

Group Testimony? The Making of A Collective Good Informant

Introduction

We gain information from collective, often institutional bodies all the time—from the publications of committees, news teams, or research groups, from web sites such as *Wikipedia*, and so on—but do these bodies ever function as genuine *group testifiers* as opposed to mere group sources of information? In putting the question this way I invoke a distinction made, if briefly, by Edward Craig, which I believe to be of deep significance in thinking about the distinctiveness of the speech act of testimony. The distinction is that between somebody's functioning as a 'good informant' and their functioning merely as a 'source of information'. The difference between these has, as he remarks, a crucial ethical aspect:

What I have in mind is the special flavour of situations in which human beings treat each other as subjects with a common purpose, rather than as objects from which services, in this case true belief, can be extracted.¹

In this paper I shall try to bring out the distinctive nature of the role of good informant in a way that helps to clarify what is at stake in asking whether there can be group testimony in the sense of genuinely collective testimony—that is to say, where the group testifier is a collective and not merely a sum of individuals testifying in one or another form of aggregated chorus. I aim to present Craig's distinction between a good informant and a source of information in a way that brings out its natural persuasiveness independently from more contentious commitments of the State of Nature story in which it is embedded, so that I can use it to lend some explanatory support to an assurance view of testimony.² I then defend the assurance view against

¹ Edward Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Clarendon Press, 1990), p. 36.

² The assurance view can take different forms, and one or another version has been advanced by Angus Ross, 'Why Do We Believe What We're Told?', *Ratio* 28 (June), 1986, 69-88; Edward Hinchman, 'Telling as Inviting to Trust', *Philosophy and*

an important line of criticism relating to its contribution to the epistemology of testimony. With an assurance approach duly defended, I go on to use its trust based conception of testimony as the model for what it would take for a group to testify. I argue that the requisite kind of collective will be one capable of forming a group commitment to trustworthiness in respect of the subject matter of the testimony, and I propose a way of establishing that kind of commitment using Margaret Gilbert's notion of joint commitment.³ Finally, I make a suggestion about why it matters in a democracy that at least some collectives, notably some institutional bodies, should be so constituted that they are indeed capable of collective testimony, rather than being able to convey information only by functioning as sources of information.

What Makes a Good Informant?

Let me start by presenting Craig's notion of a good informant, in order to make plain the initial contrast with a source of information. Craig sets about giving a 'practical explication' of why we have the concept of knowledge over and above the concept of true belief by telling a fictional story about an epistemic State of Nature. The story is presented as showing that, regardless of whatever the actual history of the concept of knowledge may be, it is inevitable that we have such a concept. This inevitability is owing to the fact that 'knower' is the descendent of a concept which would be realized in any recognizably human society in virtue of the fact that it meets absolutely fundamental human epistemic needs. This humanly necessary ancestor concept is the concept of a *good informant*, and it arises like this. In the imagined, socially stripped down State of Nature scenario, human beings are identified as having a trio of fundamental epistemic needs, the first being the most fundamental of

Phenomenological Research Vol. LXX, No. 3, May 2005, 562-87; Richard Moran, 'Getting Told and Being Believed', in J. Lackey and E. Sosa eds. *The Epistemology of Testimony* (Oxford: Oxford University Press, 2006). And, with more of an exclusive emphasis on the epistemology, see Paul Faulkner, 'Norms of Trust', in A. Haddock, A. Millar, and D. Pritchard eds. *Social Epistemology* (Oxford: Oxford University Press, 2010).

³ For the original statement see Margaret Gilbert, *On Social Facts* (Princeton, New Jersey: Princeton University Press, 1989). See also *Sociality and Responsibility: New Essays in Plural Subject Theory* (Rowman and Littlefield, Lanham, MD, 2000).

all, but giving immediate rise to the second and then the third. These needs are: first, the need to acquire enough true beliefs (and not too many false ones) for basic practical success, most obviously in terms of sheer survival; second, the need to take advantage of information possessed by others rather than relying exclusively on the comparatively poor resources of one's own faculties; and third, the need to distinguish good informants from bad. In short, human beings need information to survive, and so they had better learn how to pool it, in which case they had better develop a filtering practice so that only good information enters the pool. Such are the most basic epistemic pressures, and the upshot of these pressures is the constructed concept of the *good informant*—someone who will give an inquirer poolable information.

The good informant is a near-idealized testifier. He has three essential features: he is likely enough in the context to be right as to whether *p*; channels of communication with him are open (notably, they are not blocked by insincerity or withholding); and, finally, he bears indicator properties such that the inquirer can recognize the presence of the first two features.⁴ A good informant concerning a need to find out whether, for instance, there is a predator approaching, might be a peer who you can tell can see for miles because she is up a tree. Most crucially for our purposes, however, there is a pregnant contrast between someone's functioning as a good informant and their functioning merely as a *source of information*. Now the idea of a source of information is a thoroughly generic idea. Paradigm examples of sources of information are the usual suspects: clocks, barometers, tree stumps, and other epistemic instruments or gauges. But they can also be people. If, after watching the match, my friend returns home with a sun burned nose, he can function for me as a source of information about the afternoon's weather (evidently, it was sunny). He does so by being a feature of a state of affairs that has informational value for me,

⁴ The second and third conditions fade in the transition from imagined State of Nature to real world, whereby we come to conceive of knowledge as something possessed by an individual independently of whether she will share it, and whether one can recognize that she has it. This change takes place because while in the State of Nature our proto-concept of 'good informant' is constructed out of immediate, here-and-now subjective epistemic needs ('Who can tell me whether *p*?'), our real-world epistemic needs are more socially and temporarily extended, so that my need to know might be in respect of a future prospect, and/or it might be such that *I* need never know, so long as *someone* in the epistemic community does, or will do, as and when the need may arise.

given that I have at least some, perhaps very minimal, interest in knowing what the weather was like, and that I am in a position to interpret his telltale nose.

But (let me push the idea to reveal its limits) it is not only in virtue of their appearance that people can be sources of information. Their behaviour, including their linguistic behaviour, can also feature in informationally valuable states of affairs. Imagine this time that from upstairs I hear my friend come through the front door muttering irritably to himself ‘I forgot that sun cream!’. Here again he is functioning for me as a feature of a state of affairs from which I can learn something, perhaps by making an inference or two, or possibly just by perceiving the situation in a certain light. Next, let us alter the example again to augment still further the informational value of words uttered. Imagine this time that I hear my friend come in the front door and tell a companion, ‘I forgot that sun cream!’. Here too, even though this is a piece of testimony I overhear, my friend is functioning *for me* as a source of information and not a good informant. His words (his testimony to someone else) have informational value for me, and I infer or perhaps non-inferentially glean from them that the afternoon’s weather called for sun cream. Even though there is testimony involved in this observed state of affairs, and the speaker is functioning for his interlocutor as a good informant, nonetheless he does not play that role vis-à-vis me, but rather features for me as a source of information.

To explain why this is so we need first to point to the simple and obvious fact that he hasn’t *told me* anything—he has not addressed me at all, and in particular not in response to any real or apparent need for information on my part. But, secondly, we need to explore a distinctive aspect of the role of good informant that is very far from simple and obvious, and this is best done by returning to Craig’s tantalizing suggestion that there is something distinctive about the role of good informant; that there is a ‘special flavour of situations in which human beings treat each other as subjects with a common purpose, rather than as objects from which services, in this case true belief, can be extracted’. The special flavour seems to come from the positive ethical relationship involved when we relate to another person as a fellow epistemic agent. There is no such positive ethical relationship to my reaping informational value from a state of affairs in which noses are red or words are heard

said; but there is, by contrast, an ethical special flavour to my learning from someone's *telling me* that it was very sunny at the match today. How so?

