criteria of *explanation*, *agreement*, *mutual recognition*, and *success* are at once necessary and insufficient on both counts for the layperson to recognize a true expert (though they might be sufficient for one expert to recognize another). For how can the layperson be sure of what counts as a good explanation, or success – and how can she know whether agreement and mutual recognition constitute real evidence of expertise or simply a successful con game? Teachability seems to be the only failsafe route. Yet teachability works only by converting the layperson herself into an expert and thereby dissolving the original terms of the problem.

Nevertheless, the Socratic versions of the criteria do play a useful role. In particular, the idea that claims offered by an expert can be tested as to whether they lead to contradiction of other claims that the same expert wishes to accept – the Socratic method of the elenchus – allows the non-expert a systematic route to discrediting would-be experts. If a putative expert cannot put forward an account of her expertise (in Socratic terms, a definition) which she wishes, and is able, to go on defending after its contradiction of other claims she accepts has been manifested, then it is a good bet that she is not an expert after all. This Socratic approach to enabling laypeople to test the first-order claims of experts is one that I will develop further below.

Goldman is less troubled than Socrates by the second-order and potentially circular nature of the *agreement* and *mutual recognition* criteria. In the modern world in which expertise is far more bureaucratically elaborated than it was in ancient Athens, the rationality of relying on professional recognition and accreditation seems more secure. But the con game problem cannot simply be met by general appeals to the inevitability of relying on testimony: we may not be able to escape our reliance on testimony, but that does not underwrite all testimony as inviolate, especially testimony offered by relatively closed professions with particular sectoral interests and concerns.. And when we consider the value of *explanation* and *track record* as criteria meant to assist the layperson in judging experts, we may feel that Socratic worries are on point. One needs to be sure that the goals and standards for success of a track record are appropriate and valuable (Estlund 1993, 2008); in other words, one needs at least to understand the explanation in order to judge it. The criteria here seem to involve more demanding epistemic capacities than first appeared.

So far I have surveyed the Razian approach and attempted to supplement its Achilles' heel – the lack of attention to identifying marks of expertise – by evincing the closely parallel approaches to marks of expertise offered by Socrates and Goldman. I will now present two opposed perspectives on the success of such a project. At one extreme is the view that such identification of experts is relatively robust and straightforward; at the other is the view that it is so difficult as to be effectively impossible. I will consider Elizabeth Anderson (2011) and Scott Brewer (1997-98) as exemplars of each respective pole. My analysis will suggest that Anderson effectively makes the 'novice-2 experts' frame into a 'novice-crackpot/expert' frame – treating the choice between 2 experts as actually a choice between an expert and a fraud -- whereas Brewer effectively makes it into a 'novice-hired gun' frame in which either or both experts might be frauds, but the audience is forced to choose between them. Anderson does not discuss the ways in which lay people might need or be able to engage with the substantive claims, including the uncertainties, that the experts make; Brewer denies that they will be able to do so. (Another way of putting this is that Anderson lacks interest in cases where the competing experts are all genuine, while Brewer's skepticism means that he treats such cases as on a par with expert-fraud cases.) Both miss out the relational nature of expert-laity communication which may make it possible (contra Brewer) for lay people to evaluate substantive expert claims.

Since Brewer's skeptical argument, if sound, would block both the possibility of epistemically warranted choice between rival experts *and* the possibility of any broader lay evaluation of expert claims, it is useful to consider it first. Brewer's discussion focuses on the non-expert figures of both jury and judge in American courts of law.[5]

This forum helps to explain his focus on the 'novice-2 experts' problem. The courtroom is a context in which novices may be inclined to choose between complete packages of explanation offered by Expert A and Expert B rather than evaluating the items of the packages individually.

Brewer begins by downplaying the possibility that a putative expert's explanatory reasoning can be shown to be decisively unsound. In his view, this will happen only in rare cases of full-fledged rational incoherence or irrationality (pp.1617-18). Thus, his standard case is one in which two dueling experts are both so plausible in their arguments and explanations that it is impossible for the non-expert to distinguish between them on those grounds. That leaves the choice between experts to be made on the basis of 'credentials, reputation, and demeanor.' With these possible criteria in view, Brewer rejects *demeanor* as an epistemically sound basis for choosing an expert due to the 'market in demeanor' which rewards training in demeanor divorced from any actual knowledge. Then he, in effect, collapses *reputation* into *credentials*. His core claim is therefore the following: assessing credentials requires 'a *reasoning process*' (p.1538, emphasis original) which laypeople are incapable of properly deploying. To assess a credential one needs a full understanding of the expertise it certifies, but to gain that understanding one would need a credential, but to obtain and assess that credential, and so on, in an infinite regress.[6] Brewer sums up:

> [T]he nonexpert's lack of epistemic competence threatens to deprive her of precisely the kind of understanding she would need to be able to confirm or disconfirm a hypothesis about credentials and their capacity accurately to identify which experts are capable of producing KJB [knowledge and justified belief, treated indiscriminately: p.1601] and which are not (p.1669).

Brewer concludes that there is 'compelling reason to doubt' that non-expert judges can make a non-arbitrary assessment of the epistemic issues at stake in expert testimony. Therefore, as a standard for legal legitimacy, 'intellectual due process' requires an alternative procedure:

> The only solution (actually, it is a family of solutions) I see requires that one and the *same* legal decisionmaker wear two hats, the hat of epistemic competence and the hat of practical legitimacy. That is, whether it is a scientifically trained judge or juror or agency administrator, the same person who has legal authority must also have epistemic competence in relevant scientific disciplines' (p.1681, emphasis original).

This solution is, however, doubly unstable. First, a scientifically competent judge who is in a position to evaluate scientific credentials properly would thereby be in a position to enter into the scientific reasoning and explanations as well: she would not limit herself merely to assessing credentials. So the solution would go beyond the narrow limits Brewer sets. Second, we must ask, precisely what level and specificity of scientific training would be required to wear the 'hat of epistemic competence'? Brewer concedes that a PhD would not be required (he seems to be talking about something like a university degree in the sciences).[7] Furthermore, it is implausible that the degrees of such judges could be exactly matched to the expertise at stake in any given trial (not least because multiple disciplines may be relevant to a single case). So we must ask: how precisely would an undergraduate degree in some scientific discipline equip someone to assess the arguments of rival experts, let us say, in another discipline? Either some degree of scientific training provides a level of episto-social literacy in the nature of the academic world that fortifies one in assessing credentials, or it affords a platform on the

basis of which one can learn about new disciplines and come to understand new arguments. Both alternatives open the door to a spectrum of possibilities which cannot be arbitrarily restricted to a binary expert / lay divide. Thus, Brewer's purported skeptical refutation of lay assessment of scientific expertise actually makes room for a range of possible scenarios – including the provision of certain levels of education and literacy – in which non-experts in a given discipline may be able to make non-arbitrary judgments of the claims as well as the status of experts. Drawing a new line between experts and non-experts, as Brewer attempts to do by putting those scientifically trained as undergraduates on the 'expert' side for courtroom purposes, actually raises new possibilities as to how those two groups might be related, might interact, and might even be defined.

