

Trustworthy Groups and Organisations

Katherine Hawley, University of St Andrews

1. The Significance of Group Trustworthiness

We value trustworthiness where we find it, in friends, colleagues, and strangers. This is because it is easier to cooperate with trustworthy individuals than with untrustworthy ones, but also because trustworthiness is appreciated for its own sake. Conversely, we resent untrustworthiness when we encounter it, and try to minimise interaction with untrustworthy individuals. This preference for trustworthiness over untrustworthiness is not limited to our interactions with individuals. Who wouldn't prefer to work for a trustworthy company, to be treated at a trustworthy hospital, and to vote for a trustworthy political party?

The trustworthiness of collective entities is central to much public discourse around trust and distrust. A collapse of trust in bankers, say, is not merely a collapse of trust in individual, identifiable bankers, but takes in the profession as a whole; moreover, a collapse of trust in bankers goes hand-in-hand with a collapse of trust in the banks (de Bruin 2015). Likewise, discussion of trust or distrust in corporations, governments, news organisations, or even brands seems to concern collective entities of various kinds.

Unsurprisingly, then, trust and trustworthiness in collective contexts are widely discussed by social scientists. For example, in international relations there is debate about trust and distrust between states, other organisations or groups, and individual representatives of those collectives (e.g. Booth and Wheeler (2008)). And 'organizational trust' is a recognised research topic in management studies, taking in not only corporations, but also nongovernmental organisations and other complex group agents. Such research investigates both the influence of organisational contexts on our trust in individuals, and the ways in which trust is invested in organisations themselves (e.g. Saunders, Skinner, et al 2010).

If philosophers are to contribute to public discussions or to interdisciplinary dialogue, we need to understand how our own accounts of trust and trustworthiness can be applied in such contexts (Mäkelä and Townley 2013 make a related point). Moreover, an investigation of trust and trustworthiness in collective contexts promises to improve our understanding of

collective agency, intentionality, responsibility, belief, and knowledge; issues of trust span the theoretical-practical divide, often in distinctive and interesting ways.

But philosophers' accounts are typically focused on issues of trust in individuals, aiming to capture the nature of this attitude, the circumstances under which it may be rational, and the notion of trustworthiness which corresponds to this attitude of individual trust. (There are important exceptions to this individualistic rule; I return to some of these below.) In particular, philosophers often distinguish trusting someone from merely relying upon her, then theorise this distinction. Reliance is an attitude we may adopt towards inanimate objects: we rely upon the tent to keep us warm and dry overnight. Trust is usually taken to include reliance plus some further factor. This respects Annette Baier's insight (1986) that trust is distinctively connected with the possibility of betrayal and resentment, unlike 'mere' reliance: if the tent leaks, we should not resent the tent itself, though we might resent the person who promised to mend the tent before we set out. In (Hawley 2014) I argued that we should also distinguish distrust from mere absence of reliance, and moreover that we should not neglect the many interactions with other people in which either reliance or nonreliance is appropriate, but neither trust nor distrust is appropriate.

What is the magic ingredient which distinguishes (dis)trust from mere (non)reliance? Philosophers, inevitably, disagree: perhaps a truster imputes appropriate motives to the trustee, including perhaps a concern with or responsiveness to the needs, desires or indeed trust of the truster; perhaps a truster sees a trustee as morally obligated, committed, or accountable in appropriate ways (Simon (2013) is an excellent bibliographical guide). Across this variety of accounts, trust is understood to be directed at agents, or at persons, or in a distinctively second-personal fashion, or as an aspect of the participant stance: it is always a mistake to trust a tent, a trout, or a turkey, although we may rely upon such things. This emphasis on the interpersonal flows from the way in which the trust-reliance distinction is identified via the connection between trust and the reactive attitudes around betrayal. Not every interpersonal interaction is characterised by either trust or distrust, but every proper attitude of trust or distrust is directed interpersonally. Or so it seems.

Thus it is no accident that within this philosophical paradigm we rarely discuss trust or distrust as directed towards collective entities, since even the most inflationary accounts of collective agency hold back from treating such entities as full-fledged persons in every

respect, on a par with individual human persons. (To date, not even Mitt Romney has advocated corporate suffrage.) Given that we need a philosophical understanding of trust in collectives, how then should we proceed?

One strategy would be to consider various accounts of individual trust and trustworthiness pairwise with accounts of collective agency, responsibility, and so on, looking for rewarding and plausible combinations. This is a valuable project, which I do not pursue here. Another strategy would be to address collective issues on their own terms, without expecting to find much commonality between inter-personal trust and trust at the collective level. Baier (2013: 175-6), for example, is explicit that her account of interpersonal trust, in terms of vulnerability, competence, and goodwill, cannot easily encompass trust in organisations. We might hazard that accounts of trust which emphasise mutuality, emotional affect, and complex relationships will in general not apply to more distant forms of trust, including trust in collective entities; advocates of such accounts might see this as a strength, arguing that it would be a mistake to collapse such different attitudes.

In this paper, I adopt a different strategy, one which may complement either of the first two projects: I examine the costs of abandoning the trust-reliance distinction in collective contexts. Vindicating this distinction has been regarded as an essential criterion of success for accounts of trust in individuals. But, I will argue, we can explain and justify much of our practices around groups without using this distinction; the costs of abandoning it are low, as compared to the individual case.

Supposing my arguments are persuasive: what then? The fact that we can manage without this distinction does not entail that there is in reality no distinction between group trustworthiness and group reliability. For example, we might on closer investigation decide to adopt accounts of individual trust, and of collective belief, intention, and responsibility which together entail that groups can after all be trustworthy in the same way that individuals can be trustworthy; nothing I say in this paper will rule out that possibility.

However, if the costs of doing without the distinction are low, then we cannot use the supposed importance of the distinction as a consideration when deciding between different accounts of groups. Thus it is not a constraint on an adequate theory of groups that groups be capable of trustworthiness as opposed to mere reliability. A strong reading of my project is

that we should stop thinking in terms of either group trustworthiness or appropriate trust in groups. This might seem to spell disaster for the project of philosophical engagement with public and cross-disciplinary discussion of organisational trust. I reflect upon this challenge at the end of the paper, and gesture at a more optimistic way forward.