A first point is to acknowledge that the relation of telling between an informant and an inquirer is paradigmatically a matter of the informant responding to an inquirer's epistemic need, and so their relation is *interpersonal*—indeed it is a specifically second-personal (I-you) relation. (Note that this is a definitive feature of Craig's conception, as is enshrined in his self-conscious methodological shift from what Bernard Williams once described as the perspective of the 'examiner' so common in epistemology to the perspective of the 'inquirer', which has more credentials naturalistically speaking, given that knowledge is a name for something important that we need and all too often lack.⁵) The point holds not just for the near-idealized testimony of the good informant, but for testimony in general. C. A. J. Coady observed the interpersonal nature of testimony, as generated by its inherent aim to inform in response to epistemic need. Indeed he rendered it as a necessary condition, stating that testimony is always 'relevant to some disputed or unresolved question (which may, or may not be, *p*?) and is directed to those who are in need of evidence on the matter'.⁶ Given that when someone testifies she is giving an interpersonal response to epistemic need, great or small, then it is no surprise to find that when this response is trustworthy, there is something of both epistemic and ethical worth going on: the informant acts in a way that bends her immediate practical and epistemic ends to align them with those of her inquirer. She must pause to ask herself whether *p*, and must come out with the answer for the benefit of the inquirer. This, I propose, is how we should begin to elaborate the ethically positive sense of 'shared ends' that Craig rightly discerns as a distinctive feature of the cooperative practice of good informing.

A second point is to see that this interpersonal structure of good informing and testimony in general is connected with the fact that testifying is an illocutionary speech act—something exploited in Coady's account of the distinctive way in which

⁵ Bernard Williams, 'Deciding to Believe', *Problems of the Self: Philosophical Papers 1956-1972* (Cambridge: Cambridge University Press, 1973); p. 146.

⁶ C. A. J. Coady, *Testimony: A Philosophical Study* (Oxford: Clarendon, 1992), p. 42.

it fulfils the illocutionary point of informing an audience ('through the speaker's status as one having a particular kind of authority to speak to the matter in question, a matter where evidence is required'⁷). The mark of an illocutionary speech act is that to be fully successful it must secure uptake on the part of the addressee.⁸ In the case of testimony, then, the person in the role of inquirer (who may or may not have made an explicit inquiry) must recognize the informant's intention to perform that very speech act. You cannot fully successfully testify to me (tell me) that the bus will be arriving soon unless I recognize your intention so to tell me. Without my recognition of your intention, you may as well be airing a hypothesis, or simply saying something out loud that has no communicative significance for me, like rehearsing a line from a play. Without my uptake, whatever you may succeed in doing with your words, it won't be quite testifying.

Manifestly my own thinking is in line with those who take the speech act of testimony to be that of telling.⁹ This seems a fairly clear and natural category within human activities of informational exchange, as is suggested by the fact that telling is what the good informant does in relation to her inquirer—telling is the speech act that fits with the genealogically basic category of the good informant. There is room, however, for a range of reasonable philosophical constructions of something called 'testimony', as the literature reveals. On one side, for instance, Coady construes testimony as a narrower class than the general class of tellings. While explicitly acknowledging that there is room for more than one philosophically sensible way of delineating the concept,¹⁰ for purposes of his own account Coady takes his cues from the case of formal testimony in a courtroom, which leads him to cast testimony as the

⁷ Coady, 1992, p. 43.

⁸ J. L. Austin, *How to Do Things with Words* (Cambridge, Mass.: Harvard University Press, 1962).

⁹ Those who advance an assurance view of testimony naturally take this line (see note 2), though advocating an assurance view is not a requirement for regarding testimony as the broad class of tellings—see for instance Elizabeth Fricker, 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony', *Mind* 104 (414) 1995: 393-411.

¹⁰ Coady, 1992, p. 25.

class of *competent* tellings—the general conception obviously harmonizing with his construal of testimony’s illocutionary point as a matter of presenting one’s word as evidence.¹¹ On the other side, by contrast, Jennifer Lackey construes testimony as a significantly broader class than that of tellings, delineating testimony in relation to the two-sided, disjunctive business of either reasonably intending to convey the information that *p* from the communicable content of one’s words *or* being reasonably taken as conveying the information.¹² The upshot is that for Lackey testimony is a capacious and internally diverse category, which most notably includes cases that are not cases of the speaker telling the hearer anything, as in examples of overhearing or eavesdropping; and cases where the hearer gains knowledge that *p* from the speaker’s words even while the speaker lacks knowledge or even the belief that *p*.¹³

I believe there is a good case for taking learning-from-words as a significant category for epistemology to focus on—indeed when one looks to Coady’s original way of drawing our attention to the sheer ubiquity of knowledge acquired on the basis of the word of others, one finds all sorts of examples that are hard to construe as any kind of telling, but which are nonetheless cases of learning from words.¹⁴ However, I

¹¹ Coady, 1992, p. 42.

¹² Jennifer Lackey, *Learning From Words: Testimony as a Source of Knowledge* (Oxford: Oxford University Press, 2008), pp. 35-36.

¹³ Lackey, 2008. For her eavesdropping example, see pp. 233-4; for the example of the creationist teacher, who conveys knowledge of evolution she does not herself possess, see for instance pp. 48-53.

¹⁴ In order to enhance our general sense of the importance and centrality of testimonial knowledge, he says, for instance, ‘That I am so many years old; that I was born on such and such a date; that number H11200 does indeed correspond to the number the Australian passport authorities have in their files—none of these are facts of my individual observation or memory or inference from them. They are based, sometimes in a complex way, on the word of others’ (Coady, 1992, p.7). That is true; but they may not involve the speech act of telling (nor, *a fortiori*, competent telling). I imagine he would agree; my point is only to acknowledge that inasmuch as this kind of ubiquity argument motivates philosophical attention to testimony, it does so by first motivating attention to the epistemic significance in our lives of something more general, namely, the many diverse and indirect ways we learn from the words of others.

myself think there is ample reason to construe testimony as a more distinctive speech act than this, narrower than Lackey's disjunctivist learning from words, and broader than Coady's competent tellings. The distinctiveness of the speech act of testimony (assuming that is what we call it—no need to argue over a word) can be illuminated if we continue homing in on the ethical special flavour inherent in that near-idealized paradigm of testimony, good informing. We have already seen that it involves the informant bringing her immediate epistemic and practical ends into line with those of the inquirer. But there is something else here, more distinctive still, which will help us deepen our understanding of the contrast with someone's functioning as a source of information. In the cooperative epistemic practice that grows up in the State of Nature, inquirers do not merely rely on their good informants, relating to them as more or less complex epistemic instruments; rather they *trust* them.¹⁵ This is all but explicit in the very structure of the good informant. She is both competent and willing to tell you truly what you need to know. Similarly, a trustworthy person vis-à-vis a given practical matter has the competence and (something like) the goodwill to do what you trust them to do.¹⁶

Let us explore what is involved in such trust. In Annette Baier's influential take on the distinctiveness of trust as opposed to sheer reliance, we find effectively indicated our core distinction between a good informant and a source of information:

We all depend on one another's psychology in countless ways, but this is not yet to trust them. The trusting can be betrayed, or at least let down, and not just disappointed. Kant's neighbors who counted on his regular habits as a clock for their own less automatically regular ones might be disappointed with

¹⁵ For a genealogical argument specifically for 'affective trust' (as opposed to merely 'predictive trust'—reliance) as basic, which revises Bernard Williams' construction of truthfulness in the State of Nature, see Paul Faulkner, 'A Genealogy of Trust', *Episteme: A Journal of Social Epistemology*, Vol. 4, Issue, 3, 2007, 305-321.