We can turn once again to Aristotle to develop this point. In *Politics* book III, chapter 11 (translation in Reeve 1998), Aristotle's concern with expertise is brought out in his choice of analogy, defending the popular ability to judge experts such as doctors. (I argue in Lane, forthcoming, that the institutional manifestation of 'judging' in this chapter is electing and inspecting officials.) He gives a stark statement of the objection that experts such as doctors can be inspected or audited only by their expert peers: 'just as a doctor should be inspected by doctors, so others should also be inspected by their peers' (1282a1-3). To this he replies by distinguishing levels of education in most crafts including medicine. One might have a general education in a craft; be an ordinary practitioner of it; or be a master craftsman. All these can judge (it is implied) even the master craftsman.

Now what is striking about this passage is that Aristotle says that all three kinds of people are called ‘doctor,’ and that he seems to take the status of a ‘general education in a craft’ to apply very broadly (insofar as this conclusion is meant to defend the overall thesis endorsing participation of the general multitude in electing and inspecting officials) (1282a3-5). This can’t be a specialized education. It must be more like the involvement of ordinary people, say, in medicating their children at home and so in sharing in medical practice and concerns. Aristotle’s reply further erodes any sharp boundary between popular and expert knowledge, or what can be more properly considered popular judgment and expert knowledge. They are certainly distinct, but they fall on a continuum, and there will be certain habits of mind shared between them.

With such possibilities in mind, we may briefly turn to the opposite pole on the question of lay identification of experts: an extreme exemplified by Elizabeth Anderson (2011). Anderson operates within the broad framework of reliance on testimony by others, setting up the problem as one in which ordinary people need ‘criteria of trustworthiness and consensus for scientific testifiers’ (p.145). Her approach is to agree with Brewer’s official position (different, as I have argued, from the upshot of his view) that laypeople cannot enter into the substance of expert arguments and explanations. Nevertheless, she views credentials as only one of three routes for them to make reliable decisions about which experts to trust (and is untroubled by the thought of any regress here). Alongside credentials are two additional kinds of ‘second-order judgments’ (p.145) about whom to believe, relieving ordinary citizens of the need (of which she assumes most are not capable, p.145) of engaging in first-order assessments of the merits of scientific claims. These criteria are *honesty*, tested in practice by evidence of external

conflict of interest and also by misleading statements; and *epistemic responsibility*, which is tested in practice by the willingness to accept external peer review and to abide by evident canons of dialogic rationality in argument (not simply repeating a claim without acknowledging a prior objection, for example). Anderson herself insists that all these three routes – *credentials*, *honesty*, and *epistemic responsibility* should be understood as wholly 'second-order' (her term as well as mine). Laity need not understand the arguments of scientists in order to spot their conformity (or lack thereof) with these second-order canons.

As I argue below, Anderson's conceptions of honesty and epistemic responsibility are powerful and can be further expanded. Nevertheless, her account suffers from a more general flaw. For whereas Brewer stages his story as a duel between equally plausible hired guns, Anderson frames hers as a lopsided battle between a dominant group of credible scientists and a few 'crackpots' (pp.146-7). This seems part of what leads her to be untroubled by the prospect of rival credentials, since she envisages respectability – and widespread agreement among scientists -- as the property of only one side.[8] While this may be an accurate portrayal of the dynamics of climate change science and its deniers, it is not adequate as a general portrayal of the problem of lay judgment of scientific expertise. Not all cases in which citizens must judge expertise fit either Brewer's or Anderson's model. On the one hand, there may be a more even distribution of epistemically respectable views, and also a wider spectrum of them. On the other hand, citizens in many real life scenarios do not need merely or primarily to choose between rival experts or to root out a few obvious frauds. Rather, they need to decide how to assess and use what the spectrum of scientific expertise on a given issue reveals, a

broader problem which is likely to be accompanied by the need to cope with multiple levels of uncertainty. The presence and extent of uncertainty points to the need for a more intensive "first-order" engagement with scientific claims than Anderson's confident solution to a problem she conceives as wholly "second-order" allows.

We may draw three interim conclusions, First, the dominant framing of the problem of lay judgment of scientific expertise is skewed by its focus on the second-order problem of *identifying* experts, neglecting (or denying) the possibility and need for citizens to engage *directly* with experts' arguments and claims. Second, the internal undoing of a skeptical denial of citizens' capacity for first-order engagement points in a helpful direction, suggesting that a continuum of forms of judgment and knowledge may be cultivated. Third, the internal undoing of an overconfident anti-skepticism points in the same broad direction: any solution to the general problem cannot limit itself to second-order identifying marks, but must rather develop a richer conception of epistemic responsibility and honesty which is rooted in and makes possible an assessment of first-order scientific claims. These points can be drawn out from Table 2, which summarizes the views we have canvassed so far. (The entries for Anderson and Brewer in row (ii) are italicized to show that they are not true analogues for the Socratic or Goldman criteria there: they are second-order substitutes for expert explanation, not expert explanation itself.)

Table 2

Socrates (drawing on LaBarge 1997)	**Goldman (**2001**:** 93)	**Anderson (**2011: 146-7)	**Brewer (**1998)
(i) Teachability: expert can teach student the expertise	[no analogue]	[no analogue]	[no analogue]
(ii) Explanation offered by the expert with **test of expert claims by elenchus**	**Argument offered by the expert**	***Second-order test of Epistemic Responsibility:*** *willingness to accept external peer review and not exhibit extreme dialogic irrationality*	***Second-order: demeanor*** *[not explanations]* ***Or: second-order full-fledged rational incoherence*** *[rare; Anderson dialogic irrationality]*
(iii) Agreement among experts	**Agreement with additional putative experts**		**Reputation** [but not developed]
(iv) Mutual recognition of experts by one another	**Appraisal by meta-experts of their expertise or credentials**	**Credentials**	**Credentials** [but requires grasp of explanations: regress]
(v) Success (track record)	**Track record**	[no analogue but probably accepts]	[rejects: regress]
(vi) [no analogue]	**Evidence of biases tested by aspects of research procedure** (funding source, etc)	**Honesty**: evidence of external conflict of interest, misleading	[no analogue]

II. Modeling Scientific – Lay Communication

We are now in a position to recognize that, despite representing opposed camps on the question of whether laypeople can identify experts and expose frauds, Anderson and Brewer actually share an important unspoken presumption. Both envisage first-order scientific reasoning as effectively insulated from laypeople, who are reduced to being mere observers of external indicia about the scientists in order to decide whose package of arguments to accept, though Anderson thinks these indicia extend to eavesdropping, as it were, on debates among the scientists themselves to test whether dialogical rationality or irrationality is being exhibited. Even in Anderson's turn toward communicative norms, she identifies these norms as applying within settings of scientists debating among themselves, rather than applying directly to the communicative act between scientists and laypeople.

