

penultimate version of a paper published in *Philosophia: Philosophical Quarterly of Israel* (DOI: 10.1007/s11406-012-9408-z); final version available upon request

Coverage-Reliability, Epistemic Dependence, and the Problem of Rumor-Based Belief

Axel Gelfert
Department of Philosophy
National University of Singapore
3 Arts Link, 117570 Singapore
email: axel[at]gelfert{dot}net

Abstract

Rumors, for better or worse, are an important element of public discourse. The present paper focuses on rumors as an epistemic phenomenon rather than as a social or political problem. In particular, it investigates the relation between the mode of transmission and the reliability, if any, of rumors as a source of knowledge. It does so by comparing rumor with two forms of epistemic dependence that have recently received attention in the philosophical literature: our dependence on the testimony of others, and our dependence on what has been called the ‘coverage-reliability’ of our social environment (Goldberg 2010). According to the latter, an environment is ‘coverage-reliable’ if, across a wide range of beliefs and given certain conditions, it supports the following conditional: *If $\sim p$ were true I would have heard about it by now*. However, in information-deprived social environments with little coverage-reliability, rumors may transmit information that could not otherwise be had. This suggests that a trade-off exists between levels of trust in the coverage-reliability of official sources and (warranted) trust in rumor as a source of information.

Introduction

Crisscrossing the United States at a time of uncertainty and polarization, a frustrated observer of current affairs had this to say about the various political smear campaigns he witnessed: ‘So many absurd and contradictory things are said that it becomes necessary to renounce an opinion founded on rumors and admit only those probabilities which are approved by the strictest good sense and the most rigorous moderation.’¹ The sentiment sounds familiar enough, but the time is 1862, not 2012, and the commentator is not a political pundit, but the traveling pianist and composer Louis Moreau Gottschalk. Philosophers, too, have equated rumor with ‘public error’ (Michel de Montaigne,

¹ Quoted after Gottschalk (2006: 90).

Essays) or characterized it as ‘typically untrue information’ (Chakrabarti 1992: 424). The present paper argues that, from an epistemic point of view, rumors do not deserve such outright rejection. Under certain conditions, rumors have a positive epistemic role to play, especially when other sources of information in our social environment – those that would normally underwrite ‘coverage-supported beliefs’ – are unavailable.

The structure of the paper is as follows: Section 1 develops the notion of coverage-reliability, recently introduced by Sandy Goldberg (2010), and its relation to phenomena of epistemic dependence. Section 2 attempts a tentative definition of rumor by way of comparison with the related phenomenon of gossip and, in addition to introducing some basic terminology regarding rumor-based belief, also provides a discussion of the very real tension between epistemic authority and social power; Section 3 turns more specifically to the problem of the reliability of rumor as a source of information. Section 4 analyzes a tradeoff (between *timeliness* and *epistemic penetration*) that affects any phenomenon of diffusion of information across social networks. Finally, in Section 5, a preliminary answer is given to the question of when one should consider trusting rumors.

1. Epistemic dependence and coverage-reliability

Recent social epistemology has shown that our dependence on others for knowledge runs deeper than has traditionally been acknowledged. Not only do we all, in the course of everyday life, depend on knowledge gained from other people, but some of our most respected forms of knowledge – notably, scientific knowledge – likewise depend on the ‘division of cognitive labor’ (Kitcher 1990). In what follows, John Hardwig’s expression ‘epistemic dependence’ (Hardwig 1985) is used to refer to any situation in which a subject’s belief depends for its formation, sustainment, or reliability on the cognitive processes of one or more (other) epistemic agents. To say that a subject *H* epistemically depends on *A* (where *A* may be another person or a community) is simply to note this sort of dependence.² In the remainder of this section, two kinds of epistemic

² This is a more direct, and somewhat more liberal, definition than can be found in Hardwig’s original paper, where he describes *H*’s belief as ‘epistemically grounded in an appeal to the authority of *A* and *A*’s belief’ (1985: 336), thus leaving open the nature of the grounding relation and the precise role of the appeal to authority. At the same time, ‘epistemic dependence’ as defined here is more specific than Goldberg’s notion of ‘epistemic reliance’ (2010: 92), in that the former always involves dependence *on others*, whereas the latter also includes reliance on our own cognitive capacities, e.g. on memorial beliefs.

dependence are discussed, both of which play an important role in our epistemic lives. This focus on two broad classes of cases is not meant to deny the possibility of other, perhaps more indirect, varieties of epistemic dependence.³

The first kind of epistemic dependence to be discussed here is our extensive reliance on testimony as a source of knowledge. The epistemology of testimony has received much attention in recent years, not least the question of whether testimony-based beliefs ultimately depend for their justification on the more primitive epistemic sources of perception, memory and inference. Whether one sides with reductionists or anti-reductionists on this point, it is now widely accepted that any successful theory of epistemic dependence must be able to account for what Elizabeth Fricker calls a ‘Commonsense Constraint’: *viz.*, ‘that testimony is, at least on occasion, a source of knowledge’ (Fricker 1995: 394).

The second main kind of epistemic dependence concerns the dependence on our social environment for *coverage*. We all depend, to some degree or another, on our social environment for keeping us abreast with important developments and general knowledge of the changing world around us. In normal social intercourse, others frequently volunteer information they think we ought to know about, and most of us have epistemic routines – such as reading the newspaper – that naturally expose us to a broad range of facts. Clearly, someone who reads a newspaper for general knowledge of world affairs (rather than, say, for the purpose of mere diversion), not only expects each report to be individually truthful and reliable (enough), but also expects the selection of articles as a whole to be reasonably comprehensive, in that it should not leave out important information. Even if we cannot expect completeness from one source of information alone, we may (and do) rely on other sources – formal and informal – to fill us in on any important news we might have missed. Sandy Goldberg, who has recently investigated this form of epistemic dependence, speaks of our reliance on ‘the *coverage-reliability* of one’s community’ (2010: 154). In particular, he draws attention to ‘coverage-supported beliefs’, the existence of which he takes to show ‘that our dependence on others for what we know and justifiably believe outstrips our reliance on their testimony’ (2010: 156).