¹⁶ For the view that trust is essentially 'an attitude of optimism that the goodwill and competence of another will extend to cover the domain of our interaction with her...' see Karen Jones, 'Trust as an Affective Attitude', *Ethics*, 107 (October 1996), pp. 4-25; p. 4.

him if he slept in one day, but not let down by him, let alone had their trust betrayed. When I trust another, I depend on her goodwill towards me.¹⁷

Kant's neighbours are envisaged as using him as a source of information about the time—in this he functions for them as a human timepiece. In this relation there is no trust, only reliance. Now if we follow Baier's imagined circumstance in which Kant sleeps in late one day so that his neighbours' morning schedule goes awry, the appropriate reaction on their part would not involve any ethically loaded sense of being let down or betrayed. It would simply be as if their clock were running slow (and perhaps they might not rely on it so much in future). If, by contrast, we imagine any of them actually asking Kant the time, so that he responds to their epistemic need by telling them it is five past nine in the morning, then they are now envisaged in a relation of trust. Were Kant to give an epistemically irresponsible answer (perhaps, aware that his watch is unreliable, he nonetheless reads the time off it without warning), or alternatively a mendacious answer (perhaps, fleetingly a servant of heteronomy, he simply can't resist the opportunity to mess up her morning), then he does let down his inquirer, and in a special way that licenses her feeling the distinctive resentment of a trust betrayed, however mundane and inconsequential it may or may not be.

I have talked naturally of resentment in respect of a failure of trust, and this reminds us of something Richard Holton has explored in the same connection. Sympathetic to Baier's main point that while failures of reliance are disappointing or annoying, only failures of trust are ethically loaded lettings-down or forms of betrayal, he nonetheless presents decisive counter-examples to her claim that the presence of goodwill is what makes the difference. He argues that goodwill is neither necessary for trust ('One member of an estranged couple between whom there is precious little goodwill can still trust the other to look after the children') nor sufficient ('the confidence trickster might rely on your goodwill without trusting

¹⁷ Annette Baier, 'Trust and Antitrust', *Moral Prejudices: Essays on Ethics* (Cambridge MA, Harvard University Press, 1996); p. 99.

you’).¹⁸ Holton is surely right that goodwill is not quite the fundamental distinction between trust and mere reliance, but I believe it is closely connected to what is. I suggest that the way to amend Baier’s proposal is to substitute the idea that trusting someone to do something is to rely on them to do it *competently and for the right sorts of reason*. One may trust an estranged partner to look after the children because one can rely on him to know how to take proper care of them, and to be moved to do so by the right kinds of reasons—for instance, that he loves them. The suggestion so far, then, is that trust is a matter of relying on someone to have certain shared values and the competence to act on them; the attitude of trust, then, is fundamentally directed towards a person’s character.

This conception of the essential difference between trust and mere reliance is better suited than is Baier’s goodwill account to accommodate the fact that trust can apply not only in cases where there is mutual awareness of the trust invested so that a deal of trust is made between the two parties, but also in cases where there is neither. We should introduce a distinction between a *second personal stance* of trust and a *third personal stance* of trust. I can ask a neighbour to look after the cat while I am away on holiday, and given she agrees, I may trust her to do it. Granted our mutual awareness of my trust (an essential feature of the I-thou, second personal stance), the possibility is created of my being betrayed by her, for instance in the event that she lets the cat go hungry. But we can also trust someone third personally, typically without their awareness. Imagine instead that I know that a friend of mine has fixed his own cat-sitting arrangement with the neighbour. Although I have nothing to do with the arrangement, as it happens I do trust the neighbour to take good care of my friend’s cat—if anyone were to ask me whether I trusted her to do it, I’d say Yes. If however I learn that she has betrayed his (second personal) trust by letting his cat go hungry, then my own (third personal) trust in her is also let down, but not in the manner of betrayal, for there was no deal of trust struck *with me*. There is a difference of perspective here that wants philosophical capture, and what is needed is a third category in between mere disappointed reliance and full-on betrayal; we need a category of ethical let-down that falls short of betrayal because it lacks the I-thou

¹⁸ Richard Holton, ‘Deciding to Trust, Coming to Believe’, *Australasian Journal of Philosophy* 72 (1994), 63-76; both quotations from p. 2.

perspective. Let us say that second personal trust creates the possibility of betrayal; third personal trust creates the possibility of ethical let-down; and the ethical attitudes involved in both kinds of trust are still to be contrasted with mere reliance, whose disappointment incorporates no ethical dimension.

What makes the crucial difference between third personal let-down on the one hand and second personal betrayal on the other is that it is only in the second personal case that a deal of trust is struck, so that salient among the reasons on which I trust the neighbour to act is the reason *that I so trust her*.¹⁹ By contrast, this reason has no application in the third personal case, for my trust is simply not relevant to the commitment the neighbour makes to my friend over his cat. If I am on the right track here, then there are two closely inter-connected features of trust that generate the possibility of betrayal: the first is the I-thou interpersonality enshrined in the second-personal stance—the deal of trust was struck between the neighbour and *me*; and the second is that the nature of my side of this deal of trust was such that among the right reasons I was relying on her to be suitably motivated by was that I was trusting her.

Given this picture of the distinction between second and third personal trust, what, if anything, is to be said specifically about trust in the relation between an inquirer and his informant? Obviously, the informant is aware of the trust invested in her by the inquirer (his inquiry signals as much), and in offering her word in response to the inquirer's epistemic need, the informant supplies her side to an interpersonal deal of trust. This means that salient among the right reasons the inquirer is relying on her informant to be suitably moved by is the reason *that he is trusting her as to whether p*. In this, as others have observed, testifying is analogous to promising. So, the proposal we have arrived at concerning the special flavour of the relation between an inquirer and a good informant is this: when the inquirer asks a potential informant whether *p*, he signals his *trusting epistemic dependence* on her; and *in responding* so as to meet the inquirer's epistemic need, the informant signals her *acceptance of responsibility* to honour (and so not betray) the trust invested in her. Thus a deal of

¹⁹ This reflexive aspect of trust is observed by others; see Karen Jones 'Trust as an Affective Attitude', *Ethics* Vol. 107, No. 1 (October, 1996), 4-25. Paul Faulkner has developed the closely related idea of 'affective trust' (Faulkner 2007); see also Faulkner 'On Telling and Trusting', *Mind* 116 (464) 2007, 1-28.

testimonial trust is swiftly, implicitly, and perhaps very lightly struck. In effect, the informant thereby accepts responsibility for her inquirer's belief as to whether *p*; and this reveals how profoundly intersubjective the business of testimony is—not merely a matter of an informant's belief being transmitted to another party, but moreover of responsibility taken for the other person's epistemic status. If the informant is indeed a good informant, then she is moved by the right sorts of reasons when summoning her response. These right reasons will most generally be reasons of epistemic and ethical responsibility (broadly speaking, reasons of competence and sincerity), but salient among these right reasons is the reason that her inquirer has placed trust in her as to whether *p* as his side of their deal of trust. This latter reason is perhaps the most powerful when it comes to generating the appropriateness of a sense of personal betrayal on the part of the inquirer in the event that his informant should prove less-than-good—less than worthy of the trust he had invested in her—for it is the most immediately personally focussed reason among them.