2. Groups, Organisations, Institutions

First, however, a note on terminology. Within philosophy, the term ‘group’ is a catch-all label for a great variety of candidate collective agents. For example, Frederick Schmitt’s ‘The Justification of Group Belief’ (1994) begins with examples such as ‘the Engineering Division of the Ford Motor Corporation’, ‘the crowd that had assembled on the square’, and ‘This Court’. Deborah Tollefsen’s ‘Group Deliberation...’ (2004) focuses on small teams of scientists. Kay Mathieson’s ‘The Epistemic Features of Group Belief’ (2006) focuses on groups identified by listing their members. In her ‘Group Testimony...’, Miranda Fricker mentions committees, news teams, governments, research groups, and consultancies (2009: 272). The cover blurb for List and Pettit’s *Group Agency* (2011) begins with companies, churches and states. All of these philosophers are sensitive to the differences between various types of collective entity, but use ‘group’ as the general term.

The situation is different within the social sciences. For example, Piotr Sztompka’s influential *Trust: A Sociological Theory* (1999: ch.3) distinguishes ‘social categories’, such as gender, age, and race, from ‘social groups’ such as the football club Real Madrid, a class of students, or an army platoon, and from ‘institutions and organizations’ such as the university, the army, the courts, and the banks. Sztompka goes on to discuss technological systems, food products, and general social systems such as democracy as objects of trust or distrust. It is hard for a novice to disentangle the relevant terminology, and harder still to map these distinctions onto those which typically interest philosophers. (I take consolation from economist Geoffrey Hodgson’s remark: ‘...endless disputes over the definitions of key terms such as *institution* and *organization* have led some writers to give up matters of definition and to propose getting down somehow to practical matters instead.’ (2006: 1))

I will stick with the philosophical usage of ‘group’, as a general term for all sorts of candidate collective agents. But we should bear in mind the variety of social entities which might be considered the object of trust or distrust, and the potential obstacles to communication across disciplinary divides, given this terminological choice.

3. Trust, Trustworthiness, and Reliability

So, do we need a distinction between trust and reliance with respect to groups? I will approach this via a related question: do we need to distinguish group trustworthiness from group reliability? Trust and trustworthiness are closely entwined, of course: very roughly, to be trustworthy is to merit the attitude of trust, and to trust is to regard as trustworthy. But it is useful to begin with trustworthiness, for several reasons. In my opinion, trustworthiness is the primary site of moral evaluation in this area, though I cannot defend that position here. More pragmatically, there is very likely some confusion in our actual attitudes of trust and distrust towards groups. To understand whether our attitudes of trust and distrust are coherent, it is helpful to understand the features of groups – trustworthiness and untrustworthiness – which the attitudes target. Jones (2012) demonstrates the fruitfulness of exploring trustworthiness first, and I hope that the viability of this approach is further demonstrated in the present chapter.

This trustworthiness-first approach does not assume that only trustworthiness can justify trusting (hence my ‘very roughly’ above). There can be many reasons to trust in the absence of trustworthiness, including therapeutic trust aimed at cultivating trustworthiness, psychological self-protection, efficiency in low-risk situations, and so on. Nevertheless, even in such cases, trusting involves behaving as if the recipient were trustworthy. Moreover, we can approach theorising trustworthiness-first without making substantive commitments about how best to cultivate stronger trust relationships in society (Baier (2013); O’Neill (2013)).

Trustworthiness is to be distinguished from mere reliability. Circumstances under which either trust or distrust is appropriate, as opposed to mere reliance or nonreliance, are circumstances under which the actor’s trustworthiness is tested. Many intentional actions fall outside the scope of trustworthiness or untrustworthiness, especially where they seem independent of obligations to others. The headteacher at the school next door regularly releases noisy children into the yard at midday, and I rely upon this to remind me to take a lunch break. But the headteacher’s reliability in this matter is not an issue of trust or distrust for me, and if she decides to keep the children indoors one day, this has no implications for her trustworthiness.

We do not mark the trust-reliance distinction sharply in everyday language, and nor do we mark the trustworthiness-reliability distinction consistently. This everyday loose talk presents a problem to the theorist who hopes to explain the underlying differences between trustworthiness and reliability. Which cases are which? Exactly what are we trying to account for?

Despite these difficulties, the distinction between trustworthiness and mere reliability for individuals is worth investigating because it is morally significant. Trustworthiness is an admirable character trait, something to aspire to and inculcate in one's children, whereas mere reliability may have practical value but does not seem valuable for its own sake. As individuals, it is important for us to understand the difference between circumstances under which our behaviour is a test of our trustworthiness, and circumstances under which it is merely a sign of our predictability or reliability. There are indeed plenty of borderline cases, but there are also plenty of clear cases to work with.

The problem of distinguishing trustworthiness and reliability is all the more difficult in collective cases. Recall the variety of social entities which are spoken about in terms of trust and distrust: is a trustworthy corporation or political system trustworthy in anything like the same way that a spouse or friend can be trustworthy, or is such impersonal trustworthiness more like the reliability of a well-made car? Regarding individuals, it is worth persevering with the trustworthiness-reliability distinction. But what about groups? Do we need a morally-laden notion of trustworthiness for collective entities, or can we get by with mere reliability and unreliability?

4. Group Reliability?

Ideally, I would now propose a detailed account of group reliability, then investigate whether any additional notion of group trustworthiness is both coherent and valuable. Less than ideally, I will simply take it that we have some grasp of reliability and unreliability as these apply to groups. But a few points are worth making.

Both reliance and reliability can be differentiated with reference to tasks or domains. I rely on my alarm clock to wake me up, but I do not rely on it to pass unqueried through airport security. I rely upon my dog to behave sensibly in public (in a somewhat unspecified way), but I do not rely upon her to stay off the furniture. At the limit, reliability in a very specific

respect can be identified with success in a one-off task. I relied upon my alarm clock to wake me up today at 7am, it did so, and it was thus reliable in that particular respect, even if, over the course of several days, it proves to be more generally unreliable, and thus useless as an alarm clock. This generalising move can be understood statistically, or in terms of dispositions, or *ceteris paribus* clauses. As is familiar from the epistemological literature on the ‘generality problem’ for reliabilism, the range or comparison class can make a big difference to judgements of reliability (e.g. Bonjour (2002)). As is obvious from everyday life, reliability can come in degrees.

What about group reliability or unreliability? Sometimes when we talk of a group of people as reliable, we simply mean that all or most people in that group are reliable in the relevant respect, or perhaps, via a generic, that ‘normal’ members of the group are reliable. For example, my students can be relied upon to bring their textbooks to class, and I plan our activities on that basis; the group is reliable in this respect because enough individual group members are reliable. Sometimes, such judgements about individuals are based on their group membership, for better or for worse.