Yet lay judgment of scientific expertise depends precisely on construing and assessing that expertise as communicated to a wider public. Such communication is misunderstood if it is conceived as simply the public overhearing internal scientific debates. Rather, as science studies and the social studies of the science movement have emphasized, communication can only be understood when the norms, attitudes, and expectations of both parties are taken into account. Diverse publics do not merely accept or reject science, nor do they simply possess or lack scientific knowledge (Brossard and Lewenstein 2010, Russell 2010). Rather, they filter scientific communications through the lens of distinct frames, purposes, and preexisting knowledge and attitudes. (This may even include forms of 'lay expertise,' a contradiction in terms for the problem as both Anderson and Brewer construe it.)

The relational aspects of such sociological studies can be explicated by appeal to the 'agency model' of communication developed by philosophers Neil Manson and Onora O'Neill in their work on bioethics, *Rethinking Informed Consent* (2007). Criticizing the prevalent 'conduit' and 'container' metaphors that treat informing as a one-way dumping of undifferentiated data, they explain that such metaphors obscure the fact that 'communicating and informing are types of *action* and *interaction*, so depend on a normative framework against which such action succeeds or fails' (p. 27, emphasis original). They spell out the implications of this view as follows:

> Acts of informing (and communication more generally) *only* succeed within a rich practical and normative framework in which speaker and audience (a) *have* certain practical and cognitive commitments; (b) *know something of each other's* cognitive and practical commitments; (c) *adhere* to, and act in accordance with, relevant communicative, epistemic, and ethical norms; and (d) *assume that* the other party is acting in accordance with such norms. The conduit and container metaphors hide, or radically downplay, these essential aspects of communicative activity. (p. 40, emphasis original)

On this understanding of the nature of communication, certain norms will be integral to the success of speech acts as such: they impose epistemic requirements that are also concomitantly ethical ones. While it is not possible to give an exhaustive list of such norms since norms can be divided in various ways, Manson and O'Neill develop in the same work a useful categorization of the more significant norms that are likely to be included in any classification (this listing is a direct quotation):

> 1. Norms needed for speech acts to be accessible and relevant to intended audiences (e.g., *intelligibility, relevance*);
> 2. Norms needed for speech acts, and especially those that make truth claims, to be adequately accurate and assessable by intended audiences (e.g., *not lying, deceiving or manipulating; aiming for accuracy; not misleading in other ways; providing relevant qualifications and caveats*). (p.64, emphasis original)

These norms derive from the general nature of communication and, in particular, of communication intended to inform. In related work (Keohane, Lane and Oppenheimer, in progress), I am working with co-authors to develop a list of the particular ethical norms relevant to scientists as they communicate with elite lay audiences such as policy-makers and the policy community (in particular, we consider the communication of the results of scientific assessment studies such as those produced by the IPCC). One version of our working list includes: *honesty, accuracy, audience relevance, process transparency,* and *specification of uncertainty*. Our working hypothesis is that honesty (precluding misstatement and manipulation) is non-negotiable, whereas the other four may have to be traded off in order to achieve the overall goal of effective communication. It may be that an expanded version of Anderson's norm of *epistemic responsibility* should be added to this list, requiring that scientists demonstrate not only transparency and honesty but also a willingness to subject themselves to external peer review and to evince dialogic rationality in their responses to criticism.[9]

For present purposes, what matters is that the ethics of lay democratic judgment must broadly match the ethics of scientific communication, each giving rise to the other through a process of 'reverse engineering'.[10] We can expect them to match most fundamentally because both speakers and auditors in the case of scientific communication with non-experts share the epistemic role of being *inquirers*. Scientific speakers are sharing the results of past inquiry; lay auditors are engaging in current inquiry when they listen to and evaluate those results. They are therefore engaged in a mutually iterative process, even if a given act of communication is framed as directional, from speakers to auditors. (My adoption of the terminology of 'speakers' and 'auditors'

below should be understood in this context, as referring to roles in one particular act of communication, but not ruling out reversal and exchange of roles in others. Whether the oral connotations of these terms is the best choice, or we should instead speak of 'writers' and 'readers' in the context of the web in which so much communication now takes place, is a matter for further consideration.) That mutually iterative framework offers a richer and more dynamic canvas for lay assessments of scientific expertise than the simple one-way container model that the Anderson and Brewer accounts, like many others, presuppose.

Without going into detail for each of the norms on the side of the scientists (which is not the focus of this paper), let us develop a set of corresponding norms for their lay addressees. In each case we will find that a set of epistemic virtues is required in order to develop the ability and disposition to comply with the norm.[11] To begin with honesty: honesty is not limited only to the disclosure of external conflict of interest and the avoidance of misleading statements on the part of scientists which Anderson classed under that name. As Linda Zagzebski explains:

> [I]t is not sufficient for honesty that a person tell whatever she happens to believe is the truth. An honest person is *careful* with the truth. She respects it and does her best to find it out, to preserve it, and to communicate it in a way that permits the hearer to believe the truth justifiably and with understanding. (Zagzebski 1996: 158)

Zagzebski suggests that honesty (which she classes initially as a moral virtue) requires a range of intellectual virtues. And as their shared goals as inquirers would suggest, the virtues she describes are applicable to speakers and to auditors alike. The honest person 'must be attentive, take the trouble to be thorough and careful in weighing evidence, be

intellectually and perceptually acute, especially in important matters, and so on, for all the intellectual virtues' (Zagzebski 1996: 159).