Transposing our trust based conception of the relations between a good informant and an inquirer to the relations involved in testimony generally (principally, thereby to include cases of less-than-good informants, good-but-non-recognizable informants, and cases where the epistemic needs on the part of the addressee are very minimal and/or inexplicit), we might say that in giving a response to an expressed or implied need for information, the testifier *commits to trustworthiness as to whether p* saliently because of the expressed or implied trust on the part of her interlocutor; and, where a testifier's word is less than trustworthy, she thereby *reneges on that commitment*. The trust based picture of good informing from which I have generalized this trust based conception of the speech act of testimony has effectively led us to a version of the assurance view of testimony. Before I go on to take our trust based conception as an authoritative model for what collective testimony would have to be like, then, I shall need to attend more directly to the assurance view and defend it against a line of attack that has been mounted against the epistemological picture it is understood to incorporate. In this we shall see the continuing relevance of the difference between a good informant and a source of information as I have developed that distinction in terms of second and third personal stances of trust, and the different grades of potential ethical disappointment that these different stances entail.

Testimony As Assurance

Assurance views have been advanced in slightly different terms by different authors but I will concentrate on what unifies them, and by reference principally to papers by Richard Moran and Edward Hinchman.²⁰ Moran motivates the assurance view partly by pointing out that the idea of believing a *person* is crucial to testimony: ‘Somewhat lost in much recent discussion...is attention to the basic *relationship* between people when one person tells a second person something, and the second person believes him’.²¹ It is worth noting that one can also find the seeds of an assurance view in Coady (worth noting especially since in other respects, particularly in casting testimony as a species of evidence, his view can appear anathema to the assurance approach), for Coady anticipates the analogy between testifying and promising that Moran makes a cornerstone of his view:

We might say that both [promising and testifying] are cases in which a speaker vouches for *p* but in different ways. By testifying that *p* one vouches for an ascertained connection between how one says things are and how they are actually. By promising that *p* one vouches for bringing about such a connection by changing how things are.²²

A crucial aspect of the neglect of the interpersonal relationship between people in testimonial exchange is revealed in the fact that many accounts of the epistemology of testimony do not in fact describe anything specific to telling at all. Moran observes that arguments which aim (he has Burge and Coady in mind) to establish an a priori entitlement to believe what others tell us are not special to testimony as such, or even speech more generally. They are altogether generic, creating a presumptive entitlement to believe *others’ beliefs*, however transmitted, if

²⁰ Edward Hinchman, ‘Telling as Inviting to Trust’, *Philosophy and Phenomenological Research* Vol. LXX, No. 3, May 2005, 562-87; Richard Moran, ‘Getting Told and Being Believed’, in J. Lackey and E. Sosa eds. *The Epistemology of Testimony* (Oxford: Oxford University Press, 2006).

²¹ Moran, 2006, p. 273.

²² Coady, 1992, p. 43.

at all. And Hinchman makes the point that such arguments (he has Burge and Brandom in the frame) do not distinguish the entitlement that is made available to the hearer who recognizes his interlocutor's intention to tell him that *p* from the entitlement that may be made available to anyone else who happens to overhear.²³ By contrast, what advocates of any assurance view will emphasize is that testimony is to be delineated by reference to a certain illocutionary speech act—telling that *p*—and that (in Moran's terms) this speech act is a kind of guarantee or 'assurance that *p*' which renders the testifier accountable to his interlocutor for the truth of *p*; or (in Hinchman's emphasis) the speech act is such that its success depends upon the addressee recognizing the testifier's intention to perform just that speech act.

Moreover, this recognition generates a distinctive kind of reason, available only to the addressee, to believe what she is thereby told: in Hinchman's terms, the reason is a *prima facie* entitlement, which stops being merely *prima facie* if there is 'reasonable trust' on the part of the addressee; in Moran's terms the reason is cast as one which is 'seen by both parties as dependent on the speaker's making himself accountable' for the addressee's belief,²⁴ where the question whether this reason counts as a good or sufficient reason for belief depends upon whether the testifier has exercised epistemic responsibility and spoken sincerely. Both writers talk of testimony as a matter of an invitation to trust, which naturally confers what Moran describes as a 'right of complaint'²⁵ on the addressee in the event that the testifier should prove less than trustworthy.

²³ See Moran 2006, p. 274-76, and Hinchman 2005, p. 575-6. Hinchman also objects that Burge and Brandom's thesis that an entitlement to believe is generated by the mere fact of intelligibility is too strong, because the hearer might not reasonably trust the testifier and so will not have an entitlement to accept what she says. But I think Hinchman is misrepresenting Burge on this score, for Burge is careful to talk only of *default* entitlement here. In cases where the hearer has reason not to trust the testifier, the default lapses, and no entitlement is possessed. As I read Hinchman and Burge, their entitlement views of the epistemology are exactly analogous, though for Burge the entitlement is (so to speak) bottom-up, deriving from maximally basic and general features of intelligibility, whereas for Hinchman it is top-down, deriving from features specific to the speech act of testimony.

²⁴ Moran, 2006, p. 295.

²⁵ Moran, 2006, p. 295.

Most centrally intriguing, however, is that both writers emphasize the *non-evidential* nature of the addressee's reason to believe what she is told. Moran says 'The hope is to show that the paradigmatic situations of telling cannot be thought of as the presentation or acceptance of evidence at all...'.²⁶ And Hinchman explains the non-evidential nature of testimonial reasons as follows:

When you have evidence of a speaker's reliability you don't *need* to trust her: you can treat her speech act as a mere assertion and believe what she says on the basis of the evidence you have of its truth. You can ignore the fact that she's addressing you, inviting you. You can treat her as a truth-gauge.²⁷

This non-evidential aspect of the assurance account's epistemology of testimony has prompted a trenchant line of criticism from Jennifer Lackey and Frederick Schmitt. Both writers raise as a principal objection that an assurance or invitation to trust cannot in and of itself produce an *epistemic* reason to believe what one is told, but at best a practical or ethical reason. I find their line of criticism helpful in homing in on an obscurity in the assurance view as it stands, but an obscurity I aim to clarify by reference to the distinction between second and third personal trust. To anticipate, I shall aim to vindicate the claim that a testifier's word is not (in the normal case) *treated as evidence* by the addressee, so long as we understand that claim in a certain way, namely by reference to the distinction between a good informant and a source of information.

Schmitt's critical discussion is sensitive to the fact that the assurance view as presented by Moran may receive a range of slightly different readings, and the objection I am focussing on relates to what Schmitt calls the 'actual agreement' variant. First summarizing the central tenet of the actual agreement variant, Schmitt puts his objection like this:

²⁶ Moran, 2006, p. 274.

²⁷ Hinchman, 2005, p. 580.

By accepting my invitation to trust me, you commit yourself to trusting me and you have a reason to do so in virtue of this commitment...