In this paper, I will set aside this type of group reliability, focusing instead upon situations in which the group as a whole acts, and is reliable or unreliable in doing so. We might see *bringing-their-textbooks-to-class* as a kind of group action: this is something the students do together, simply in virtue of each individual bringing his or her textbook to class. But many interesting group actions stand in more complex relationships to the actions of individual group members. For example, the students create a relaxed environment in class, occupy all the seats, and disagree about the nature of free will. I will not attempt to specify the range of possible group actions, nor to engage with the rich, extensive literature on collective or group agency (e.g. List and Pettit 2012). My task in this paper is to investigate whether we need to think of groups as trustworthy actors, or merely as reliable actors. Thus I need to assume that groups can act, reliably or unreliably.

Trustworthiness can be discussed both in the context of practical action, and in the context of testimony: we trust people to do stuff, and we trust people to speak truthfully. Such discussions are enriched when they are integrated with one another. Speaking truthfully is a kind of action, and moreover trusting people to do stuff often involves trusting them to follow through on their words. Nevertheless, trust in testimony does have some distinctive features

which merit special attention. Moreover, there is a small but valuable recent literature on group testimony, so I will begin my discussion with testimony.

5. Trustworthiness and Individual Testimony

There is much we might hope for when others speak. Embarrassingly often, we just hope to be entertained, and on this front we can distinguish someone who is reliably entertaining from someone who can be trusted to be entertaining. For example, someone who is often unintentionally entertaining is reliable in this respect, but not trustworthy.

More soberly, we might hope to learn something when others speak – appreciating what others say is a key source of true beliefs. An individual can serve as a reliable guide to the truth in all sorts of ways. Someone with an open countenance reliably reveals her thoughts via blushing or other ‘tells’, whilst someone else’s eye-bags provide information about her recent lifestyle. Verbal behaviour can also reliably indicate the truth. For example, the questions I ask reveal my preoccupations, my strenuous denial of interest in some topic suggests quite the opposite, and what I say in my sleep indicates my true fears.

This type of reliable openness or readability does not amount to trustworthiness. Why not? A full answer would involve a substantive, and therefore controversial, account of the difference between trust and reliance. But the short answer is that unreliability in these respects does not constitute a betrayal, we do not in general owe such reliability to others, and others are not entitled to resent us if we lack such reliability. That is, one can be unreliable in these ways without being untrustworthy. If I can dance the night away and still look fresh in the morning, this does not make me untrustworthy; likewise if I rarely blush, talk only nonsense in my sleep, and ask questions in a neutral manner. Such traits help me to mislead others if I wish to, and can thereby facilitate untrustworthiness. But unreadability does not constitute untrustworthiness in its own right.

In contrast, in core cases of testimony where an individual asserts something to an audience – or, if you prefer, where an individual tells an audience something – reliability with regard to truth-indication does constitute trustworthiness. As elsewhere, the notion of reliability here is very rough-and-ready. Reliability comes in degrees; moreover we should distinguish between reliability on a given occasion, reliability across a subject matter, and reliability quite generally. My main purpose here is to contrast reliability with trustworthiness, and

individual with group cases: we do not need a full understanding of reliability in order to make some comparative judgements.

But however we understand reliability, reliability in assertion (or telling) is a matter of trustworthiness, because assertion involves undertaking responsibility for what is said. In contrast, speaking truthfully or falsely in one's sleep does not contribute to either trustworthiness or untrustworthiness, because we do not hold people responsible for sleep talk. This aspect of assertion is acknowledged by a wide range of philosophers who otherwise take quite different views of what assertion involves. For Peirce, 'to assert a proposition is to make oneself responsible for its truth' (1932: 384). For Searle an assertion of p 'counts as an undertaking to the effect that p represents an actual state of affairs' (1969: 66). And for Williamson, 'To make an assertion is to confer a responsibility (on oneself) for the truth of its content' (2000: 268-9).

More ambitiously, Brandom (1983) argues that asserting that p involves commitment to withdrawing the assertion if it proves to be mistaken, and/or to defending the assertion against reasonable challenges. Such commitments last beyond the moment of assertion itself. But one can recognise that assertion involves an undertaking of responsibility without recognising any such longer-lasting commitments: this is illustrated by the quotation from Williamson above, which comes just a few pages after he explicitly rejects Brandom's account. Williamson goes on to say that one discharges this responsibility 'by epistemically ensuring the truth of the content' (2000: 269), i.e. by satisfying the condition that one knows what one asserts, at the moment of assertion. This encapsulates Williamson's preferred 'knowledge norm' on assertion, but those who prefer other norms – perhaps a truth norm, or a justified-to-believe norm – can also cast these in terms of discharging one's responsibility as an assertor.

So I will take it that, in the individual case, reliability or unreliability in the provision of (apparent) information becomes a matter of trustworthiness or untrustworthiness when the individual makes an assertion, or tells an audience something, thus somehow taking responsibility for what is said. Can we apply these ideas to collective entities?

6. What is Group Testimony?

What is it for a group to be reliable as a guide to the truth? We can often obtain information by observing the behaviour of groups of people: if the crowd is gathered around laughing, I infer that there is something funny going on. Sometimes we obtain information from the verbal behaviour of a group. Surowiecki (2004) illustrates the ‘wisdom of crowds’ with a phenomenon noted by Francis Galton: when a large number of people at a country fair were asked to estimate the weight of an ox, the average estimate was very accurate, indeed more accurate than the estimates of individual experts.

Interesting and useful though this phenomenon is, it does not involve group assertion, or anything which involves trustworthiness as opposed to mere reliability. Neither the crowd collectively nor any individual member of the crowd can sensibly be held accountable for the accuracy of the average estimate; neither the crowd nor any individual takes responsibility for the truth of the average judgement.

In other circumstances, however, a group or organisation produces something which looks much more like assertion or testimony. Seemingly just as an individual may express herself by employing a spokesperson, issuing a written statement, or speaking in the ordinary way, a group may issue a statement via a spokesperson, or publish a report with the group’s imprimatur. Not all groups are capable of doing this: issuing a statement or report seems to require a certain degree of internal structure and organisation. Moreover it is sometimes unclear whether an individual is authorised to speak on behalf of a group: consider how journalists turn to informal ‘community leaders’ to represent the collective view.

Nevertheless, we often understand groups or collective entities as issuing statements in these ways. This might suggest that we need to distinguish the trustworthiness of a group, with respect to its assertions or testimony, from the mere reliability we see in ‘wisdom of crowds’ cases. Does this provide a clear case of group trustworthiness as something over and above group reliability?