Accuracy and process transparency are more one-sided in the way that they apply differentially to speakers and auditors, yet even here the communicative relation connects them. For speakers, accuracy and process transparency must be gauged in relation to the need to be intelligible to their intended audience. But speakers cannot, of course, always know that audience precisely or delimit it in advance. They have a responsibility to make clear their intended audience, so that auditors can gauge the trade-offs that have been made in aiming to communicate with the audience that the speaker envisioned. But auditors therefore have a corresponding responsibility to assess speakers in relation to their intended audiences, taking account of the limitations and possible misconceptions in the speaker's knowledge of the intended audience and of any divergence between that intended audience and the actual audience. This principle is especially important for cases of 'overheard' speech or 'over-the-shoulder' seen writing, such as leaked emails or leaked accounts of strategy sessions in formulating scientific assessments. In the 'Climategate' emails, for example, advice about how to formulate communication was 'overheard' in a way that cast doubt on the content of the communication itself. There are norms, of course, for formulating communication, and it is a separate question whether the emails breached those norms. My point here is that every 'speaker' goes through an internal process of fine-tuning her message before making it public; to fail to recognize the trial and error nature of such fine-tuning is to apply the wrong interpretative frame, which is likely then to distort the substantive and normative judgment of the fine-tuning process itself.

Of most interest for this paper are the norms for communicating uncertainty, norms that are very much in play in the L'Aquila case. Neither Anderson nor Brewer takes any account of uncertainty as an important feature in scientific discourse. Anderson seeks to give people reason to trust evident experts against crackpots; Brewer is skeptical that people can tell the difference between plausible fakes and real experts, or choose between rival hired guns, at all. For neither is uncertainty central to the story. Yet uncertainty at multiple levels will be ineradicable from virtually any real-world incarnation of the fundamental problem we are considering. For these purposes, 'uncertainty' is closely related to risk assessment. Although in some contexts scientists distinguish between uncertainty (in particular, Bayesian uncertainty, where frequencies of events are not known) and risk (where probabilistic assessments are made on the basis of known frequencies), this distinction is complicated by the fact that what is called 'risk modelling' is undertaken precisely to offer a way of thinking about probabilities in the absence of frequency information. Many discussions of uncertainty and risk in practice tie them closely together, as in the important work on heuristics and biases in risk assessment called *Judgment under Uncertainty* by Kahneman, Slovic, and Tversky (1982).

For heuristic purposes, we may distinguish three different varieties of uncertainty as follows (far more finely grained classifications are also possible). First, there are uncertainties that are intrinsic to the scientific phenomenon being studied. For example, there are uncertainties in the weather system which no method of study could hope to eliminate. Second, there are uncertainties that are conditional on the present state and methods of scientific inquiry. These include both uncertainties about which model to use,

and about how to set the parameters of a given model ('structural' or 'model' uncertainty and 'parameter' uncertainty, respectively). Third, there are what I will call 'competitive uncertainties,' a special form of uncertainty which arises from the phenomenon of disagreement, whether between one scientific team and another in the same subfield, or between one subfield of science and another, or so on. These forms of uncertainty may be present at the same time or not. Competitive uncertainty often arises in part from some underlying intrinsic or conditional uncertainties.[12]

Any or all of the three varieties of uncertainty distinguished above may be viewed differently by the speakers (presumed to be the scientific experts) and their auditors. Competitive uncertainty among scientists, for example, may or may not lead scientists to have less confidence in their own findings so far as they go. But lay audiences who hear about competitive uncertainty between scientists on certain points may become more uncertain about all the claims those scientists make, even those which scientists agree in regarding as relatively certain.

Despite these potential divergences between speakers and auditors in responding to uncertainties, I want to stress some important and deeply shared commonalities among both experts and non-experts in their typical responses. Expert speakers and lay auditors share a general tendency to what has been called an 'overconfidence bias' (Sterman 2011: 816) in relation to uncertainty. A review article on climate change in *Nature* concludes that 'uncertainty breeds wishful thinking' and 'promotes optimistic biases,' leading individuals to "often misinterpret the intended messages conveyed regarding the probabilistic nature of climate change outcomes – and tend to do so over-optimistically' (Markowitz and Shariff 2012: 244). That review is focused on general human

psychology and so relates most immediately to lay people. But two other psychologists specifically report that experts share the same tendency to over-confidence about uncertainty as laypeople: 'There is clear experimental evidence that both experts and laypeople are systematically over confident when making judgments about, or in the presence of, uncertainty' (Morgan and Mellon 2011: 709; see also Kusch 2007).

Both experts and laypeople, then, are in need of norms that check and restrain their overconfidence in assessing uncertainty, and that can guide them in making more accurate assessments (for a model, see Kloprogge, van der Sluijs and Wardekker 2007). Those norms include the explicit recognition of uncertainty in speaking and in receiving communications about scientific knowledge, and explicit self-reflection as to whether one is responding to such uncertainty correctly. There are special difficulties here in knowing how to measure such recognition in relation to the presumed goals of one's audience (as a speaker) and in making allowance for the speaker's likely inability to connect directly with one's own individual goals (as an auditor). Here, the norms of accuracy and of audience relevance must sometimes be traded off, in order to communicate in a way that will enable an auditor to overcome biases or resistances and to receive the full force of the communication. The role of rhetoric in such communication is well described by Victoria McGeer and Philip Pettit: 'The central axiom of rhetoric can be summarised very simply: in persuading others of our point of view, it is often not enough just to make a good case for that point of view; it is also necessary to move or bend your hearers, letting them feel the force of what you have to say' (2009: 65).

Most important is discrimination between the varying kinds of scientific uncertainty and their sources, and refraining from taking one kind of uncertainty to bleed

into another. Aristotle counseled long ago that one can only expect the level of precision for a given domain of expertise to be consonant with the kind of subject matter (*Nicomachean Ethics* I.3, in Barnes 1984). Applying this Aristotelian insight, we may add that the very nature of modern science consists in model and parameter uncertainty since science advances by refining and challenging models and testing the appropriate parameters for them. The existence of such conditional (as well as some intrinsic) uncertainties is, *prima facie*, evidence of proper science being done rather than a reason to doubt its findings.

Competitive uncertainties for their part need to be addressed using the enriched conception of epistemic responsibility and honesty, as well as the necessary but not in themselves conclusive methods of credentials and agreement among scientists (this last somewhat neglected above, but important for both Socrates and Goldman). Most dangerous is the view that any uncertainty at all, in the form of competitive uncertainty provoked by fomenting conditional uncertainties and calling attention to intrinsic uncertainty, is ground for doubt of the results being put forward. For example, a widely reported 1998 'action plan' by a group calling itself the 'Global Climate Science Communications Team' stated this as its goal:

> Victory will be achieved when: Average citizens "understand" (recognize) uncertainties in climate science; recognition of uncertainties becomes part of the "conventional wisdom"; media "understands" (recognizes) uncertainties in climate science; media coverage reflects balance on climate science and recognition of the validity of viewpoints that challenge the current "conventional wisdom"; industry senior leadership understands uncertainties in climate science, making them stronger ambassadors to those who shape climate policy; and those promoting the Kyoto treaty on the basis of extant science appear to be out of touch with reality.
>
> (Congressional Record, vol.144, no.48 (House, April 27, 1998))