The trouble with this argument is that we cannot infer from the fact that your accepting my invitation to trust me gives you a reason to trust me, that it also gives you an *epistemic* reason to believe me that *p*. The reason your accepting my invitation gives you is not itself an epistemic reason to believe *p*; and there is no basis for thinking that giving you a reason that is not itself epistemic could give you an epistemic reason.²⁸

Certainly, if this were what Moran's assurance view turned on, it would indeed be hard to see how it might produce an epistemic reason for the addressee to believe what he is told. But I do not think Moran is committed to the idea that the epistemic reason to believe is generated by any commitment on the part of the addressee; rather the epistemic reason is generated by the commitment made by the *testifier*. The picture, as I see it, is rather this: in offering her word that *p*, the testifier invites the addressee to trust her that *p*, and so incurs a commitment to trustworthiness regarding *p*. Now, provided there is sufficient background reason for the addressee to assume the testifier has exercised epistemic responsibility and is speaking sincerely, then *the testifier's (ethical-epistemic) reason to be trustworthy regarding p generates a correlative (purely epistemic) reason for the addressee to take her at her word*. For example, if at a dental check up my dentist tells me 'You have a wisdom tooth coming through', the fact that he tells me this (invites me to trust him on it), *given* there is sufficient background reason to regard him as trustworthy (on balance reason to trust both his diagnostic competence and his sincerity towards me as a patient), then his commitment to the proffered diagnosis provides me with a further epistemic reason to believe him (to accept his invitation to trust his word). In trusting him, I seal the deal of trust between us, and in this case it is a good deal—I am not epistemically cheated. That my trust is reasonable allows the distinctive assurance based or invitation-to-trust based reason to apply, so that I have a trust based reason for accepting his word that I have a wisdom tooth coming through. Contra Schmitt, then,

²⁸ Frederick Schmitt, 'The Assurance View of Testimony', in A. Haddock, A. Millar, D. Pritchard eds. *Social Epistemology* (Oxford: Oxford University Press, 2010), p. 227.

what the assurance supplies in such circumstances of trust *is* an epistemic reason—it is a reason that bears on the likelihood that what he has told me is true.

Let me now consider Lackey's way of levelling the general objection that assurance provides no epistemic reason to believe what one is assured of. Classifying the assurance view as an interpersonal view of testimony [IVT], she puts the challenge in terms of a dilemma:

The proponent of the IVT faces a dilemma: either the view of testimony in question is genuinely interpersonal but epistemologically impotent, or it is not epistemologically impotent but neither is it genuinely interpersonal. Either way, the IVT fails to provide a compelling alternative to existing theories in the epistemology of testimony.²⁹

The first leg of the charge is essentially that if we take the pure assurance in isolation, without the background condition of sufficient reason to assume competence and sincerity, then we quickly see that it is simply not up to the job of providing an epistemic reason for acceptance, for one might be assured of things by a testifier who is not in the least trustworthy, so that their assurance gives one zero reason to accept their word. The thought is this shows that the fact of assurance in and of itself has no epistemic value. The second leg of the charge is that if the assurance theorist should respond by pointing out that it is only conditionally upon a background of sufficient reason to trust that the assurance related reason is generated, then this reveals that the assurance element is epistemically superfluous, for it is the reasonable trust that is doing all the epistemic work. Now given that both Moran and Hinchman do indeed argue (in different ways³⁰) that the trust based reason only justifies acceptance

²⁹ Lackey, 2008; p. 222.

³⁰ Hinchman casts reasonable trust as a matter of the trusted party being worthy of the trust, and there being no evidence available that it is untrustworthy (Hinchman, 2005; p. 578). Moran talks chiefly of the testifier's having discharged epistemic responsibility and being sincere (Moran, 2006, p. 295). Paul Faulkner gives an independent account in terms of norms of trust; see Faulkner, 'Norms of Trust', in A. Haddock, A. Millar, and D. Pritchard eds. *Social Epistemology* (Oxford: Oxford University Press, 2010). For my own part, I see the assurance view as compatible with the account of the epistemology of testimony I have advanced elsewhere, such that

conditionally upon the trust being reasonable, the first of these objections simply does not hit home. Indeed Lackey levels it only at Moran, and only in the first instance, subsequently acknowledging that he does in fact state that assurance only provides a reason to trust given the speaker has exercised epistemic responsibility and sincerity. So it is really only the second objection that needs addressing, namely the objection that, precisely *because of* the requirement that there be a background of reasonable trust, the assurance based reason is revealed as justificationaly superfluous.

One appreciates the lucidity of the objection: where there is already a background of sufficient reason to assume competence and sincerity, there *is* no more epistemic work to be done in terms of meeting the threshold for warranted acceptance—the addressee is already epistemically permitted to accept the word of the trusted testifier. I think this is true. But that does not mean there is no further source of warrant to be identified and explained, and we have already seen that if we look and see what goes on in the relations of trust involve in the speech act of telling, what we find is that there is indeed a further epistemic reason generated. The proper purpose of the assurance view is to reveal the reason-giving interpersonal aspect of the speech act of telling, and to characterize the distinctive *kind* of epistemic reason that it can generate—albeit superfluously from the narrow point of view of sheer permission to accept. The assurance view is simply not structured to quell the epistemologist’s traditional anxiety that we might not be warranted in accepting that *p* from someone’s say-so. For where we are not already warranted by the various background considerations, receiving assurances from the very speaker whose competence and/or sincerity are in the balance is simply not going to help. (Indeed assurances *per se* would normally have a negative effect: imagine a dentist who starts up his drill and pauses to assure you ‘I really am a qualified dentist, you know’.) In short I think any assurance account should be clear that it does not aspire to play any part in a story that aims to calm the traditional epistemological anxiety. Rather, its contribution lies in its ability to reveal that in the normal spontaneous case of testimony, the addressee does not treat the speaker’s word that *p* as evidence for the

the virtuous hearer exercises a well-trained testimonial sensibility which allows her reliably to pick up on positive and negative cues relating to the speaker’s trustworthiness (see *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford: Oxford University Press, 2007), ch. 3).

fact that *p*. Somebody else might so treat it (that somebody told someone that *p* is just sort of thing that *can* properly be treated as evidence), but it is characteristic of the addressee of testimony that she does not.

Let me explain this idea further by way of Lackey's own discussion of the case of an eavesdropper. Her purpose being to illustrate the epistemic insignificance of assurance, Lackey compares the epistemic status of an addressee with that of an eavesdropper, and proposes convincingly that, other things equal, the addressee is epistemically no better off than the eavesdropper. Thus the example of Earl who eavesdrops on Ben telling Kate that their boss is having an affair with an intern called Irene. Lackey rightly presses the question:

[I]f Kate and Earl are equally properly functioning as recipients of testimony, have the same relevant background information, both about Ben as a testifier and about the proposition to which he is testifying, and form their beliefs about the boss and Irene solely on the basis of Ben's testimony, then what could distinguish their beliefs epistemically?³¹

This question deserves an answer; and the beginnings of the answer is that Kate and Earl may well have the same amount of warrant, but since Kate was told that *p* while Earl was not told it but rather acquired the belief that *p* from the discursive environment on his own initiative, their warrant is generated by different sorts of epistemic reasons to believe what Ben has told Kate. This difference, as Lackey goes on to anticipate, shows up primarily in the fact that Kate (but not eavesdropping Earl) has a right to resent Ben for misleading her, while Earl has no such right. Now Lackey is unimpressed with this response as a way of explaining the distinctiveness of Kate's reason to believe, arguing that an eavesdropper might properly feel just as much resentment as the addressee when it comes to a testifier whose word turned out to have been untrustworthy—for instance if Earl believed what he overheard Ben tell Kate, only later to discover that Ben was lying. Lackey is pressing an important point, and while I believe there is ultimately a satisfactory distinction to be drawn here between the different kinds of resentment that are proper to the two cases, we clearly

³¹ Lackey, 2008; p. 234.

need to avail ourselves of further resources in order to draw it convincingly. What exactly is the difference between the resentment that Earl might appropriately feel and that which Kate is entitled to feel; and how does this distinction of resentments correlate with the different kinds of warrant that they possess?

To proceed with this question we must draw on our amplified version of the distinction between a good informant and a source of information, amplified, that is, by our understanding of the different stances of trust that belong to each case. Recall the earlier example of my second personal trust in a neighbour to take care of my cat, as contrasted with my third personal trust in a neighbour to look after a friend's cat, where I am an outsider to their arrangement. In the first case, when she fails to take proper care of my cat, I may appropriately feel the interpersonal sting of betrayal ('How could you? *You told me* you'd look after Mimi!'). By contrast, in the second case, when I learn that she has failed to take care of my friend's cat, my trust is disappointed so that I may well feel ethically let down (and she may well sink in my esteem), but this is not the kind of ethical let-down that belongs to betrayal. What makes the difference here is the difference between the second and third personal stances of trust. Both kinds of trust are real trust (not mere reliance) as is revealed in the fact that both can be culpably breached and one or another kind of ethical resentment properly felt. But only second personal trust is played out in the key of betrayal.