On Goldberg’s account, a belief *p* is *coverage-supported* if it is formed or sustained by appeal to the ‘*truth-to-testimony* conditional’ (Goldberg 2010: 157): ‘*If ~p were true I would have heard about it by now.*’ Coverage-supported beliefs express, by and large, a commitment to the *status quo*: We trust that, by and large, states of affairs –

³ One example that immediately springs to mind is our dependence on experts and epistemic communities for semantic resources; see (Burge 1986) and (Goldberg 2007).

at least those we hold stable beliefs about – will continue to obtain, and any significant changes will be deemed newsworthy enough, by sources in our social environment, to merit reporting. John McDowell, in a similar context, notes that many of our knowledge claims concern what are ‘reasonably durable’, but in fact ‘impermanent states of affairs to whose continued obtaining we have only intermittent epistemic access’ (McDowell 1998: 422). Yet, on pain of skepticism, it would be quite ill-advised to make an inference from the mere fact of their impermanence to the impossibility of having knowledge of such matters.

Goldberg’s own hypothetical example involves Smith, to whom it occurs – out of the blue – ‘to wonder whether the Prime Minister announced a new major change in foreign policy last week’:

Believing that if the Prime Minister had done so he (Smith) would have heard about it by now, Smith forms the belief that the Prime Minister did not announce a new major foreign policy change last week. (Goldberg 2010: 155)

The fact that one can generate any number of coverage-supported beliefs in this way – by raising the question of whether, unbeknownst to oneself, the world has changed in significant ways, and immediately concluding that no such change has occurred, since if it had one would have heard about it – should not obscure the fact that, whatever validity the truth-to-testimony conditional (*If $\sim p$ were true I would have heard about it by now*) has, does in no way depend on its being about $\sim p$ rather than p .⁴ In both cases, we exhibit an unspecific reliance on our social environment to be both

- (i) reliably apprized of the relevant facts in a certain domain, and
- (ii) disposed to offer reliable reports regarding the obtaining of these facts (when they are believed by the source to have obtained). (Goldberg 2010: 157).

It is important to emphasize what the truth-to-testimony conditional, and the associated notion of coverage-reliability, are not. They are not a defense of the general view that absence of evidence *per se* constitutes evidence of absence. Importantly, the truth-to-testimony conditional depends for its heuristic value on prevalent conditions in the subject’s social environment. If a subject H forms a belief p by relying on her social environment for coverage in a domain of interest D , then she epistemically depends on there being a source A in her environment such that A

⁴ Indeed, Goldberg formulates the truth-to-testimony conditional directly in terms of p rather than its negation.

(i) will (investigate and) reliably determine whether p , (ii) will be reliable in reporting the outcome of that investigation, and (iii) will satisfy both of the previous two conditions in a timely fashion (more on which below).⁵

Appeals to the truth-to-testimony conditional, then, must be carefully circumscribed, and relying on one's community for coverage is only a legitimate belief-forming strategy if one's community is indeed coverage-reliable.⁶

Two further contrasts merit brief attention. First, coverage-supported beliefs should not be confused with what one might call 'shared platitudes' – that is, collectively shared knowledge that goes without saying. To be sure, '[w]hat is shared does not in general call for comment; it is too dull, trite, or familiar to stand notice' (Davidson 1984: 200), and public silence on an issue may, on reflection, reinforce one's own beliefs by validating their status as platitudes that go without saying. But coverage-reliance need not be restricted to beliefs one already holds, let alone beliefs that are collectively shared. Certainly from the perspective on an individual epistemic subject H , there is such a thing as *newly formed* coverage-supported belief.⁷ Given suitable social conditions, H can come to justifiably hold the novel belief that p , simply by reflecting on the truth-to-testimony conditional that *if $\sim p$ were true, she would have heard about it by now*.

Finally, it is worth comparing coverage-reliance with our earlier notion of epistemic dependence. Clearly, a subject H 's relying on her community for coverage is an instance of epistemic dependence, given that it renders her epistemic status dependent on social processes that supervene on the cognitive processes of one or more (other) epistemic agents – not least those of her potential informants. At the same time, there is no suggestion in the theoretical framework of coverage-reliability that, when an individual relies on her community for coverage, it is somehow the community, rather than the individual, that should be deemed the 'knower'. Coverage-reliance thus adds a new kind of epistemic dependence – a dependence that runs deeper and is more distributed in character – to the more familiar case of reliance on testimony, but it remains compatible with our practice of primarily crediting individuals with knowledge.

⁵ This is Goldberg's definition of a source A 's being *coverage-reliable* in a given domain D (Goldberg 2010: 159).

⁶ There are also obvious restrictions on what kinds of statements the truth-to-testimony conditional can be reasonably expected to cover. When it comes to certain necessary truths (e.g., an – as yet undiscovered – mathematical truth) or universally quantified statements, appealing to the truth-to-testimony conditional will be of limited use.

⁷ On this point, see (Goldberg 2010: 157).