Jennifer Lackey (2014) argues that, when a group ‘speaks through’ a spokesperson, the group’s testimony simply is the spokesperson’s testimony, and thus the epistemic significance of the group’s testimony just is the epistemic significance of the individual testimony. Of course, the epistemic significance of the spokesperson’s statement – whether it is reliable, for example – will depend in large part upon the way in which the content of the

statement was generated by the group and provided to her. But, as Lackey points out, this is often true for individual testimony too. The epistemic significance of what I say to you in my personal capacity depends in large part upon the epistemic standing of the people from whom I acquired my ‘information’.

Suppose we accept Lackey’s deflationary account of the epistemic significance of group testimony. Does this mean that group trustworthiness in testimony is to be identified with the individual trustworthiness of the spokesperson? No: the spokesperson is not making an assertion, and does not undertake responsibility for the truth of what she says. The spokesperson is not like an embedded journalist: she does not, in general, make an assertion *about* the group. Instead, she speaks on behalf of the group, about the relevant subject matter. This distinction between reporting views and expressing them is familiar from ordinary individual assertions: in saying that p , I do not report that I believe that p , but instead I purport to express my belief that p . In making my assertion, I undertake responsibility for the truth of p , not merely for the truth of the claim that I believe that p .

Does the spokesperson undertake responsibility for the truth of what she says on behalf of the group? No. Her responsibility is to execute her duties as a spokesperson, and we can judge her as trustworthy or untrustworthy in that capacity, but this is not to judge her trustworthiness as an assertor, even in this circumscribed setting. Indeed, Lackey’s view is that the spokesperson is not governed by standard norms of assertion, but rather by norms specific to her status as spokesperson (Lackey, forthcoming). So although there is a sense in which what the spokesperson says is what the group says, this does not mean that the spokesperson makes an assertion which is the group’s assertion.

We can judge the spokesperson as reliable or unreliable with regard to the truth in this matter – and Lackey may well be correct about the epistemic significance of what the spokesperson says – but this is not to evaluate the spokesperson’s trustworthiness with regard to this subject matter. If any entity is taking responsibility for the truth of what is said, it is the group rather than the spokesperson.

Does the group take responsibility for what is said on its behalf? One approach to this question would be to take up general issues about collective responsibility, to understand first what it is for a group to *be* responsible for something, then what it is for a group to *take*

responsibility, and specifically to take responsibility for the truth of a statement. But I will adopt a different approach, asking not what it would be for a group to take responsibility for the truth of what is said, but whether we can manage without that notion: what costs would we incur if we did not think of groups in this way?

We saw in the individual case that the distinctions between trust and reliance, and between trustworthiness and reliability, seem to be of central importance in our dealings with one another. The notion of reliance doesn't seem sufficient to account for our interpersonal interactions and attitudes. But what about the group case: can we manage without the notions of group trustworthiness, group responsibility for the truth, and group assertion? I will investigate this by looking first at the epistemic aspects of group testimony, and then at the ethical aspects of group testimony. I will argue that in each case we can vindicate our practices without resort to a distinction between trustworthiness and reliability; this does not settle the question of whether groups can be trustworthy, but it does shift the burden of proof.

7. Group Trustworthiness in Testimony: Epistemic Aspects

A number of authors, including Hinchman (2005), Moran (2006), McMyler (2007) and Faulkner (2011) have argued that trust, as opposed to mere reliance, plays an important role in the epistemic significance of individual testimony. There are significant differences amongst these various authors, but some common themes emerge. On this picture, the speaker's undertaking of responsibility for what is said provides the intended audience with a distinctive epistemic reason to believe what is said. Conversely, the audience's trust, not mere reliance, in the speaker plays a key role in justifying beliefs acquired through testimony. The speaker offers his or her assurance to the audience, and, when circumstances are favourable, this assurance gives the audience reason to trust and reason to believe.

If assurance plays a key epistemic role in our acquisition of knowledge from testimony, then so does the trust-reliance distinction, and thus also the trustworthiness-reliability distinction. When we learn from what someone tells us, complete with assurance, we regard her as trustworthy in this matter. According to the assurance view, we are then in a better epistemic position than when we merely learn from someone's verbal behaviour, regarding her as merely reliable.

Thus an assurance account of the epistemic significance of group testimony offers the best prospect for vindicating a trustworthiness-reliability distinction in that domain. Assurance accounts of individual testimony emerged in reaction to more standard views, and they have been criticised on various counts (e.g. Lackey 2008); I will not try to adjudicate this debate in the individual case. Rather, I will explore what non-assurance and assurance accounts of the epistemic significance of group testimony might look like, before arguing that assurance accounts are less motivated in the group case than in the individual case.

In her (2007), Deborah Tollefsen advocates a non-assurance reductionism about the justification of the beliefs we form on the basis of group testimony. Such beliefs are justified insofar as we can monitor the reliability of the group's testimony more generally; the view is reductionist in that the justification of testimonial beliefs is reducible to justification gained from other sources. On this view, group pronouncements are treated as just more evidence, in the mix with group behaviours and features of all sorts; there is no epistemic role here for group trustworthiness over and above group reliability. Similarly, Lackey's identification of group testimony with spokesperson testimony allows us to assess the epistemic merits of group testimony by assessing the testimony of the spokesperson. Lackey herself adopts a non-assurance view of the epistemic significance of individual testimony (2008), meaning that, on this view, there is no epistemic role for assurance in group testimony.

What would an assurance account of the epistemic significance of group testimony look like? In a later paper, Tollefsen advocates a more demanding notion of what it is for a group to testify: '...the fact that groups issue intelligible statements either in writing or via a spokesperson seems to me now not sufficient to say that they, themselves, are testifiers.' (2011: 12) Tollefsen turns to assurance accounts, seeing testimony as essentially interpersonal, 'deeply tied to epistemic responsibility' (15). For groups to testify, on this picture, they must be able to acknowledge their epistemic responsibilities; this echoes the connections I have highlighted between trustworthiness and fulfilling responsibilities.

One of Tollefsen's real-life examples involves a panel of scientists, tasked to issue a definitive statement about the genetic hazards of radiation. Although the scientists had underlying disagreements about the facts, Tollefsen argues that 'group members realized the fact that others were depending on the group to speak its mind and that they had a responsibility to say something definitive and take collective responsibility for what was said'

(2011: 17). It is striking that Tollefsen refers to the individual group members as acknowledging responsibility for getting the group as a whole to take responsibility for what is said.