Now let us apply this to Earl *et al.* What Earl does is draw informational value from a state of affairs in which Ben tells Kate that *p*. Since Ben does not tell him anything, Earl does not relate to Ben as being any kind of informant and accordingly does not stand in the relevant second personal relation of trust to him. But in this state of affairs, Ben's role is more active than it would be if his informational role were merely a matter of passively manifesting, for instance, a sun burned nose. We might crystallize this active/passive distinction by saying that Ben is functioning as an *active source of information*, rather than a *passive source of information*: he is exercising agency in a way that, deliberately or not, will generate informational value for an appropriately equipped observer. To be clear, an active source of information need not have the idea of an observer in mind—on the contrary, he may be quite unaware that anyone is observing and gleaning information. Thus Ben tells Kate that

their boss is having an affair, thereby functioning for her as an informant, and Earl is a secret observer naturally equipped to harvest the informational value that Ben, through an exercise of agency, has sown in the state of affairs. Ben is unwittingly functioning for Earl as an active source of information: a special kind of evidence which renders appropriate a sense of trust-let-down in the event of failure, but which is a species of evidence all the same.

What this means is that while *both* Kate and Earl may indeed be said to trust Ben, only Kate has that distinctive second personal stance of trust towards Ben that renders appropriate the sting of betrayal if it transpires that he was deceiving her. Earl may be epistemically and ethically let down, but the merely third personal stance of trust he has towards Ben permits no genuine sense of betrayal.³² Thus I can agree with Lackey that, other things equal, Kate and Earl may feel equal quantities of resentment against Ben if it turns out he culpably misled them, but they are not the same sorts of resentment. They derive from different stances of trust, which are associated with relationships that generate subtly different kinds of epistemic reasons to believe: the second personal stance of trust that Kate has towards Ben is her side of a two-sided deal of trust with Ben which can generate a non-evidential epistemic reason for her to believe what he tells her. It supplies a reason for her to defer to him, not in the thin sense of merely setting a relatively high probability on the correctness of what he claims, but in the thick, ethically loaded sense of letting him take epistemic responsibility for her belief so that she may typically *hold him to that responsibility by holding him to his word*. This thick sense of deferral, which involves this ethically demanding acquiescence in an asymmetrical structure of epistemic authority, is only available in relationships of second personal trust.

By contrast, the third personal stance of trust that Earl has towards Ben is more simply probabilistic. It is an attitude of trusting reliance on Ben's testimony

³² The same points apply to the case of overheard trustworthy soliloquy (see Lackey 2008, p. 224). Such a case is fully explained on my version of the assurance view, according to which when someone learns that *p* by overhearing it said to no one, they do so by drawing out the informational value that the speaker has actively injected into the state of affairs. The speaker thus functions as an active source of information for the overhearer; which means that in the case where the soliloquy proves false, the overhearer's third personal trust will have been let down.

considered as an evidential feature of the observed state of affairs that has informational value for Earl (Earl is seeing/hearing for himself evidence that bears on the question of the boss's affair). If Earl defers, it is only in the thin sense available within a relationship of third personal trust. His relationship to the testifier does not permit him to take up the demanding attitude of holding him to the natural responsibility he has assumed by proffering his word. Compare a slightly different example. Imagine a visitor to an art gallery who finds herself perambulating within earshot of a tour group so that she may earwig the expert testimony of the guide as he conveys his group from picture to picture with his highly informed art historical flow.³³ Obviously this is a situation of high third personal trust (one is listening in on expert testimony) and so the listener recognizes the social cues that indicate she is in an asymmetrical structure of epistemic authority, and duly gives great weight to the tour guide's words. So there is certainly deference here of the thin, broadly probabilistic, evidential variety; but there is no handing over localised epistemic responsibility for her beliefs about the pictures, no *holding* the tour guide to a responsibility assumed, for the guide never assumed such responsibility vis-à-vis *her* beliefs; only vis-à-vis those of his group. In short, there is no deal of second personal trust between guide and earwigger, no I-thou relationship to make for thick deference and its partner possibility of betrayal.

In such a case, the tour guide is functioning as a highly accurate active truth gauge in relation to the art, rather as he might be if his words were merely an unattributed recorded voice in the headset, or an anonymous accompanying booklet. Our earwigger defers to him as an active source of information, harvesting informational value from his words. I fully acknowledge that this difference between what I am calling thin and thick deference cannot be entirely clear-cut, and there will no doubt be borderline cases—at the very least it will sometimes, for instance, be unclear who, or how many, are being addressed in a given case of telling (again, compare promising). I introduce the thin/thick deference distinction in order to further amplify the distinction I have already made between the two stances of trust—second as opposed to merely third personal—and it is this difference of trusting stances that I contend makes sense of Moran and Hinchman's claim that the relationship between a

³³ I thank Marie McGinn for this example.

reasonably trusted testifier and an addressee supplies the addressee with a ‘non-evidential’ epistemic reason to believe what she is told. Such a reason is the product of thick deference and its associated attitude of holding the testifier to their word. None of this is going on when people treat others’ utterances as evidence.

We should note, however, that the addressee of a piece of testimony is not *compelled* to defer in the manner I describe as thick, even if she ends up accepting what the speaker tells her. She may reject the invitation to second personal trust, or at any rate not accept it—so that she does not hand over responsibility for her belief on the matter. She may instead reflect on the evidence at large and come to accept (or not) what is said by way of that route. Imagine that when Ben tells Kate their boss is having an affair, Kate demurs at first because she is aware that Ben is an unscrupulously inventive gossip, but then she considers that on balance Ben wouldn’t gossip irresponsibly about Irene, because he really cares for her and is primarily motivated by concern for her wellbeing. Kate believes Ben; she on balance trusts Ben’s word, but now more in the third personal sense, and despite being the addressee of his testimony. She is now more like eavesdropping Earl: they both receive and assess Ben’s word as part of a body of evidence. It is a proper, indeed essential, feature of a rational testimonial practice that addressees are capable of moving in and out of second and third personal relations of trust (and distrust) in this way—moving, that is, from a perspective of *holding* to one of *evidential assessment*—so that they can be appropriately responsive to incident cues for doubt. The very reasonability of second personal trust depends on our counterfactual responsiveness to indicators that we cannot, after all, reasonably take our interlocutor at his word, and should now take a step back to assess matters in evidential mode. When we do this, we do treat the teller’s word as evidence; but prior to this stepping back, we did not—rather we deferred in the thick sense of handing over localised epistemic responsibility for our belief as to whether *p*.