Here, the sheer recognition of uncertainties in climate science is presented as reason to prevent action. This stance is taken in respect of any uncertainties at all, without discrimination or reflection on their sources or their implications. A recent PBS 'Frontline' documentary (Upin 2012) surveyed cases of local political decision-making in which such strategic appeals to uncertainty were employed to block research or action on climate change. To avoid allowing these indiscriminate appeals to uncertainty, appropriate recognition of scientific uncertainty must be accompanied by robust explanations of the extent to which such uncertainty is appropriate to the nature of the science, and of just which implications and findings it does and does not call into question (on approaches to uncertainty in the media, see Friedman, Dunwoody and Rogers 1999). This analysis provides another reason why our general accounts of lay judgment must attend to the way laypersons evaluate and determine the implications of *acting* on the claims that experts make, including experts' claims about uncertainty. That Brewer and Anderson ignore the implications of uncertainty for judgment and action provides another reason to move beyond their narrow frames.

IV. Cultivating Common Norms Within Ethical Limits

A shared tendency to overconfidence in the presence of uncertainty is but one of many other cognitive tendencies and biases that experts and laypeople share. Analyzing the climate-economy system, John Sterman argues that even 'highly educated adults with substantial training in Science, Technology, Engineering and Mathematics (STEM) suffer from systematic biases in judgment and decision-making and in assessing the dynamics

of the climate-economy system. There is no reason to believe policymakers are immune to these problems' (2011: 814). These biases, resulting from common heuristics, are now well-known from the pioneering research of Daniel Kahneman and Amos Tversky and their followers. Sterman summarizes them thus:

> We violate basic rules of probability and do not update our beliefs according to Bayes' rule. We underestimate uncertainty (overconfidence bias), assess desirable outcomes as more likely than undesirable outcomes (wishful thinking), and believe we can influence the outcome of random events (the illusion of control). We make different decisions based on the way the data are presented (framing) and when exposed to irrelevant information (anchoring). We credit our personal experience and salient information too highly and underweight more reliable but less visceral data such as scientific studies (availability bias, base rate fallacy). We are swayed by a host of persuasion techniques that exploit our emotions and our desire to avoid cognitive dissonance, to be liked, and to go with the crowd.... (Sterman 2011: 816)

This catalogue does not even include other factors that Sterman discusses, such as the general failure to reason in accordance with sound scientific method; the effects of unconsciously processed conditions (e.g., weather) on our judgments; general ignorance; and faulty mental models. Most important for present purposes, Sterman acknowledges that 'Scientists and professionals, not only "ordinary" people, suffer from many of these judgmental biases' (816). Compare the assessment of geologists attempting to develop methods of expert elicitation to reduce such biases: 'all humans – experts included – are subject to natural biases when trying to estimate probabilities or risks mentally' (Curtis and Wood 2004: 127; see also Polson and Curtis 2010).

The widespread recognition of cognitive biases in both experts and laypersons may tempt despair about the possibility of overcoming these challenges, but there is a silver lining to this analysis. Scientists are able to overcome or minimize these

weaknesses by engaging in learning through the scientific method. Indeed, Sterman suggests that some part of the gulf between scientific experts and laypeople who reject or resist their testimony may arise simply from the fact that scientific experts are engaged in an 'iterative, interactive learning process' in which the latter are not (823). But Stermin recognizes this is not an irremediable gulf. One possible solution, which Sterman develops for the case of climate change, involves engaging laypeople in a form of reasoning developed by considering 'interactive, transparent simulations of the climate' (824). While such simulations are not the only tools scientists use in research, they are among their tools, and because they belong to the context of discovery as well as to the context of communication of results, engagement with these simulations can unite the two contexts in a mode of active understanding.

This kind of iterative engagement can afford members of the public an active sense of being what Aristotle called the 'users' of the products of expertise, users who are entitled to a particular authority in judgment. In another part of the discussion in *Politics* III 11 referred to earlier, Aristotle defends the claim that 'there are some crafts in which the maker might not be either the only or the best judge', that is, in the case of crafts 'where those who do not possess the craft nevertheless have knowledge of its products.' He gives three examples. A head of household judges a house better than its builder; a captain judges a rudder better than its carpenter; and a guest judges a feast better than its cook. Note that the primary verb at work here is the specific verb for *judging* (*krinein*: forms at 1282a18, a21), with the verb for *knowing* (*gignôskein*: forms at 1282a19, a20) being a general one which can mean knowing, but can also mean perceiving, recognizing, or judging, rather than a more specialized verb denoting a specifically theoretical kind of

knowledge or understanding. As judges, the users are acquainted with the products of the arts and so are able to judge their merits although they lack the *technê* necessary to produce them (1282a19). By being invited to engage in such iterative learning processes, the "users" of scientific expertise – the broader public -- will be able to grasp the general structure and nature of scientific reasoning, to gain a feel for uncertainty ranges and their implications, and to understand and assess the context from which scientific expert judgments emerge.

Simulations and similar forms of engagement, in other words, may offer a less demanding but still adequate version of the undergraduate education in science that Brewer required. Though more evidence is needed, such mechanisms may be able to furnish laypeople with the kind of habits of mind and scientific literacy which will enable them to judge experts, and expert claims, without going so far as to become experts themselves. This poses a more dynamic version of a point also defended by Alvin Goldman (2001: 94-97, 107-09), who argues that nonexperts can assess the indirect argumentative justifications offered by (competing) experts. Although laypeople cannot grasp the experts' premises as premises for themselves, Goldman suggests, they can nonetheless come to have reasons to judge one expert's views to be more likely to be true than another's.

Goldman claimed that this is in principle possible. We can go further, in identifying a skill and virtue of good judgment which will make more people likely to be able to develop such reasons. Zagzebski describes something close to this when she identifies a higher-order virtue of cognitive integration, which she celebrates as a matter of 'good intellectual character' (2012: 275). Philip Tetlock develops a similar idea in this

space by describing good judgment as a 'meta-cognitive skill' (2009: 23). Tetlock considers this skill as one that helps to explain the differential accuracy of experts (in his research, putative political experts) in terms of correspondence of their judgments to reality and the coherence of those judgments. He argues that the use of good judgment by 'foxes' who use good judgment in attending to a wide range of factors and being open to a wide range of explanations, explains their higher accuracy in making judgments and avoiding undue defensiveness about their errors as compared to 'hedgehogs' who cling to a single big idea. This is a kind of judgment of experts that ordinary citizens could learn to make. It is not restricted to second-order assessments of credentials and credibility, but rather reaches into first-order scientific explanations, attending to the range of factors considered and to the way that rival explanations are entertained or dismissed, and so going beyond the extreme cases of spotting sheer dialogic irrationality to which Anderson limits her second-order criterion. Such an assessment of explanation or argument, and of its success, does not require the auditors to share exactly the same cognitive competence and resources being assessed. But it does require an ability to engage in relevant forms of reasoning, and an ability to assess patterns of reasoning which may at least in principle be displayed by experts and non-experts alike.