Miranda Fricker (2012) asks what is required for a group to be able to offer assurance in its testimony, and suggests 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.’ (272) The notion of joint commitment here is based on that developed by Margaret Gilbert across many works (e.g. her 2006). According to Fricker, group members make – initially to one another – commitments to be jointly trustworthy. The resulting joint commitment helps constitute a group which can then itself offer assurance to its audience. Much like Tollefsen, Fricker gives a key role to individual commitments, within the collective context, in generating assurance on behalf of the group.

Fricker’s picture has a certain appeal when we think about the ‘good’ cases, in which well-intentioned people get together in the hope of forming a trustworthy group. Fricker shows how it can accommodate cases in which such a group becomes untrustworthy because members do not fulfil their commitments to joint trustworthiness, perhaps because of personal corruption.

However other sorts of cases do not fit so well. Imagine a nefarious group, formed by members who make a mutual commitment to jointly deceive the public (insert your own example here). Such groups are constituted by the individuals’ commitments to deception, not by joint commitments to trustworthiness, so do not fall under Fricker’s account. Are these groups ‘suited to enter into the second-personal relations of trust that characterize testimony’? Not if we construe the question literally: nobody should trust such groups. But if some groups are trustworthy, then surely these nefarious groups qualify as untrustworthy, not merely unreliable (Lackey makes a related point in her draft ‘Group Lies’, fn.9). Such groups seem as suited to enter into second-personal relations of distrust, resentment, and betrayal, as more admirable groups are suited to enter into second-personal relations of trust.

Moreover Fricker’s account does not allow for the possibility of trustworthy organisations whose members – for example, employees – follow appropriate procedures, but are not

motivated by a desire for or commitment to group trustworthiness. Perhaps Fricker would be happy to conclude that such a group does not produce genuinely assurance-based testimony, but is in practice more similar to a information-producing machine, or a source of information, rather than an informant, in the terms she shares with Craig (1990).

Despite these concerns, some groups do seem to operate in the ways Fricker describes. But what is the epistemic significance of this? Fricker herself argues that, in the individual case, assurance provides a different, but not essentially stronger reason to believe what is said than does merely overheard testimony. She writes that:

an addressee [who receives assurance] and an eavesdropper [who does not] could have exactly the same background reasons to trust [an individual] 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. (268-9)

Less sympathetically, Lackey (2008: 249) suggests that the difference between trust and reliance 'may very well be psychologically, morally, or even pragmatically relevant', whilst denying that it has epistemic significance.

One can recognise the significance of assurance and trustworthiness in the individual case without downgrading the *epistemic* significance of the ways in which we can learn from others' words without being mediated by assurance. Moreover, the various considerations which may make assurance views of epistemic significance attractive in the individual case do not easily transfer to the group case.

For example, Moran argues that the fact that individual speakers freely choose what to say is what makes testimony epistemically valuable to the audience, indeed more valuable than a hypothetical opportunity to directly inspect the speaker's beliefs. But this is puzzling given a non-assurance view of testimony which regards speech as mere evidence of the speaker's beliefs:

If speech is seen as a form of evidence [of the speaker's beliefs], then once its intentional character is recognized (that is, not just as intentional behavior, but

intentional with respect to inducing a particular belief) we need an account of how it could count as anything more than *doctored* evidence (2006: 278).

On the assurance account, argues Moran, the fact that the speaker opts to ‘stand behind’ her words, actively offering them as reasons to believe, is an advantage, not a disadvantage to the audience. (Keren 2012 challenges Moran on this front.)

The analogous ideas are much less compelling in the group case. First, it is not obvious that non-assurance accounts of the epistemic significance of group statements must see them as evidence of group belief, especially since the very notion of group belief is disputed. Second, group statements do not obviously have the status of freely-chosen words, in any relevant sense. Instead, we can see groups as producing statements via the functioning of various internal mechanisms; we may then consider whether those mechanisms lend themselves to the production of truth or false statements. Nonassurance accounts of group testimony are not challenged by the considerations Moran raises for the individual case.

Advocates of assurance-style accounts of individual testimony sometimes lean on ideas about respect for others as interlocutors. For example Fricker (2007, chapter 6) argues that, under some circumstances, regarding others as mere sources of information rather than as informants involves an ethically dubious brand of objectification. And Hinchman (2005) motivates his account with reference to the insulting ‘slights’ and ‘rebuffs’ we create when we do not accept what we are told. But these considerations bear much less weight in group cases: objectifying group testifiers is problematic or insulting only to the extent that it in fact amounts to objectifying or insulting individual members of the group, perhaps on the basis of their group membership.

Tollefsen (2007) argues that Burge-style anti-reductionism about the epistemology of testimony does not fit well with group statements: there is no default of accepting what groups say, in the absence of evidence of their reliability. Tollefsen sees this as favouring a reductionist rather than an anti-reductionist approach to the evidence offered by testimony, considering an assurance view only in her later work. But I think similar considerations tell against an assurance account of group testimony: even if we needed the epistemic boost provided by assurance to get the practice of testimony going in the first place, the resources thereby provided through individual testimony seem sufficient to underpin an evidential account of the epistemic significance of group testimony.

To summarise: we can acknowledge the epistemic significance of group testimony without requiring a notion of group assurance, responsibility or trustworthiness; moreover, the arguments used to motivate assurance accounts of the epistemic significance of individual testimony do not easily transfer to the group case. This does not demonstrate that assurance views are false in the group case, but it does place a heavy burden of proof upon their supporters, over and above the burden of proving the strength of assurance accounts in the individual case.

8. Group Trustworthiness in Testimony: Ethical Aspects

I have argued that we do not need the distinction between group trustworthiness and mere group reliability in order to make sense of our epistemic practices around group testimony. What about our ethical practices? The epistemic significance of assurance is contentious even in the individual case. But no-one doubts the ethical significance of trustworthiness for individuals; lying is a paradigmatic untrustworthy act, with paradigmatic ethical significance, and reckless testimony which pays no regard to truth or falsity is likewise ethically culpable. So it might seem that, if we could not attribute lies (as opposed to mere false statements) to a group, nor regard groups as either trustworthy or untrustworthy, then we would lose an important aspect of accountability both in law and in morals.