The picture of the epistemology of assurance that has emerged from the discussion overall is that an addressee and an eavesdropper could have exactly the same background reasons to trust a testifier’s word; but that in telling his addressee that *p* a testifier offers her a second personal trust based reason to believe his word that *p*; whereas he (wittingly or unwittingly) offers the eavesdropper a third personal

trust based reason to believe that p . These are both epistemic reasons, for they both bear on the likely truth of p , and they may deliver the same strength of warrant; but I have tried to vindicate the idea that they are subtly different sorts of epistemic reason, deriving from the two different stances of trust. My proposal has been that this distinction between second and third personal trust is what, at root, explains and clarifies Moran and Hinchman's defining idea that an assurance can generate an epistemic reason to believe that is not evidential in kind: an assurance, or invitation to trust, is essentially interpersonal and so inspires the distinctive *second personal* stance of trust that can generate a non-evidential epistemic reason to believe. The assurance view reveals that the testifier's very commitment to truthfulness is shored up by the implicit threat of a powerful kind of blame ('You betrayed me!'). This is one of the many ways that the addressee's holding the testifier to the trust invested is intertwined with the relevant epistemic relations: in holding testifiers to the truthfulness of their word in the manner of second personal trust, addressees far and wide massively boost the veritistic energies that drive a well-functioning testimonial practice. Just as sanctions of blame and other forms of resentment help us sustain good moral motives, so do they shore up trustworthy testimony.

Having made my case of support in favour of the main elements of an assurance view of testimony according to which a testifier makes a deal of second personal trust with her addressee, we are at last in a position to ask our question: Is a group capable of doing *that*? Can a group be so constituted that it is capable of standing in the distinctively second personal relations of trust that are proper to testimony?

Can Groups Be Testifiers?

One way to approach this question would be to start by considering summative and non-summative models of groups. One sort of summative view would be that a group can only testify that p if a majority of its members so testify; or that a group can only

testify if an individual spokesperson testifies on the group's behalf.³⁴ There is, as ever, a range of possibilities here. But my strategy is not to proceed from the summative/non-summative problematic, to argue for the existence of genuinely group intention, belief, knowledge, and so on, building up to group testimony from there. Rather I want to take for granted the existence of these collective states and capacities, on which there is a large and long-established literature, and to proceed instead from the view we have formed of the exigencies of the speech act of testimony, so that we can see what species of group would be up to the task. I have been arguing in support of a view of testimony according to which the speech act centrally involves the testifier making a commitment of trustworthiness to the addressee (worthiness, that is, of his second personal trust) with respect to *p*. Now groups can be constituted in various ways, and at least one kind of group agency can surely be constituted in a relatively non-committal manner, so that it might involve individuals integrating their plans *pro tem* either with an individual intention motivating them, or a more distinctly group intention motivating them, and where in many cases these individuals may quit the arrangement without incurring appropriate rebuke.³⁵ I do not doubt that these are real phenomena of collective activity, for I believe there are many grades of collectivity, but there is a founding view in the literature which posits commitment as the mainspring of what it takes to constitute a certain robustly and stably collective agent: I have in mind Margaret Gilbert's account of 'plural subjects'.

On Gilbert's account, a plural subject is formed by each member of a group expressing (perhaps tacitly) willingness, under conditions of common knowledge, to X-ing as a body. Such are the conditions of joint commitment: 'A and B (and...)... constitute a *plural subject* (by definition) if and only if they are jointly committed to

³⁴ Jennifer Lackey argues for the latter—a 'deflationary account', which she describes as reductionist but non-summative ('A Deflationary Account of Group Testimony', forthcoming).

³⁵ See, for instance, Michael Bratman, *Faces of Intention: Selected Essays on Intention and Agency* (Cambridge: Cambridge University Press, 1999), Part Two. And Raimo Tuomela identifies different gradations of joint intention in the 'I-mode' and in the 'we-mode'. See, for instance, 'Collective and Joint Intention', *Mind and Society* (2000) 1 (2), 39-69; or *The Philosophy of Sociality: The Shared Point of View* (Oxford: Oxford University Press, 2007).

doing something as a body...'.³⁶ It is a central feature of Gilbert's account that the *commitment* in joint commitment shows up in the fact that if a member fails to honour it, for instance by simply ceasing the joint activity, she will be appropriately subject to some level of rebuke. (In such cases the commitment is reneged upon and is not dissolved. This is to be contrasted with various ways in which members may free themselves from a joint commitment, properly terminating their part in it.³⁷) Taking my cues from the features of a good informant, I have argued that, paradigmatically, a testifier's word is a response to epistemic need regarding *p* on the part of the addressee, such that the testifier thereby *commits to second personal trustworthiness concerning whether p*, according to which saliently among the right reasons governing her conduct will be the fact that the addressee is trusting her on it—for *that*, most specifically, is the trust of which the testifier is committing to being worthy. This second personal relation of trust means that where a testifier fails to be trustworthy, she reneges on a commitment—her side of the deal of trust—so that the addressee is thereby betrayed. Now the kind of commitment at stake in the formation of a plural subject is just the right kind of commitment to generate the possibility of such a betrayal. Gilbert says:

I suggest that a joint commitment is the clearest possible context for interpersonal betrayal. If Deb fails to conform to a joint commitment she is party to with George she has—to some extent—betrayed him. She would betray many people at once should she fail to conform to a joint commitment she is party to with many others.³⁸

We can see, then, that the particular assurance account of the speech act of testimony that I have been promoting involves the same materials that operate in constituting plural subjectivity. With this conception of a group or collective, we are clearly in the right register.

³⁶ Margaret Gilbert, *A Theory of Political Obligation* (Oxford: Clarendon Press, 2006); p. 144-5. See also n. 3.

³⁷ Gilbert, 2006; pp. 141-144.

³⁸ Gilbert, 2006; p. 150.

Let us look a bit more closely, then, so that we can see how it might work for such an entity to testify. In the first instance, we must take note that the intersubjective commitments that make up joint commitment are all between the individuals in the group; they are not, in the first instance, commitments on the part of the group to another party, for example, an inquirer or addressee of testimony. But still, a collective that is constituted by way of joint commitment *is* therein well suited to making commitments to parties outside the group, for whatever intersubjective commitments are at work in constituting a given plural subject—ordinarily it might be some complex of intentions and/or beliefs, for instance—*commitments* are just the sort of group-making relations we need in order to guarantee the stability of group-level attitudes that any group commitment to an outside party will require. (Stability of group belief is an obvious case in point—Deborah Tollefsen has emphasized the importance of stable group belief precisely in connection with group testimony.³⁹) My proposal is that we construe a group testifier as constituted, at least in part, by way of a *joint commitment to trustworthiness as to whether p* (or whatever range of *p*-like questions might delineate the body’s expertise, formal remit, or informal range of responsibility). Indeed, it simply seems to me to be true that many collective bodies—committees, news teams, governments, research groups, consultancies—are partly constituted by way of such a commitment to second personal epistemic trustworthiness. Indeed I suspect that any collective body whose role is likely to involve telling people things would need to have this as one of its constitutive commitments, because precisely that sort of epistemic trustworthiness would be integral to the group’s role. At any rate, my main claim can be the weaker one: that any group partly constituted by way of a joint commitment to trustworthiness (regarding some relevant range of questions) is pre-eminently suited to enter into the second personal relations of trust that characterize testimony.

A further difference between individual and group testimony that we should mention before considering an example is that the testimony of collectives, especially if they are institutional bodies, will often have a plural or general addressee rather

³⁹ See Deborah Tollefsen, ‘Groups as Rational Sources’, in H. Schmid, D. Sirtes, and M. Weber eds. *Collective Epistemology* (Springer, 2011).

than an individual one. The addressee might be ‘the public’, for instance, or ‘service users’, or ‘anyone who wants to know’. But this presents no difficulty for our model. Let’s consider an example in which the testifier is a group, and the addressee too. Imagine a committee appointed by the government to inform the public about the health risks related to certain food additives. Its members jointly commit to trustworthiness, at least tacitly (it comes with the job) regarding what they will tell the public by means of the report they are charged with producing. A member of the public may hear or read the information that p , grasping that he is indeed being told that p by the committee that authored the report. (He recognizes the committee’s intention to tell the public what is asserted in the report.) Here we have a clear example of a collective body testifying that p to a member of the public, by having offered its testimony to the public in general. In offering its testimony, the committee invites the public to trust its word on the matter, and in grasping this (whether he happens to accept the committee’s word or not) the citizen plays his reciprocating part in the group’s successful illocutionary act of testifying. If he does trust, then a deal of trust is sealed. If that trust is reasonable, then the committee’s word gives him a non-evidential epistemic reason to believe, and this reason supports acceptance in concert with the evidential reasons that are supplying the reasonable trust.