Thus, a common self-awareness, a common engagement in learning (even if not at the same level of epistemic complexity and sophistication), and a common good intellectual character – with habits of epistemic virtue and associated skills -- can bridge the capacities of the expert and nonexpert (on the role of rhetoric in improving one's personal judgment, see McGeer and Pettit 2009). Pushing further open the door inadvertently left ajar by Brewer, who allowed judges with a basic scientific education to

judge outside their discipline, we can suggest that there is no sharp binary line between expert and non-expert that is pre-given. Rather, in different disciplines and contexts, non-experts may be able to develop the repertoire of skills, habits and dispositions which can enable them to judge certain scientific claims well. Institutions of public policy, public deliberation, and public communication need to cultivate these capacities and virtues further, embedding and supporting them in a common culture of inquiry (Christiano 2012, Koppl 2005, Anderson 2006).[13] There is no reason in principle to give up on the possibility of democratic judgment of expertise – but there is every reason in practice to try to create conditions in which it is more likely to be exercised well.

V. The L'Aquila Case

The conviction of six seismologists and a government official on charges of manslaughter in an Italian court in L'Aquila on 22 October 2012 provides a tense and tragic case for consideration of these issues. All were members of the National Commission for Forecasting and Predicting Great Risks, which was called to hold a special meeting in L'Aquila on 31 March 2009, in which they were 'asked to assess the risk of a major earthquake in view of many shocks that had hit the city in the previous months' (Nosengo 2012). After the meeting, one of the scientists appeared at a press conference alongside Bernardo De Bernardinis, then vice-director of the Department of Civil Protection and the seventh convicted in the recent trial, and two local officials, and De Bernadinis offered press interviews both before and after the meeting. In the course of these public communications, he characterized the recent wave of seismic tremors in the region as "certainly normal" and posing "no danger" (Hall 2011: 268), comments

which, according to the later prepared minutes, do not accurately capture the remarks made by scientists at the meeting itself. Before the meeting, De Bernardinis also said that "the scientific community continues to assure me that, to the contrary, it's a favourable situation because of the continuous discharge of energy," a claim from which several of the scientists have dissociated themselves and which does not appear in the official minutes. Those minutes were not written until after the earthquake, nor was the customary formal statement of the commission after a meeting issued at the time, so that the press conference and interviews were 'the only public comments to emerge immediately after the meeting' (Hall 2011: 268).

The comments by De Bernardinis above seem to be clearly inaccurate, though their inaccuracy pertains to him rather than to the views of the scientists recorded at the meeting. Their inaccuracy would constitute one ground for considering him, at least to have engaged in misleading conduct. But let us say that the commission was accurate in its actual meeting, at which, according to the minutes, seismologist Boschi said, "It is unlikely that an earthquake like the one in 1703 [which had destroyed L'Aquila] could occur in the short term, but the possibility cannot be totally excluded" (Hall 2011: 267). That claim raises a difficulty for how citizens and government should respond to events that are unlikely but possible. If scientific assessment is limited to probabilities that are themselves characterized by uncertainty, the question of the appropriate response remains to be determined by considerations of value and public policy. Here it is worth noting that another civil servant was investigated in Italy in 1985 for an opposite kind of failure. As one of the convicted men's lawyers was reported in *Nature* as saying: '[T]he then-head of civil protection, Giuseppe Zamberletti, was investigated for instigating a public

panic when he ordered the evacuation of several villages in northwest Tuscany after a seismic swarm [the same phenomenon that had afflicted L'Aquila in the months before the big quake]; on that occasion, no major quake occurred' (Hall 2011: 269, citing Alfredo Biondi).

The actual indictment, however, went beyond the question of the objective findings of risk by the commission. It charged the commission with failing to meet its legal charge 'to avoid death, injury and damage, or at least to minimize them' (same). On this view, the commission's duties were not limited to making an assessment of the probabilistic risk of an earthquake (one which De Bernardinis seems to have reported inaccurately). Contrary to much press report, they were not indicted for failing to predict the terrible earthquake which struck on 5 April 2009, but for failing in the role of a scientific commission charged not only with informing but also with advising action to meet the goals of avoiding or minimizing death, injury and damage from natural hazards. Specifically, they were charged with having given "incomplete, imprecise, and contradictory information" to the residents of L'Aquila, including failing to take into account the nature of its fragile buildings and dense urban population in quantifying the risk that an earthquake of a certain magnitude would pose (Hall 2011: 266). On this view, the duty of the commission was not only to estimate the probabilistic risk of an earthquake, but also to estimate the degree of damage that such an earthquake could cause, so estimating also what we might call the 'social risk' of a quake and advising and reinforcing messages of earthquake preparedness in its light. A very small likelihood of a quake could still pose a significant risk in this particular town for these reasons: the prospect of great harm increases the risk posed even by a relatively unlikely event. (One

might however question whether a panel of six seismologists and a civil servant, without a single civil engineer, was best suited to make that judgment; and some critics have suggested that even such an assessment would still have concluded that 'the hazard level in L'Aquila...was insufficient, by two to three orders of magnitude, to justify evacuation of even the weakest buildings' (Vidale 2011: 324).)

The overall charges against the convicted men might be broadly summed up in the claims that they failed to give a message of earthquake preparedness and to remind residents of the potential very great damage that an earthquake could cause. This leaves open the question of how residents could and should have assessed the prospect of risk to their lives and property in light of the uncertainty attending earthquake prediction in general. One of the most striking and disturbing features of the case is the testimony offered by L'Aquila residents that the absence of such warnings -- and the reassuring, optimistic messages received from the press conference and interviews -- led them actually to disregard the accumulated 'lay expertise' of the town about how to respond to tremors. Whereas their parents and grandparents had fled buildings on the least tremor, and many of the aged are said to have done so in April 2009, the more educated members of the L'Aquila public claim to have been swayed by what was, and wasn't, in the publicly transmitted press release and press conference message to read it as a message of reassurance, that it would be irrational to fear unduly, and so to stay inside on the fatal night (Hall 2011).