The issue of corporate or collective moral responsibility is complex and contested (e.g. Isaacs 2011; van de Poel, Ryackers, and Zwart 2015), and I cannot hope to engage it here in any depth. Instead, I will focus on an issue which emerges more directly from the literature on trust and trustworthiness: reactive attitudes. Following Baier, philosophers distinguish trust and reliance as directed at individuals because of the intimate connections between trust and certain reactive attitudes, connections which are missing in the case of reliance. And certainly some of us seem ready to react to groups and organisations with attitudes like loyalty, gratitude, resentment, and a sense of betrayal. If groups are not genuinely trustworthy or untrustworthy with regard to their testimony, merely reliable or unreliable, then such reactions would seem to be mistaken.

I think many of us do direct trust-and-distrust-related reactive attitudes towards groups, and that these are indeed misplaced if groups can be neither trustworthy or untrustworthy

(Thompson forthcoming discusses reactive attitudes in connection with groups).

Nevertheless, there are two related reactions which can be appropriate, even without group trustworthiness or untrustworthiness. First: individual people are liable to praise and blame, gratitude and resentment, for their personal trustworthiness and untrustworthiness in helping generate group statements. Second: it is appropriate to react positively (negatively) to living in a society in which important groups and institutions are reliable (unreliable) producers of true statements. I will examine these ideas in turn.

Imagine a situation in which a number of people undertake to help ensure that a group issues reliable statements. This may include members of the group, who can fill different roles in their organisation, but may also include non-members, for example people who founded or designed the group, or people now tasked with regulating or overseeing its activities. Such individuals may be more or less trustworthy in playing their various roles, and can sensibly be judged, praised, or resented for their actions in that capacity.

This picture is reminiscent of Fricker's joint commitment account of group assurance, but I draw different conclusions from it. Each individual takes responsibility for helping ensure that the group produces a true statement, but in line with my arguments of the previous section, none need take responsibility for the truth of the group's statement in the way required for assertion, and nor does the group itself take such responsibility. Even if we endorse a Williamsonian knowledge norm on assertion, for example, we need see nothing problematic in a situation where none of the group members knows the truth of the group's statement. On this picture, a group statement is like a reading provided by a complex machine: we can hold the designers, operators, and maintainers responsible for ensuring that the machine provides accurate readings, without regarding any individual as asserting the content of the reading.

So reactive attitudes connected to 'trusting' a group can sensibly be directed at individuals connected to the group. But is this really feasible where the audience does not know who these individuals are? Yes: you easily resent the person, whoever it was, who wrote graffiti on your front door, and your feelings are quite different about the wind which inconveniently blew litter into your garden. Likewise, you can resent the individuals who contributed to the publication of a misleading report, even if you do not know who those individuals are.

So far, I have focused on situations in which we can regard individuals as trustworthy to the extent that they help ensure the reliability of a group's statements. But, as I suggested in response to Fricker, other situations do not fit this pattern. There could be situations in which the activity of various people leads to a reliable group 'statement', although the individuals themselves are merely reliable or unreliable, rather than trustworthy or untrustworthy. Perhaps the 'wisdom of crowds' cases are like this, at least where crowd members have no particular responsibility to be careful or accurate. But, appropriately, if we appreciate that we face such a situation, we are not tempted to adopt reactive attitudes towards the crowd.

Less innocuously, there seem to be situations in which individuals perform well in their roles exactly to the extent that a group produces unreliable statements, precisely because that is the group's purpose. Perhaps this true of a PR agency, a lobbying group, or a propaganda 'machine'. It is natural for an audience to feel resentful if they discover the unreliability of their sources, yet there may be no individual who is untrustworthy in the sense of not doing her job properly. Is this a situation in which we must either take up reactive attitudes towards the group itself, or else abandon the participant stance?

No. In many such circumstances, there are individuals who may legitimately be resented for creating or perpetuating an agency whose main function is to mislead. Whether this involves a failure of trustworthiness, rather than some other sort of moral or political failure, may depend upon which detailed account of trust and trustworthiness we espouse, but reactive attitudes reach beyond the domain of trustworthiness and untrustworthiness in any case.

In other circumstances, there may be no individual(s) to blame, and yet still it seems reasonable to feel resentful about societal structures: this is the kind of generalised resentment I alluded to above. For example, we might well feel angry that we live in a society where many media outlets fail to provide us with reliable information, even supposing (implausibly) that there are no specific individuals who can be held accountable for this. This raises large issues within political and social philosophy which I cannot explore here; a natural first step would be to use Iris Marion Young's notion of 'structural injustice' (e.g. 2011).

I have suggested that we can retain much of our ethical practice around group testimony without needing to invoke notions of group trustworthiness and untrustworthiness, as

opposed to group reliability and unreliability. But the connection with reactive attitudes was not the only ethical motivation for distinguishing trustworthiness and reliability in the individual case: I also suggested that the former, unlike the latter, is a candidate virtue.

Fricker (2012), drawing on Lahroodi (2007), offers an account of institutional epistemic virtues such as fair-mindedness. She focuses primarily on the ways in which individuals, acting as members of collectives, can display virtue and vice. She also discusses virtue and vice in institutional structures, arguing that an institution in isolation cannot be virtuous or vicious, only as it is populated by individual people. Responding to Fricker's earlier (2007), Elizabeth Anderson emphasises the need for structural remedies to testimonial injustice, over and above remedies involving individual action and cultivation of virtue. She says 'when the members of an organization jointly commit themselves to operating according to institutional principles that are designed to achieve testimonial justice...this is what it is for the organisation itself to be testimonially just' (2012: 168-9).

These are compelling thoughts: we need to address structural problems and collective contexts if we are to combat injustice and create better institutions. However, we do not need the distinction between trustworthiness and mere reliability at the group level in order to pursue these projects. The distinction matters at the individual level precisely because we cannot always require others to be reliable in respects which matter to us. I cannot require that you read all my published works (unless that's you, Mum – hi!), and if you do not do so, then you are unreliable in that respect, but in no sense untrustworthy. There is no general obligation upon individuals to be reliable, which is why we need the language of trustworthiness to highlight those particular respects in which individuals are obliged to be reliable. But we can require of our institutions that they be reliable in the respects that matter to us, or at least we can require this of our public institutions.

9. Trustworthiness and Group Action

I have focused on trustworthiness and reliability in the context of group testimony, arguing that we can abandon the former notion without incurring significant costs. Does the same go for trustworthiness in the context of group action? Some of the issues I considered – for example the connections between trustworthiness, epistemic responsibility, and assertion – seem specific to the case of testimony, so the transition will not be seamless. Matters are complicated by the fact that, even in the individual case, different philosophical accounts of

trust distinguish between trust and mere reliance in quite various ways: what it takes for an issue to become a question of trust may be to do with the perceived or actual motivations, interests, commitments, or obligations of the actor, according to different theories.