2. Rumor-based belief, gossip, and the information grapevine

The social world allows for a vast range of informal pathways of communication, as is reflected in the large number of natural-language terms that refer to such interactions. In English, there is a dizzying array of such terms, including ‘small talk’, ‘hearsay’, ‘rumor’, ‘gossip’, ‘grapevine’, ‘scuttlebutt’, and ‘chatter’ (where the latter may refer to either ‘idle talk’ in general or, in the context of strategic communication, more specifically to intercepted – or otherwise overheard – informal communication among suspected parties). We routinely make fine-grained distinctions depending on how, and with whom, we communicate, and about which subject matter. The pragmatics of ‘gossip’, ‘rumor’ and their near-synonyms depends heavily on a variety of factors, including the content of a given claim, the social function of the utterance, its intended audience and conversational context. As is only to be expected, there is considerable overlap in the phenomena the terms are meant to describe: small talk may touch on many ‘neutral’ issues – for example, the weather – but it may also include gossip and can be used as an occasion to pass on rumors; gossip, though typically occurring on a limited scale among groups of friends, colleagues, and acquaintances, can in turn give rise to more large-scale rumor-mongering. Terms such as ‘rumor’ and ‘gossip’ are often used conjointly, and sometimes interchangeably, depending on context. What this suggests is that informal pathways of communication share a certain ‘family resemblance’, which cannot easily be broken down into non-overlapping domains (say, of ‘rumor’ *as opposed to* ‘hearsay’), each with its own distinct set of necessary and sufficient conditions. This is not to say that conceptual distinctions and clarifications are futile – indeed, the remainder of this section will attempt just that – but is meant to emphasize that any such attempts need to be well-motivated and need to be grounded in the empirical reality of informal communication and its representation in language. While a ‘family-resemblance’ approach based on ‘typical’ features of specific kinds of informal communication might give rise to the worry that there can be no counterexamples ‘since any such example may be regarded as atypical’ (Ben-Ze’ev 1994: 12), it seems to me preferable to an approach based on stipulative definitions, which would bring with it the risk of linguistic revisionism. For, even in the absence of hard-and-fast necessary and sufficient conditions, it seems nonetheless possible to identify stable conceptual differences between, for example, ‘rumor’ and ‘gossip’. In the remainder of this section, I shall compare and contrast these two, so as to bring into sharper focus some of their epistemologically salient characteristics, thereby

contributing to a tentative definition of what renders some informal communications instances of ‘rumor’.

Gossip, which has received considerably more philosophical attention than rumor (e.g., Goodman & Ben-Ze’ev 1994, Code 1995, Adkins 2002), has been variously defined by social scientists as ‘evaluative talk about a person who is not present’ (Eder and Enke 1991: 494) and ‘the process of informally communicating value-laden information about members of a social setting’ (Noon and Delbridge 1993: 25). While these definitions are refreshingly neutral with respect to the moral and ethical status of gossip – which traditionally has been the main bone of contention among philosophers – they are also incomplete. The deliberations of a search committee behind closed doors will typically involve evaluative talk about a person – the candidate – who is not present, yet they do not count as gossip. Similarly, the weekly meetings of members of a self-help group may bear all the hallmarks of a ‘process of informally communicating value-laden information’ about its members – indeed, its very point may be that all members share information about themselves and others as freely and openly as they can – yet information shared among all members of the group would not be considered gossip.

In a paper on organizational studies, Nancy Kurland and Lisa Pelled (2000) define gossip as

informal and evaluative talk in an organization, usually among no more than a few individuals, about another member of that organization who is not present. (2000: 429)

While this definition is specific to workplace gossip in larger organizations, it adequately captures the fact that one cannot gossip about the very person one is talking to (except by mistake); by replacing ‘membership in an organization’ with ‘membership in a circle of friends, colleagues, or acquaintances’ the definition can easily be extended to encompass non-workplace gossip. What all forms of gossip have in common is that they concern individuals and their (often, though not always, private) conduct, or ‘personalities and their involvements in events of the community’ (Paine 1967: 283). Rumor, by contrast, is not restricted to information or speculation about specific individuals (or groups of individuals) and their conduct – even if the primary function of rumor is often described as ‘group sensemaking’ (DiFonzo 2010: 1125). Indeed, some of the earliest empirical studies of rumor psychology, for example by Indian psychologist Jamuna Prasad (1935) in the immediate aftermath of the 1934 Bihar earthquake, found that rumors in disaster situations frequently concern natural phenomena – such as the physical extent of destruction, loss of life, or predictions of further impending calamities (such as aftershocks or tsunamis). Even where rumors

concern the social world – as in the well-studied case of stock market rumors (e.g., Schindler 2007) – they typically purport to communicate factual information of broader significance; unlike gossip, rumor is not restricted to the doings of specific human beings.

As is to be expected, given the ‘family resemblance’ characteristics of informal communications, there are borderline cases that exhibit features of both gossip and rumor. Examples include the so-called ‘celebrity gossip’ that fills the pages of the yellow press and is widely disseminated via the internet. While content-wise much celebrity gossip is of the same kind as its less illustrious cousin, in that it deals mainly with – real or fictitious – romantic relationships (and their subsequent break-ups), pregnancies, rivalries, health issues, minor misdemeanors, and so forth, its dissemination is not restricted to specific social networks. There is also no pretence that the target of the gossip should be shielded from the fact that (s)he is being gossiped about. Whereas everyday gossip is received ‘through the grapevine’ – via informal networks of communication – celebrity gossip is typically communicated via specialized outlets which, even if they are known to be unreliable purveyors of unauthorized information, lack the informality associated with gossiping among acquaintances. When gossip spreads beyond a given network of acquaintances, for example because it becomes pervasive in an organization or society, or when its perceived significance transcends narrow concerns of individual conduct (for example because it is seen in the light of larger social or organizational concerns), it may well become indistinguishable from rumor *simpliciter*. As Ralph Rosnow puts it in the form of a rhetorical question:

Sometimes it is impossible to separate rumor from gossip. When, according to unattributed hearsay, a female executive is alleged to have “slept her way to the top”, is this rumor or gossip? (Rosnow 1988: 14)

The intimation of moral impropriety is often a driving force behind gossiping; its ethical aspects are also what has attracted the most attention among philosophers – more so than its epistemological standing. (See Westacott 2000, and references therein.) Immanuel Kant, in the *Metaphysics of Morals*, goes so far as to suggest that ‘the intentional *spreading (propalatio)* of something that detracts from another’s honor – [...] *even if what is said is true* – diminishes respect for humanity’, leading him to conclude that we have ‘a duty of virtue not to take malicious pleasure in exposing the faults of others’ (Kant 1996: 212; emphasis added). Within this philosophical tradition, the moral and ethical aspects of gossip and rumor have tended to outweigh the epistemological aspects of reports received ‘through the grapevine’; it is this imbalance I wish to address.