One final feature of this disturbing case is what seems to have occasioned the extraordinary meeting of the commission in March 2009 in the first place. This was the activity of a local resident, Giampaolo Giuliani, who had begun to make unofficial public

earthquake predictions on the basis of radon gas level measurements (Hall 2011: 267). It has been alleged that the head of the Department of Civil Protection, Guido Bertolaso, called the extraordinary meeting to assuage the public unease and confusion being caused by Giuliani. This may explain why the procedures were not as normal, failing to offer a formal statement and so opening the door to the misinformation purveyed in the press conference and press interviews instead. If this is so, it means that in the attempt to refute a man whom they saw as a fraud and crackpot, the Italian authorities arguably neglected to furnish citizens with a sufficient public account of the reasoning and uncertainty of scientific experts that they could assess. In other words, the Italian state fell prey to the limitation of overemphasis on refuting crackpots that I have diagnosed in the philosophical literature. Because they failed to broaden the agenda of their public communications to provide citizens with a sound basis for exercising democratic judgment of the claims made by the genuine experts, the Italian state – perhaps in the composition of its risk commission, but also in failing to provide its full scientific view and instead offering only summary and arguably inaccurate reassurances -- failed its citizens to disastrous effect. This catastrophe highlights the need to revise the current the philosophical discussion of lay expertise to be more attentive to the challenges of uncertainty and the way citizens must assess expert claims in order to act. The tragedy in L'Aquila points to the high stakes of the problem and the need to develop approaches that recognize the value of lay judgment and responsible scientific communication.[14]

Works Cited

Anderson, Elizabeth (2006). 'The epistemology of democracy,' *Episteme* 3: 8-22.

----- (2011). 'Democracy, Public Policy, and Lay Assessments of Scientific Testimony,' *Episteme* 8: 144-164.

Baker, Sherry, and David L. Martinson (2001). 'The TARES Test: Five Principles for Ethical Persuasion,'*Journal of Mass Media Ethics* 16: 148–175.

Barnes, Jonathan (ed.) (1984). *The Complete Works of Aristotle: The Revised Oxford Translation*, vol. 2 (Princeton: Princeton University Press).

Brewer, Scott (1997-98). 'Scientific expert testimony and intellectual due process,' *Yale Law Journal* 107: 1535-1679.

Brossard, Dominique, and Bruce V. Lewenstein (2010). 'A Critical Appraisal of Models of Public Understanding of Science: Using Practice to Inform Theory,' in LeeAnn Kahlor and Patricia A. Stout (eds.) *Communicating Science: New Agendas in Communication* (New York: Routledge), 11-39.

Brown, Mark B. (2009). *Science in Democracy: Expertise, Institutions, and Representation* (Cambridge, Mass: MIT Press).

Christiano, Thomas (2012). 'Rational Deliberation among Experts and Citizens,' in John Parkinson and Jane J. Mansbridge (eds.) *Deliberative Systems: Deliberative Democracy at the Large Scale* (Cambridge: Cambridge University Press, 2012), 27-51.

Curtis, Andrew, and Rachel Wood (2004). 'Optimal Elicitation of Probabilistic Information from Experts,' in Andrew Curtis and Rachel Wood (eds.) *Geological Prior Information: Informing Science and Engineering* (London: Geological Society, no. 239), 127–45.

Estlund, David (1993). 'Making truth safe for democracy,' in David Copp, Jean Hampton, and John E. Roemer (eds.) *The idea of democracy* (Cambridge: Cambridge University Press), 71-100.

----- (2008). *Democratic Authority* (Princeton and Oxford: Princeton University Press).

Friedman, Sharon M., Sharon Dunwoody, and Carol L. Rogers (eds.) (1999). *Communicating Uncertainty: Media Coverage of New and Controversial Science* (Mahwah, N.J: Lawrence Erlbaum Associates).

Gentzler, Jyl (1995). 'How to Discriminate between Experts and Frauds: Some Problems for Socratic Peirastic,' *History of Philosophy Quarterly* 12: 227-246.

Goldman, Alvin I. (1999). *Knowledge in a Social World* (Oxford: Clarendon Press).

----- (2001). 'Experts: Which Ones Should You Trust?,' *Philosophy and Phenomenological Research* 63: 85-110.

Hall, Stephen S. (2011) 'Scientists on trial: At fault?,' *Nature* 477: 264-69.

Hardwig, John (1985). 'Epistemic Dependence,' *The Journal of Philosophy* 82: 335-49.

Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.) (1982). *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge and New York: Cambridge University Press).

Keohane, Robert O., Melissa Lane and Michael Oppenheimer (in progress). 'The Ethics of Scientific Communication under Uncertainty.'

Kloprogge, Penny, Jeroen van der Sluijs, and Arjan Wardekker (2007). *Uncertainty Communication: Issues and Good Practice* (Utrecht: Copernicus Institute for Sustainable Development and Innovation).

Koppl, Roger G. (2005), 'Epistemic Systems,' *Episteme* 2: 91-106.

Kusch, Martin (2007). 'Towards a Political Philosophy of Risk: Experts and Publics in Deliberative Democracy, in Tim Lewens (ed.) *Risk: Philosophical Perspectives* (London and New York: Routledge, 2007), 131–155.

Kutz, Christopher (unpublished). 'Epistemethics.'

LaBarge, Scott (1997). 'Socrates and the recognition of experts,' in Mark McPherran (ed.) *Wisdom, ignorance, and virtue: new essays in Socratic studies* (Edmonton: Academic Printing and Publishing), 51-62.

Lane, Melissa (1998-99). 'States of Nature, Epistemic and Political,'*Proceedings of the Aristotelian Society*: 1-24.

----- (forthcoming). 'Claims to Rule: The Case of the Multitude,' in Marguerite Deslauriers and Pierre Destrée (eds.) *The Cambridge Companion to Aristotle's* Politics (Cambridge: Cambridge University Press).

Manson, Neil C, and Onora O'Neill (2007). *Rethinking Informed Consent in Bioethics* (Cambridge and New York: Cambridge University Press).

Markowitz, Ezra, and Azim F. Shariff (2012) 'Climate Change and Moral Judgement,' *Nature Climate Change* 2: 243–247.

McGeer, Victoria, and Philip Pettit (2009). 'Sticky Judgement and the Role of Rhetoric,' in Richard Bourke and Raymond Geuss (eds.) *Political Judgement: Essays for John Dunn* (Cambridge and New York: Cambridge University Press), 48-73.

Moore, Alfred, and John Beatty (2011). 'Political Authority and Scientific Authority: What Does Deference Mean?", paper presented at UBC Workshop on Scientific Authority in Democratic Societies (on file with author).

Morgan, M. Granger, and Carnegie Mellon, (2011). 'Certainty, uncertainty, and climate change,' *Climatic Change* 108: 707-721.