So far so good. Now let us imagine that six months later it transpires that there was corruption in the committee—that some members had accepted lavish hospitality from one of the companies whose products were under scrutiny and subsequently softened their line on the risks associated with those products. Our citizen is betrayed. He trusted the committee’s word, but they abused his trust, and doubtless that of many other members of the public. All the defining trust related features of testimony are present in this case, and the kind of group that is generated on the joint commitment account is of course robustly collective (non-summative). I propose our example provides us with a paradigm of the fullest kind of group testimony. What we see in the idea of a joint commitment to second personal epistemic trustworthiness is a collective version of the good informant. The formation of a collective body whose role is to be seen to stand in the second personal relations of trust requisite for testimony *is* the making of a collective good informant.

Good informants are of course to be contrasted with sources of information; and earlier I introduced a distinction between active and passive sources of information in order to make sense of the ethically lighter, or at least less direct, kind of trust (third personal trust) that surely holds, and can be let down, in cases where we rely on another's active informational input into a state of affairs. Among individuals we considered an example of eavesdropping. The eavesdropper learns from words but in a way that does not involve his being told; he trusts the speaker but his stance of trust is only third personal, so that the kind of ethical let down he may experience falls short of betrayal. The eavesdropper has, so to speak, judged for himself on the basis of the evidence available to him, evidence which includes someone's testimony to someone else.

Let us imagine a correlative example, in which the active source of information is a collective or institutional body. We want a case where members of the public rely on an institutional body to provide information in a reliable way but where nothing is taken on second personal trust. Consider, for instance, a news consultancy that simply collates information from newspapers and online news. It does not itself tell anyone anything, but functions as a conduit for news reports, with sources duly cited. If it is a well-established organization, we may well trust it; that is, our stance towards it might be that of third personal trust. Such an organization is an active source of information—it exercises group agency in order to inject informational value into states of affairs that can be observed by others for purposes of harvesting that informational value. (Perhaps it gathers the information on a web site, and we access it there.) Now, what if some of that information proves false, or culpably misleading? It depends. Given what I have already said about group testimony, then if the originator of the news story is at fault, and provided the originator is properly construed as having testified to us, then we are betrayed by the originator. But if the irresponsibility is on the part of our trusted news gathering organization—perhaps it carelessly messed up all the dates so that it put some of last week's news into this week's summary—then we are let down, though not betrayed. We are let down by a collective active source of information that we had trusted (third personally) to take more epistemic care. We often depend on collectives' epistemic activity in just this way, and it is important not to conflate it with the kind of thick deferral at stake in group testimony.

In a discussion of *Wikipedia*'s status as a potential testifier, Deborah Tollefsen explores the interesting possibility that *Wikipedia* could gradually evolve into the sort of entity that is stable enough and self-regulating enough to qualify as a group testifier.⁴⁰ On her own account of group testimony, this would involve giving an assurance, though she travels lighter than I am willing to when it comes to making out how giving an assurance is something a group can do. On my scheme of things it would take a good while longer before *Wikipedia* might develop into a sufficiently committal form to be capable of standing in the requisite relations of second personal trust. Consequently, although I agree with much of Tollefsen's interesting and instructive argument, my own take on *Wikipedia* is rather that in its current state—and precisely in virtue of the increasing reliability and self-regulation that Tollefsen rightly emphasizes—it provides us with a nice example of a collective active source of information. Modern society has many institutional bodies that can function for us in this way—news media outfits, research groups, polling organizations, educational bodies, financial advice authorities, investigative committees. And, like individuals, many of these bodies might also be able to function for us as testifiers, depending crucially on whether the requisite joint commitment to epistemic trustworthiness is *pro tem* among the commitments that constitute the group. On my view, then, only when *Wiki* contributors can jointly so commit, will it qualify as a group testifier.

Do We Need Group Testimony?

My overall aim in this paper has been to explore how Craig's original distinction between a good informant and a source of information can be developed at the collective level, with a view to vindicating the idea of group testimony. With the distinctions between good informants, and passive and active sources of information in place, I hope also to have cast some light on the different forms of dependence we are bound to have on the good epistemic conduct—we might say, good epistemic

⁴⁰ Deborah Tollefsen, 'Wikipedia and the Epistemology of Testimony', *Episteme: Journal of Social Epistemology* (2009) Vol. 6 (1); 8-24. For her earlier paper on group testimony—to my knowledge the first in the literature—see 'Group Testimony', *Social Epistemology* (2007) 21 (3), 299-311.

character—of both individuals and collectives. Let me finish, however, with a brief reflection about the distinctive importance of institutional testimony to the democratic polity. Onora O’Neill has argued that (at least in the UK institutional context) there is reason to worry that relations of trust are being eroded or distorted by a seemingly unstoppable tide of managerial styles of institutional accountability, with its obligatory targets and bureaucratic new-speak. O’Neill gives a powerful critique of this upstart managerial regime in institutional life and notably of the colonization of educational institutions by the audit culture that imposes from the world of business a rigid idea of how institutional bodies should report on their activities. She laments the imposition of the managerial mode and calls for a more ‘intelligent accountability’:

Intelligent accountability, I suspect, requires more attention to good governance and fewer fantasies about total control. Good governance is possible only if institutions are allowed some margin for self-governance of a form appropriate to their particular tasks, within a framework of financial and other reporting.⁴¹

Intelligent accountability thus requires that institutional bodies be allowed to report on their activities to other relevant bodies—in short, to testify group to group. For an institution to render itself accountable in part by way of testimony is different from its being trusted third personally to supply information about targets met or unmet, or tick boxes about things done or not done. I interpret O’Neill’s conception of intelligent accountability as *inter alia* calling for the relations of second personal trust to be in place that are needed to sustain institutional testimony as an essential part of how public bodies may ‘intelligently’ render themselves accountable to each other, and ultimately to the public they serve. In general, one of the things citizens of a democracy should expect is that institutional bodies will sometimes metaphorically look them in the eye and tell them things—telling them who won the election would surely be a founding example. Lying about who won the election would be a collective betrayal of the first order. Without that possibility of institutional testimony, and the second personal relations of trust that are required for it, the

⁴¹ Onora O’Neill, *A Question of Trust – The BBC Reith Lectures 2002* (Cambridge: Cambridge University Press, 2002), p 58.

democratic ideal of how institutions make themselves accountable to the citizens they serve, and the collective understanding of what is at stake in institutional truthfulness, would be very much diminished.⁴²

⁴² I gave an earlier version of this paper at *Collective Epistemology: The Epistemic Life of Groups*, a conference at the Institute of Philosophy, School of Advanced Study, University of London (March 2011), at the Rutgers Epistemology Conference (May 2011), at a workshop at the University of Reading (July 2011), and a colloquium at Monash University, Melbourne. I am grateful to many people participating in these occasions for their helpful questions and comments, and particularly to Robert Audi, Jeremy Fantl, Paul Faulkner, Sandy Goldberg, Peter Graham, John Greco, Ted Hinchman, Karen Jones, Jennifer Lackey, Berislav Marusic, Ben McMyler, David Owens, Jim Pryor, David Sosa, and Deb Tollefsen.