The differences discussed so far between gossip and rumor are compatible with gossip being a subspecies of rumor. In terms of its content, gossip is more specific than rumor, in that the former is mostly about the personal, private and moral aspects of individuals and their actions, whereas the latter also serves as a medium for claims about impersonal states of affairs (e.g., natural disasters) and larger – social or non-social – developments (such as economic crises) that are not necessarily attributable to individuals. As far as its intended audience is concerned, gossip is again more restrictive than rumor. Not only does ordinary gossip typically occur only between acquaintances, but to the extent that it excludes the target of the gossip from the conversation, it is by necessity directed at a restricted audience. No such in-principle restriction exists for the audience of a rumor (even if in many cases it will be prudent, or otherwise preferable, for a rumor-monger to address only a limited audience). While both of these observations are compatible with the idea that gossip is simply a more restricted kind of rumor, several authors have suggested that the differences between rumor and gossip run deeper than this. For one, whereas the gossiper may on occasion present himself/herself as privy to first-hand evidence, rumors tend to be reported in a non-committal, second-hand way ('I have heard that...'). As Rosnow and Fine put it, without elaborating on the point, 'the basis of gossip may or may not be a known fact, but the basis of rumor is always unsubstantiated' (Rosnow and Fine 1976: 11). The suggestion seems to be that certain instances of first-hand testimony may well count as gossip, whereas rumor, by necessity, is based on hearsay that cannot – at least not momentarily – be confirmed by independent evidence. Thus, if Mary truthfully reports to her colleagues that she saw John from IT and Jane from HR passionately kissing in the office one evening, this might make for a juicy piece of workplace gossip, but it would not constitute rumor – at least not yet. If others chime in, perhaps adding further observations in support of the claim that John and Jane are having an affair, and the claim begins to circulate – no longer being attributable to Mary – what started off as a piece of (first-hand) gossip would then have 'matured' into a full-blown rumor.

The absence of independent evidence also seems to be what underlies another frequently cited definition of rumor, due to Gordon Allport and Leo Postman. In their classic study *The Psychology of Rumour* (1947), they define rumor as

a specific (or topical) proposition for belief, passed along from person to person, usually by word of mouth, without secure standards of evidence being present. (Allport & Postman 1947: ix)

Gary Fine echoes this when he holds that 'rumor's foundation is a lack of evidence – without regard for topic', contrasting it with gossip, which 'specifies the topic – the moral doings of other humans' (Fine 1985: 223). Both formulations – absence of

‘secure standards of evidence’ and ‘lack of evidence’ – are problematic, insofar as they overstate the extent to which rumors lack evidential support. Not only do many rumors have a basis in fact, but, as I shall argue in Section 4, rumor may also be the first – sometimes the only – route by which recipients learn about relevant evidence. More often than not, rumors arise not from an in-principle lack of evidence, but because relevant evidence and sources of independent confirmation are currently inaccessible to the relevant social networks. From the perspective of the recipient, mere existence of independent evidence is irrelevant, if it is out of reach – perhaps because it is temporarily unavailable, kept secret, or unintelligible to anyone lacking significant expertise. Rumor-mongering, thus, does not arise from a *suspension* of ‘secure standards of evidence’, but instead indicates a collective need for corroborating (or disconfirming) evidence to be made available. Sociologists of rumor, too, although typically concerned more with ‘the social context within which rumors develop, spread, and gain meaning’ (Miller 2006: 507) and with their social function as ‘group sensemaking’, have acknowledged that rumors are ‘first of all information claims; they purport to inform the hearer’ (DiFonzo 2010: 1125). This also explains why rumors often evolve ‘informational credentials’ – such as references to authoritative sources, explanatory detail, quantitative data, etc. – that ‘lead listeners to believe that the rumor is accurate’ (Fragale and Heath 2004: 225). Importantly, it has been argued that, as the result of heuristics governing the self-attribution of beliefs, ‘rumor-tellers attribute rumors to credible sources not just so that they will be believed by listeners, but also because they are believed by storytellers’ (Fragale and Heath 2004: 226).

Instead of building lack of evidence *in general* into the definition of rumor, it would therefore be preferable to characterize rumor as hearsay that is ostensibly propagated and presented as informative – whether or not its status as hearsay is made explicit to the recipient – and that cannot presently be verified, or refuted, by independent evidence. Note, however, that while this characterization does a better job of stressing the ‘local’ (instead of ‘in-principle’) unavailability of independent evidence in many rumor contexts, it risks overshooting the mark by being too inclusive. For one, there are many claims that may be presented as informative or factual, yet which no amount of independent evidence could possibly suffice to verify or falsify. (Perhaps certain religious doctrines fall into this category; or the claim that the universe contains a prime number of galaxies.) What rumors typically require is at least the presumption that relevant evidence *can be had* – and, often enough, is in fact enjoyed by epistemically privileged outsiders. It also needs to be considered what should count as ‘independent evidence’ in the first place. Clearly, it cannot be required that independent evidence must always be non-testimonial in character, nor that it must always allow for

the direct ascertainment of the facts. Otherwise, much of what we take ourselves to know – for example, historical knowledge of events long past – would fail this test. Yet our knowledge of history – patchy and biased though it may be – is generally regarded as of a different kind as information gained on the basis of rumor.