But the ethical issues I discussed in the preceding section do seem to apply not just to the case of testimony but to group behaviour more generally. The trust-reliance distinction is motivated by the connections between trust and certain reactive attitudes, and also through the moral significance of trustworthiness, over and above reliability. I suggested above that although we do in fact direct trust-related reactive attitudes towards groups, with regard to their testimony as well as their actions more generally, we could more appropriately direct these towards individuals, known or unknown, group members or not, who are in various ways responsible for the functioning of the group. To generalise this account, we would need to think carefully about what we can reasonably expect of group behaviour, and also about the ways in which individuals can be responsible through inaction, as well as through action.

Moreover, there are significant differences in our thinking about trustworthiness in individuals and trustworthiness or reliability in groups. Regarding an individual as trustworthy is typically caught up with respect for her as autonomous in some way: someone who reliably follows a benevolent despot's orders under duress is not displaying her trustworthiness. Trustworthy behaviour often, though not always, reflects a determination to fulfil obligations or commitments which were voluntarily acquired. Likewise, it can be important to recognise when an individual's behaviour is *not* a matter of trustworthiness or untrustworthiness: if the headteacher does not release the noisy children at noon, so I am not reminded to take my lunch break, I should recognise that she has not betrayed me in any way.

Do we need to regard groups and organisations in the same way? It may depend upon the group, its constitution and supposed purpose, including questions about whether an organisation is a element of the state or a private entity, and whether membership of a particular group is chosen or imposed. To what extent do certain groups have obligations to individuals, or to other groups? As these considerations indicate, issues of trustworthiness in group action quickly involve larger questions about collective responsibility and commitment; a fruitful line of inquiry would be to investigate whether or not these questions can be sidestepped by thinking in terms of reliability rather than trustworthiness.

9. Onwards and Outwards

I have considered the question of whether we need a notion of group trustworthiness, over and above group reliability, in order to understand the epistemic and ethical significance of group testimony, and my tentative answer was ‘no’: reliability is all that we need. Likewise, I suggested, we do not need this distinction in order to understand our attitudes to group behaviour more generally. My discussion was not comprehensive, and I do not regard my arguments as conclusive. Nevertheless, I have begun to cast doubt on the viability of the trustworthiness-reliability distinction as applied to groups.

Suppose that in fact this distinction does no useful work. Does this mean we should abandon talk of group trustworthiness? Should we regard public and interdisciplinary discussion of organisational trustworthiness (e.g. Hawley 2012: chapter 8) as simply confused?

At one level this is a strategic issue about choice of terminology, and about whether it is worth struggling to impose quasi-technical terms onto broader discourse. A related issue arises even in the individual case. As I noted earlier, although the trust-reliance distinction is rightly valued by philosophers discussing interpersonal trust, it is not consistently marked in ordinary language: we often talk of trusting or distrusting inanimate objects, for example, or indeed trusting someone to get things wrong. One response is to change terminology. Instead of contrasting trust and reliance, we might follow Hollis (1998: 10) in distinguishing normative trust from predictive trust, or, as Faulkner (2007: 880) prefers, affective trust from predictive trust. Whilst these choices may have independent philosophical merit, they do not help us avoid jargon in the public realm.

More substantively, whatever our terminology, if the trustworthiness-reliability distinction lacks merit in the group case – or, in other terms, if there is a type of trustworthiness which individuals but not groups can exemplify – then this is of significance to wider debates. Even in the individual case, many public concerns about ‘trustworthiness’ are really about reliability, but they are pressing concerns nonetheless.

Moreover, many public concerns about trust and groups are in fact best construed as (genuine, pressing) concerns about trusting individuals in a group context. I have discussed some aspects of this above, in connection with individual responsibilities for group behaviour. But in addition distrust or trust of an individual may be caused or rationalised by

beliefs about her group membership. Fricker (2007) shows how such judgements can reflect ethically culpable prejudices, but not all group-based trust or distrust is problematic in this way. Indeed, if we construe 'group membership' very loosely, almost any trust or distrust can be construed as group-based: I trust you because of your good track record, i.e. your membership of the group of people who have good track records.

This notion of group membership may seem trivial, but other groups are more robustly unified: they are held together not just by some common feature of the members, but by informal or formal relations and structures within the group. Such relations and structures may genuinely help to cause, promote or constitute the trustworthiness of group members in relevant domains, either by generating sanctions and rewards for individual members, or by enforcing trustworthiness-based entry barriers. Professional self-regulatory bodies are supposed to fit this pattern: membership of the body indicates professional competence. On the other hand, of course, groups may cause, promote, or indicate the untrustworthiness of individual members in certain respects, for example by creating incentives for members to put group loyalty above trustworthiness to non-members, or by reducing members' competence in certain respects.

The ways in which different group structures and identities can enhance or discourage individual trustworthiness in different domains are investigated empirically in organizational studies, and by social psychologists. Consider peer pressure, or the variety of phenomena which could be studied under the heading of 'institutional' or 'corporate' culture or climate. For example, such local cultures can influence which respects of trustworthiness are seen as most important. Is it a priority to defend your team, or the company, or to 'champion' the client, or blow the whistle where necessary? It is easy to get into a situation where the demands of trustworthiness point in different directions, and different cultures may indicate different resolutions of such dilemmas. Another dimension of variety may be the importance placed upon individual trustworthiness as opposed to other virtues or goals.

Finally, although I have distinguished our attitudes to groups from our attitudes to individuals in group contexts, this distinction is not always sharply marked in practice. We often interact with individuals as members, leaders, spokespersons, or representatives of groups, and we may do so with varying degrees of trust. But it can then be indeterminate whether our attitude of (dis)trust targets the individual, the group, or both; if both, there is a further

question as to which comes first, epistemically speaking. Do I trust the customer services representative because I have found the company to be reliable, or does my relationship with the individual come first? There will of course be no single answer to such questions. In fact, this gives us another good reason to start with the notion of trustworthiness or reliability – what does it take for an individual or group to exhibit these features? – rather than with our often indistinct attitudes of trust and distrust, reliance and nonreliance.