Nosengo, Nicola (2010). 'Italy Puts Seismology in the Dock,' *Nature News* 465.7301: 992–992.

----- (2012). 'Italian Court Finds Seismologists Guilty of Manslaughter.' *Nature* News, 22 October, corrected 23 October, online.

Polson, Debbie, and Andrew Curtis (2010). 'Dynamics of Uncertainty in Geological Interpretation,' *Journal of the Geological Society* 167: 5–10.

Raz, Joseph (1986). *The Morality of Freedom* (Oxford: Clarendon Press).

----- (2005-06). 'The Problem of Authority: Revisiting the Service Conception', 90 *Minnesota Law Review*: 1003-44.

Reeve, C.D.C. (trans. and ed.) (1998). *Aristotle* Politics (Indianapolis: Hackett).

Russell, Nicholas J. (2010). *Communicating Science: Professional, Popular, Literary* (Cambridge and New York: Cambridge University Press).

Sterman, John D. (2011). 'Communicating climate change risks in a skeptical world,' *Climatic Change* 108: 811-826.

Tetlock, Philip E. (2009). *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton: Princeton University Press).

Upin, Catherine (dir.) (2012). 'Climate of Doubt,' *Frontline*, broadcast on PBS, 23 October 2012.

Vidale, John E. (2011). 'Italian quake: critics' logic is questionable,' *Nature* 478: 324 [correspondence].

Zagzebski, Linda Trinkaus (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge and New York: Cambridge University Press).

----- (2012). *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief* (Oxford and New York: Oxford University Press).

[1] I will work with the following definitions proposed by Scott Brewer in an important article, albeit one with which I will later take issue on other grounds. Drawing on the pioneering work of Alvin Goldman, Brewer defines the relevant terms as follows: 'An *expert* is a person who has or is regarded as having specialized training that yields sufficient epistemic competence to understand the aims, methods, and results of an expert discipline. An *expert discipline* is a discipline that in fact requires specialized training in order for a person to attain sufficient epistemic competence to understand its aims and methods, and to be able critically to deploy those methods, in service of these aims, to produce the judgments that issue from its distinctive point of view. A *nonexpert* is a person who does not in fact have the specialized training required to yield sufficient epistemic competence to understand the aims, methods, and judgments of an expert discipline, or to be able critically to deploy those methods, in service of the discipline's aims, to produce the judgments that issue from the discipline's distinctive point of view (Brewer 1998: 1589, emphasis added). I will generally use *layperson* in place of *nonexpert*.

[2] Alfred Moore and John Beatty (2011) explore a contrast between acceptance of an authority's claims and belief in that authority's claims which is helpful in clarifying the epistemic state to which judgment may give rise. 'Judgment' however has the advantage of also describing an epistemic faculty and activity, and in picking out a small class of beliefs that are 'reflectively available,' as per McGeer and Pettit 2009: 49.

[3] This approach is potentially compatible with an alternative starting point, that of the epistemological literature on the inescapability of reliance on testimony; the advantage of the Razian framework is that it makes the link to action, and the parallel between practical and theoretical experts and hence authorities, more explicit.

[4] Many other authors discussing modern expertise begin with Socrates, for example Brown 2009: 9.

[5] I prescind from his discussion of the special context of rules of evidence, relevance, and admissibility in American courts of law, including the problematic transition from the Frye rule to the Daubert rule (which in practice, he suggests, often reduces back to the former).

[6] Compare the regress posed by Socrates in the *Meno*: to seek knowledge, one must be able to identify the knowledge one seeks, but that requires having the knowledge that one is seeking. For Socrates, the solution to the regress is a universal innate capacity underwritten by the metaphysics of recollection. While I will not appeal to this metaphysics, the idea that there are potential capacities that can be developed to solve the regress is a useful clue.

[7] A challenge to both Brewer and me comes from the fact that over 30,000 people with at least an undergraduate degree in the sciences signed a petition against the scientific credibility of global warming (the so-called Oregon Petition, as described in Upin 2012).

[8] Anderson ignores the problem of whether one can rightly presume a sufficient degree of at least partial independence for numbers to lend genuine additional credence on Bayesian principles, discussed by Goldman 2001: 98-104.

[9] Compare the norms suggested for a rather different set of communicators, advertisers and public relations practitioners – truthfulness, authenticity, respect, equity, and social responsibility – offered by Baker and Martinson 2001.

[10] I owe this term to Bob Keohane.

[11] I focus here on epistemic virtues, but there are also extremely relevant ethical virtues which may serve as a guide to the epistemic, as explored in an unpublished paper by Christopher Kutz. The fruitfulness of thinking about virtue in this context was first suggested to me by Michael Lamb.

[12] While this three part classification is my own, its constituent elements are very much informed by a collaborative typology of uncertainty initiated by Felix Creutzig and further developed by him and other members of the Princeton Institute for International and Regional Studies research network on Communicating Scientific Uncertainty to which I belong. For the project in general, see http://www.princeton.edu/piirs/research/research-communities/communicating-uncertainty.

[13] Thomas Christiano (2012) develops an illuminating account of the 'discursive relations among experts and citizens' (p.51). He is between Brewer and Anderson in his optimism that such relations can support good citizen judgment: he holds that expert views can constrain legitimate options in policy-making, but not determinatively select the best (it can ensure that policy is made 'consistent with the theories that remain acceptable to the expert community' but not necessarily with the best such theories, p.51).

[14] I am indebted to the PIIRS 'Communicating Uncertainty' research network for stimulating conversations, and especially to Bob Keohane and Michael Oppenheimer for work on a related co-authored project for the same network; Julie Rose and Michael Lamb for extremely valuable discussion and research assistance in the first and second phases of the paper respectively; the students in my Princeton graduate seminar on 'Knowledge and Politics' in fall 2010, with whom I first developed my thinking on certain aspects of this topic; and to others including Bob Brulle, Craig Murphy, Ken Schultz, Tim Schroeder and Joel Watson for conversations at CASBS in fall 2012. I am grateful to Eric Beerbohm and Larry Lessig for the opportunity to present the paper for its first outing in a public lecture at the Edmond J. Safra Center for Ethics, Harvard University; Christopher Kutz for the invitation to present a revised version in the GALA seminar series of the Kadish Center for Morality, Law and Public Affairs at Berkeley Law School; and Jon Krosnick for the invitation to present related work to the Stanford Political Psychology Research Group. Participants in all these fora provided extremely helpful comments and challenges. Support for my research has been generously provided by Princeton University; in 2012 by a Fellowship of the Guggenheim Foundation; and in the 2012-13 academic year by a Fellowship at the Center for Advanced Study in the Behavioral Sciences at Stanford University.