A first indication of what makes unavailability of independent corroboration special in the case of rumors, may be gleaned, once again, from comparing rumor with gossip. What distinguishes rumor from more ‘localized’ social phenomena (such as gossip among acquaintances) is that, as Jörg Bergmann puts it, it virtually always involves ‘unauthorized messages that are always of universal interest and accordingly are disseminated diffusely’ (Bergmann 1993: 70). Rather than lack of evidence *in general*, it is the lack of independent confirmation by authoritative sources which is characteristic of rumor: Sometimes we have enough first-hand empirical evidence at our disposal to come to a considered judgment regarding an instance of hearsay we encounter, but often we need to rely on institutional testimony – e.g., reference works, experts, encyclopedias, scientific bodies, and trusted news sources – for authoritative confirmation. When either (or both) of these – independent first-hand evidence and official sources – are unavailable, or contradict one another, conditions are conducive to the development and spread of rumor. (Perhaps it is this lack of subjectively definitive confirmation that motivates the *Oxford English Dictionary*’s definition of rumor as ‘general talk or hearsay, not based on definite knowledge’.) Rumor, thus, may be described as the propagation of ostensibly informative hearsay, usually on a topic of broader interest, and typically communicated via informal pathways in the absence of independent corroboration by either first-hand evidence or official (authoritative) sources.

Referring to official sources as ‘authoritative’ requires disambiguation, since ‘mere’ officialdom and epistemic authoritativeness can, of course, come apart: Whether or not ‘the authorities’ – that is, those who have been tasked with, or are simply recognized as performing the social function of, providing the public with information – are indeed sufficiently competent, sincere, and forthcoming in their handling of relevant information to count as *epistemically* authoritative, is an entirely contingent matter. A prime example of how merely occupying a position of power over the creation and dissemination of information does not, thereby, translate into epistemic authority, is the notorious case of Lysenkoism – the Soviet doctrine, proposed in 1928 by Trofim Lysenko as a direct assault on ‘orthodox’ genetics, that organisms pass on acquired characteristics to the next generation – which, in 1948, was decreed by the Communist Party to be ‘the “only correct position” in matters of biological science’ (Wrinch 1951:

488).⁸ Lysenko's theories had little, if any, basis in empirical evidence (and were quickly proven to be false by independent scientists), but for political expediency were promulgated as true by the Soviet establishment until well after Stalin's death in 1953.

However, the intrusion of social power over the creation and dissemination of knowledge into matters of epistemic authority is not exclusive to totalitarian state systems. Various social conditions of communication and inquiry can lead to misalignments between epistemic authority and social power (that is, between authoritativeness and officialdom). Thus, Daniel Bell, in his classic *Work and its Discontents* (1956: 28), noted how a new culture of management and corporate communications had given rise to a change 'in the culture as a whole, from authority to manipulation as a means of exercising dominion', and Miranda Fricker, more recently, has laid out in detail how systemic misalignment between 'indicator-properties' of (actual) epistemic authority and social markers of (presumptive) credibility can give rise to what Fricker calls 'epistemic injustice' (Fricker 1998: 176).

Others have been more optimistic about the overall contribution of extant processes of public communication to the reliability of our knowledge. Thus, in a paper on the epistemology of conspiracy theories relating to the 9/11 terrorist attacks, Neil Levy has argued that, if 'the alternative is cutting oneself off, to a greater or less extent, from the very resources one needs to correct one's beliefs, [then] accepting the official story is almost always rational' (Levy 2007: 190). Yet, as a general claim this seems overly optimistic. While it may sometimes be rational to accept the 'official story' – since, for example, we may realize that the situation is such that only the official source (or those acting on behalf of it) could reasonably be expected to know – this will confer *epistemic* authority on the official source if and only if it is *in fact* sufficiently competent, sincere, and forthcoming with the relevant information. And, all things considered, it is *epistemic* authority (deriving from actual competence and reliability) which is what we are aiming for when assessing the merits of an informant as a source of knowledge. Whether, in any given setting, the distribution of social power is such that occupants of 'official' positions are, in fact, epistemically authoritative is a contingent matter – and is something that anyone wishing to rely on official sources had better be able to judge. This, presumably, is why even Levy acknowledges that, for an official source to be epistemically authoritative, it must be 'properly constituted' (Levy

⁸ I am grateful to an anonymous referee for suggesting this example to me as a way of disaggregating the different senses of authority involved.

2007: 182). However, more needs to be said about when a source of information is *in fact* properly constituted, and when it is reasonable to *presume* that it is.⁹

As the preceding discussion has made vivid, there is an intrinsic tension between the two senses – institutional and epistemic – of ‘authority’ (and its cognates).¹⁰ It might be tempting to try to resolve this tension simply by *defining* epistemic authoritativeness in terms of reliability (or some admixture of purely epistemic desiderata, such as competence, sincerity, etc., on the part of the source), and likewise characterize institutional (‘official’) authority purely in terms of power relations. But the issue is not merely one of terminological ambiguity. Rather, it is an epistemological (and very real) challenge that arises from the fact that, as inquirers with imperfect information, we can only ever have direct access to certain (fallible) indicator-properties of epistemic authority on the part of our informants, since absolute measures of their competence, reliability and track record are typically out of our reach. Since we live in a social world where many indicator-properties are intimately tied to certain social roles (e.g., through processes of accreditation and credentialing), *ascriptions* of epistemic authoritativeness (though not ‘absolute’ reliability and competence *per se*) typically supervene on social facts. Recent work on the nature of epistemic authority has come to recognize the concomitant pragmatic dimension of ascriptions of epistemic authoritativeness which underlies their reason-giving role. As Arnon Keren (2007: 375) puts it: ‘To have epistemic authority on *p* is to have the ability to entitle others to form an opinion by giving others a preemptive reason for believing that *p*.’ By deferring to an epistemically authoritative informant, one’s opinion ‘becomes sensitive to evidence that the trusted person has’ (and at least partly insensitive to other factors; *ibid.*); in turn, presenting oneself as an epistemic authority amounts to an invitation to trust and conveys one’s readiness to vouch for the truth of one’s assertions.