Acknowledgements

Material from this chapter was presented at the New Philosophical Perspectives on Trust conference held at the Blavatnik School of Government in Oxford; at a workshop on Trust at the University of Manchester; as a Women in Philosophy lecture at the University of Graz; at a seminar at KCL; and (remotely) to the Cambridge Trust seminar. Many thanks to the organisers and participants on all those occasions for their hospitality and philosophical engagement. In addition I am grateful to Paul Faulkner, Jennifer Lackey, Tom Simpson, Chris Thompson, Deborah Tollefsen, and Leo Townsend, for written comments and questions. This work was supported by a Major Research Fellowship from the Leverhulme Trust, which I very gratefully acknowledge.

References

- Anderson, Elizabeth S. (2012): ‘Epistemic Justice as a Virtue of Social Institutions’, *Social Epistemology*, 26.2: 163-173.
- Baier, Annette (1986): ‘Trust and Antitrust’, *Ethics* 96.2: 231-60.
- Baier, Annette (2013): ‘What is Trust?’ in Archard, Deveaux, Manson and Weinstock (eds.) *Reading Onora O’Neill*, Routledge: 175-185.
- Bonjour, Laurence (2002): ‘Internalism and Externalism’, *Oxford Handbook of Epistemology*, ed. Paul K. Moser, Oxford University Press: 234-263.
- Booth, Ken and Wheeler, Nicholas J. (2008): *The Security Dilemma: Fear, Cooperation and Trust in World Politics*, Palgrave Macmillan.
- Brandom, Robert (1983): ‘Asserting’, *Noûs*, 17: 637-50.
- Brown, Jessica and Cappelen, Herman (eds.) (2011): *Assertion: New Philosophical Essays*, Oxford University Press.
- Craig, Edward (1990): *Knowledge and the State of Nature*, Oxford University Press.
- De Bruin, Boudewijn (2015): *Ethics and the Global Financial Crisis*, Cambridge University Press.

- Faulkner, Paul (2011): *Knowledge on Trust*, Oxford University Press.
- Fricker, Miranda (2007): *Epistemic Injustice*, Oxford University Press.
- Fricker, Miranda (2009): 'Can There be Institutional Virtues?', *Oxford Studies in Epistemology* 3, eds. Gendler and Hawthorne.
- Fricker, Miranda (2012): 'Group Testimony? The Making of a Good Informant', *Philosophy and Phenomenological Research*, 84.2: 249-76.
- Hinchman, Edward (2005): 'Telling as Inviting to Trust', *Philosophy and Phenomenological Research*, 70.3: 562-587.
- Hodgson, Geoffrey (2006): 'What are Institutions?', *Journal of Economic Issues*, 40.1: 1-25.
- Hollis, Martin (1998): *Trust Within Reason*, Cambridge University Press.
- Holton, Richard (1993): 'Deciding to Trust, Coming to Believe', *Australasian Journal of Philosophy*, 72: 63-76.
- Isaacs, Tracy (2011): *Moral Responsibility in Collective Contexts*, Oxford University Press.
- Jones, Karen (2012): 'Trustworthiness', *Ethics* 123.1: 611-85.
- Lackey, Jennifer (2008): *Learning From Words*, Oxford University Press.
- Lackey, Jennifer (2014): 'A Deflationary Account of Group Testimony', in Lackey (ed.) *Essays in Collective Epistemology*, Oxford University Press.
- Lackey, Jennifer (forthcoming): 'Group Assertion', *Erkenntnis*.
- Lackey, Jennifer (draft): 'Group Lies'.
- Lahroodi, Reza (2007): 'Collective Epistemic Virtues', *Social Epistemology*, 21.3: 281-97.
- List, Christian and Pettit, Philip (2011): *Group Agency*, Oxford University Press.
- Mäkelä, Pekka and Townley, Cynthia (2013): 'Editors' Introduction', *Trust: Analytic and Applied Perspectives*, edited by Mäkelä and Townley, Amsterdam: Rodopi.
- Mathieson, Kay (2006): 'The Epistemic Features of Group Belief', *Episteme* 2.3: 161-75.
- McMyler, Ben (2007): 'Knowledge at Second Hand', *Inquiry*, 50.5: 511-540.
- Moran, Richard (2006): 'Getting Told and Being Believed', in Lackey and Sosa (eds.) *The Epistemology of Testimony*, Oxford University Press: 272-306.
- O'Neill, Onora (2002): *Autonomy and Trust in Bioethics*, Cambridge University Press.
- O'Neill, Onora (2013): 'Responses' in Archard, Deveaux, Manson and Weinstock (eds.) *Reading Onora O'Neill*, Routledge: 219-39.
- Peirce, C.S. (1932): *Collected Papers of Charles Sanders Peirce vol 2*, edited by C. Hartshorne and P. Weiss, Harvard University Press.

- Saunders, Mark N. K., Skinner, Denise, Dietz, Graham, Gillespie, Nicole, and Lewicki, Roy J. (eds.) (2010): *Organizational Trust: A Cultural Perspective*, Cambridge: Cambridge University Press.
- Schmitt, Frederick (1994): 'The Justification of Group Beliefs', in *Socializing Epistemology*, Schmitt (ed.), Rowman and Littlefield: 257-88.
- Searle, John (1969): *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press.
- Simon, Judith (2013): 'Trust' in Pritchard, D. (ed.) *Oxford Bibliographies in Philosophy*, New York: Oxford University Press.
- Surowiecki, James (2004): *The Wisdom of Crowds*, Little Brown.
- Sztompka, Piotr (1999): *Trust: A Sociological Theory*, Cambridge University Press.
- Thompson, Christopher (forthcoming): 'The Moral Responsibility of Group Agents'.
- Tollefsen, Deborah (2006): 'Group Deliberation, Group Cohesion and Scientific Teamwork: is there Room for Dissent?', *Episteme*, 3: 37-51.
- Tollefsen, Deborah (2007): 'Group Testimony', *Social Epistemology* 21.3: 299-311.
- Tollefsen, Deborah (2011): 'Groups as Rational Sources', In Hans Bernhard Schmid, Daniel Sirtes & Marcel Weber (eds.), *Collective Epistemology*, Ontos.
- Van de Poel, Ibo, Royakkers, Lambèr, and Zwart, Sjoerd D. (2015): *Moral Responsibility and the Problem of Many Hands*, Routledge.
- Watson, Gary (2004): 'Asserting and Promising', *Philosophical Studies*, 117.1-2:57-77.
- Weinstock (2013): 'Trust in Institutions', in Archard, Deveaux, Manson and Weinstock (eds.) *Reading Onora O'Neill*, Routledge: 199-218.
- Williamson, Timothy (2000): *Knowledge and its Limits*, Oxford University Press.
- Young, Iris Marion (2011): *Responsibility for Justice*, Oxford University Press.