Both gossip and rumor are primarily spread ‘through the grapevine’ – that is, via informal networks of communication – and it is this that puts them at odds with other forms of public communication, which can typically be traced back either to their official source (a press release, say). This raises the question of who can be held epistemically responsible should one’s trust in the veracity of claims received ‘through

⁹ As I will argue in the next section, certain features of the network through which information is likely to have travelled will have a bearing on whether information thus received is likely to be epistemically authoritative.

¹⁰ Miranda Fricker shares this assessment when she writes that epistemic practice ‘is *intrinsically* prone to’ mismatches between epistemic authority and social power, thus leading to the ‘ever-present risk that the norm of credibility will be socially manifested in a discriminatory manner’ (Fricker 1998: 176).

the grapevine' ultimately turn out to have been misplaced. Rumors and gossip do, of course, get taken up by the media, as in the aforementioned case of celebrity gossip, which leads to some blurring of the line between official and unofficial sources of information. However, the primacy of unofficial pathways in the communication of rumors and gossip is implicitly acknowledged even by those who, in some official capacity or other, are engaged in counter-rumor strategies (or 'reputation management'; see Helm et al. 2011): For example, one difficulty in combating extremist ideologies, which often feed on rumors about their perceived enemies, is gaining trusted access to the informal networks of communication among the relevant social fringe groups.¹¹

As a philosophical problem, what little attention rumor has attracted has mostly focused on rumor as a 'pathology of testimony' (Coady 2006: 253) or on our moral response to rumor as a norm violation (Jones 2005). As is already evident from the discussion so far, my focus in this paper is on the epistemic status of rumors, not on the moral questions associated with either the potential harm caused by rumors or the (perhaps morally reprehensible) activity of rumor-mongering. From an epistemic point of view, the problem of rumor closely parallels that of testimony as a source of knowledge – yet with the added constraints that arise from the fact that rumor is the propagation, via informal pathways of communication, of hearsay (typically of uncertain origin) in the absence of authoritative corroboration. The question then is how recipients should respond to rumors they encounter, and under what conditions rumor-based beliefs are justified. The expression *rumor-based belief* here refers to any beliefs *which one reasonably and directly forms in response to what one reasonably takes to be a rumor and which are essentially caused and sustained by it*. This definition of 'rumor-based belief' closely parallels Duncan Pritchard's definition of 'testimony-based belief' as

any belief which one reasonably and directly forms in response to what one reasonably takes to be testimony and which is essentially caused and sustained by testimony. (Pritchard 2004: 326)

This should not come as a surprise, for unless prefaced with an explicit qualifying remark (such as the phrase 'rumor has it'), a rumor will typically appear to its recipient as a *prima facie* instance of testimony.

Nothing said so far implies that rumors are *essentially* unreliable. While, according to the tentative definition above, rumors are invariably linked to hearsay, lack of first-hand evidence for a claim does, of course, not entail that it is necessarily unreliable. It would be both unwise and revisionist of normal linguistic usage to *define*

¹¹ See (Packer 2006).

rumors as unreliable, since we know from experience that rumors often have a basis in fact and, indeed, are often true. In predominantly oral cultures, the power of rumors to disseminate information is widely acknowledged; thus, the Athenians, in 467 BCE, went so far as to dedicate an altar to PHEME (rumor), as a tribute to the speed with which the news of the victory at Eurymedon had reached Athens.¹² In much the same way that the Athenians learnt about this vital military success by relying on informal networks of communication, so in modern contexts, too, rumor can be a surprisingly reliable source of information. As Theodore Caplow noted, summarizing his study of World War II military rumors over a two-year period: ‘[E]very major operation, change of station, and important administrative change was accurately reported by rumor before any official announcement had been made’ (Caplow 1947: 301). What is needed, then, is an improved theoretical understanding of the conditions under which rumor may function as a source of news rather than disinformation – not only in military contexts, but across the whole range of the social world.

3. Filters, cascades, and the reliability of rumor transmission

By definition, rumors propagate. But just how do they spread? The present section discusses three theoretical models and their significance for assessing the epistemic status of rumors. As will become clear, how one conceives of the predominant mode of rumor transmission will influence one’s take on the reliability of rumor as a source of information.¹³

The first model is implied by Allport & Postman’s earlier definition of rumor as a ‘proposition for belief, passed along from person to person’; on this model, the dominant mode of transmission of rumors is serial in character. The reliability of a rumor, on the *serial transmission* (ST) model, is simply a matter of the reliability of the original report and the faithfulness with which subsequent instances reproduce it. The idea of serial transmission informed a number of lab experiments conducted by Allport and Postman, in which an ‘eyewitness’ to an event recounted the situation to another, who in turn described it to another, and so forth – until the report reached a research associate who carefully noted the content and analysed any deviations from the original

¹² See (Webster 1954: 14).

¹³ Importantly, reliability here refers not merely to the faithful reproduction of an original claim (as will be discussed in the case of serial transmission, below), but to the epistemic standing of the claim itself: it is, after all, the epistemic qualities of an information source that we are concerned with.

account. Transmission of rumor was thus idealized as a linear series of simple, unidirectional acts of communication; more complex interactions, which might have involved the recipient questioning the speaker's reasons for the telling, or corroboration by reports from multiple speakers, were abstracted away in the study. The artificiality of the ST model engendered much criticism, not least by sociologists who pointed out some obvious 'inadequacies of mundane realism of the experimental procedure', which did not allow interlocutors 'to repeat, seek clarification, question, or interact with each other' – thus imposing restrictions that were 'unlike what occurs in everyday conversation' (Bordia 1996: 149).

At a conceptual level, too, the ST model's focus on rumors as 'stand-alone statements', passed on from one person to the next, was quickly challenged by other researchers. Thus, Tamotsu Shibutani argued that rumors 'cannot be identified in terms of any particular set of words', but only 'by abstracting from dozens of communicative acts'; rumor is 'a recurrent form of *communication through which men caught together in an ambiguous situation attempt to construct a meaningful interpretation of it by pooling their intellectual resources*' (Shibutani 1966: 17; italics original). On what may be called the *sense-making* (SM) model, rumors are collective attempts at addressing information deficits, whose provisional nature – much like that of hypotheses – means that subjects may rationally reassess their epistemic stance towards them in the light of new evidence. While this characterisation highlights the social function of rumors – as is reflected in the locution 'collective sense-making' – the SM model has been credited with drawing attention to 'the distinctively recirculatory character of rumor' (DiFonzo 2010: 1128). The idea that rumors are not merely being passed on, moving along a linear chain of interlocutors, but tend to be repeated and may thus be encountered more than once by any group member, was already pointed out by H. Taylor Buckner (1969), who also noted that

two group level variables operate to promote or retard the spreading or repeating of rumors: the structure of the group or public through which the rumor is spreading, and the involvement or interest the group has in the topic. (Buckner 1969: 63)

Multiple encounters of a rumor – which may be mistaken for independent evidence by any given recipient – are more likely to occur within closed groups than within diffuse groups or publics. Given that many rumors are related to a specific social role (e.g., to the workplace), an 'individual member of a diffuse group is likely to hear a role-related rumor while engaging in that role activity' (Buckner 1969: 64), but not in others. The SM model thus allows for a greater range of transmission scenarios – including the co-

presence of serial chains of interlocutors and loose multiple-interaction networks – as well as highlighting the phenomenon of ‘re-circulation’ of rumors.¹⁴

More recently, Cass Sunstein has argued that rumors spread across social networks via ‘informational cascades’ (Sunstein 2008: 6), as a result of which ‘affective factors, and not mere information, play a large role in the circulation of rumors of all kinds’ (Sunstein & Vermeule 2009: 215). In joint work with Timur Kuran, the notion of an ‘informational cascade’ is defined as follows:

An informational cascade occurs when people with incomplete personal information on a particular matter base their own beliefs on the apparent beliefs of others. [...] In response to their communications, other individuals, who lack reliable information, may accept that belief simply by virtue of its acceptance by others. (Kuran & Sunstein 1999: 685)

On this *informational cascade* (IC) model, the transmission of rumors is determined less by the concrete social interactions between individuals but by the prevalence of a belief, or belief system, in a social network – as perceived by a given member. Perceptions about which beliefs are shared by members of a social network one belongs to are shaped not only by concrete interactions with other individual, but may be shaped and amplified by exposure to mass media. Certain public figures or ‘opinion-makers’ may become what social network theorists call ‘supernodes’, by spreading information – including information about the relative prevalence of beliefs – across otherwise isolated social networks. Affective factors may further modulate this mode of transmission, thereby giving rise to runaway ‘informational cascades’, even in the absence of new information or substantive social interaction between individual members (unlike, say, in the case of strongly interacting ‘closed groups’ discussed by Buckner, 1969).¹⁵

One guiding intuition behind the ST model is that rumor as a source of information can only be as good as its initial input. This assumption may be challenged on two grounds. First, to the extent that the ST model takes belief in the reliable

¹⁴ Caplow (1947: 301) argues that, given the right circumstances and sufficient time, ‘[t]his re-circulation tends to eliminate variation’ in the content of a rumor; instead of disintegrating, the content of a rumor may thus become more stable over time, regardless of the rumor’s truthfulness.

¹⁵ A recent example would be the ‘birther’ movement in the United States, which – egged on by Conservative talk show hosts – demanded proof that Barack Obama was, in fact, born in the United States and which thrived both on doubts about Obama’s commitment to ‘American values’ and a shared sense of identity among (otherwise socially disconnected) ‘birthers’.

transmission of a truthful report (that p) along a chain of interlocutors to be our primary justification for forming a corresponding belief that p , it is based on the methodological supposition that any inference to the reported fact must *begin from* the existence of a reliable chain of testifiers, rather than from the report itself. This recalls the criticism leveled by Elizabeth Anscombe against David Hume's account of historical testimony:

Belief in recorded history is on the whole a belief *that there has been* a chain of tradition of reports and records going back to contemporary knowledge; it is not a belief in the historical facts by an inference that passes through the links of such a chain. (Anscombe 1973: 4)

Why, one might ask, should the case of rumors be any different? Isn't a truthful rumor simply an instance of 'oral history'? If one had knowledge of the facts concerning the reliability of the chain of transmission, such knowledge would be of obvious relevance. As recipients of testimony, however, we typically find that such information is out of our reach: If we had this much information about how a given report reached us, there would presumably be no need to take anyone's word for it in the first place.

A second criticism of the ST model's assumption that communication of information can only be as good as its initial input is based on recent arguments that testimony can sometimes generate knowledge, rather than merely transmitting it.¹⁶ As Jennifer Lackey (1999) has argued, a hearer may obtain knowledge that p from a speaker's testimony that p , even when the speaker lacks such knowledge, due to the non-transmission of doxastic defeaters on the part of the speaker. Indeed, as Lackey argues, it is entirely possible to conceive of a chain of testifiers, all of whom unanimously testify that p (where p is in fact true), yet none of whom can be credited with knowledge, due to their having undefeated doxastic defeaters of the claim in question. And yet, if such a chain were to reach its final link, Sarah, who 'has a defeater-defeater for believing that p ', it seems reasonable that Sarah 'could none the less come to know that p via the testimony of other members in the chain' (Lackey 1999: 487). Note that, in both cases, the reliability of the testimony is not what is at issue – it is not the lack of reliability that undermines the speaker's knowledge, but the presence of a doxastic defeater. Furthermore, as several authors have argued, unreliable (or otherwise unsafe) testimony, too, can generate reliable testimonial belief, for example in scenarios, where the hearer – but not the speaker – relies on the 'obtaining of a condition whose presence increases the reliability of an otherwise-unreliable piece of testimony' (Goldberg 2005: 306). One such example would be the 'silent monitoring' (Goldberg 2005: 304) of the speaker by a third party – who would speak up against the

¹⁶ I am grateful to Sandy Goldberg for pointing out the relevance of this case.

speaker's testimony if it were unreliable – which a hearer need not be aware of in order to gain knowledge on the basis of the speaker's (uncontradicted) testimony.

Both the 'collective sense-making' and the 'informational cascade' models allow for more complex patterns of communicating rumors than serial transmission. Yet, one might worry that the more one departs from a 'preservative' model of transmission – whatever its intrinsic limitations – and the more one begins to allow for embellishment and 'collective re-shaping', the bleaker the prospects might seem for conceiving of rumors as reliable. It would be hasty to conclude, though, that worries about specific modes of transmission render rumors intrinsically unreliable. For one, it would fail to explain why rumors are routinely believed in the first place. Furthermore, as David Coady has recently argued,

if you hear a rumour, it is not only *prima facie* evidence that it has been thought plausible by a large number of people, it is also *prima facie* evidence that it has been thought plausible by a large number of reliable people. And that really is *prima facie* evidence that it is true. (Coady 2006: 47)

Coady refers to anti-rumor campaigns carried out by the U.S. Army during World War II, which aimed at 'dispersing the normal channels along which rumours passed' (2006: 48). Ironically, by breaking up stable networks of transmission, the campaigns had the effect of *decreasing* the reliability and veracity of the rumors. Theodore Caplow explains this paradoxical effect as follows:

Distortion in terms of wishes and avoidance seems to be an individual rather than a group characteristic. As channels solidified, this phenomenon became comparatively rare, because of the exclusion of persons associated with previous invalidity. When they were broken up [in the course of the anti-rumor campaigns], wish fulfillment again became conspicuous. (Caplow 1947: 301)

Because the "survival" of rumors is 'partly dependent on their being disseminated by people widely known to be reliable sources', reliable interlocutors effectively act as 'filters', removing unsubstantiated rumors from circulation.¹⁷ Chances are that, by the time we encounter a rumor, it will have gone through various stages of 'filtering', thus making it more likely to merit belief, or so the argument goes. The precise form that such filtering might take depends on a number of factors, including social context. 'Silent monitoring' of the speaker by a third party is effectively a mechanism of rebuttal

¹⁷ See (Coady 2006: 47).

and – provided the hearer has no reason to distrust the third party – any such rebuttal would likely reduce the weight that the hearer gives to the speaker’s pronouncements.

The Caplow-Coady argument from filtering does not (and was not intended to) generalize to the wider claim that rumors are always, or even typically, reliable.¹⁸ Much depends on the dynamics of communication within a social network. For example, while it may be true that less reliable rumor-mongers are more likely to be disbelieved, they may also be more likely to spread rumors, thereby keeping a rumor in circulation.¹⁹ Indeed, empirical estimates of the general reliability of rumors vary widely, from negligible (Weinberg & Eich 1978) to over eighty percent (Rudolph 1973), with rumors in some organizational case studies reported to be ‘nearly 100% true’ (DiFonzo & Bordia 2006: 262). Whether or not a given rumor is reliable depends both on its origins – its basis in fact – and its mode of transmission before reaching the hearer. It also crucially depends on the social setting, since rumors in established organisational settings are typically more accurate than those that circulate outside such settings.²⁰ What the Caplow-Coady example does show is that, as a rumor penetrates more deeply into the community, its reliability need not diminish, but can in fact accrue. In short, the case regarding the reliability of rumor is still open.

4. Timeliness and epistemic penetration

The notion of coverage-reliability is meant to reflect our experience that a great deal of information reaches us, if not automatically, then certainly *as a matter of course*, in the normal conduct of everyday life. As consumers of news media, we are exposed to reports on a range of issues, from science to society and world affairs, thereby allowing us to gain knowledge about a variety of topics without much *special* effort on our part. To be sure, there is nothing conceptually necessary about our being in this happy epistemic condition – that is, if indeed we are so fortunate. Much depends on the quality

¹⁸ To mention just one possible criticism, Coady’s choice of wartime rumors as an example is a highly specialized one. In times of war, when one’s own survival, and possibly that of one’s community, is at stake – in other words, when ego-involvement is high and the action-guiding role of rumor dominates – any pressure towards truthfulness and accuracy of second-hand reports is likely to be greater than during less trying times.

¹⁹ An example would be ideologically motivated rumor-mongers, who often make it their ‘mission’ to seek out new audiences, thereby offsetting any filtering that might occur within stable audiences.

²⁰ On this point see (DiFonzo & Bordia 2002: 5).

of the epistemic sources we have previously been exposed to. Had we not had easy access to reliable sources of information throughout a significant portion of our epistemic lives – beginning with a comprehensive school education, and continuing in later life with exposure to trustworthy news sources and well-informed interlocutors – we could not now legitimately appeal to the truth-to-testimony conditional for the justification of coverage-supported beliefs.²¹ Changes in informational habits and media environment – for example, which news sources are routinely consumed – will inevitably determine the degree to which reliance on coverage is justified. Recent studies have shown, strikingly, that in contemporary media environments, ‘Some News Leaves People Knowing Less’: Even after controlling for partisanship, education and other demographic factors, In a 2011 study of 612 randomly selected residents of the U.S. state of New Jersey, it was found that those who regularly consumed certain cable TV news shows, when quizzed about their general knowledge of current events, performed significantly worse than those who consumed no TV news at all.²²

The very formulation of Goldberg’s truth-to-testimony conditional – ‘*If $\sim p$ were true, I would have heard about it by now*’ – introduces, crucially, a temporal marker. Indeed, when we rely on our social environment for coverage, we do so not only with the expectation that truths will *eventually* reach us, but that they will do so *in a timely fashion*, commensurate with our current epistemic interests. This idea is present in our folk conception of ‘being well-informed’, which implies not only having a large number of justified true beliefs on a topic, but also being ‘up-to-date on events of note’ (McDowell 1998: 422). The notion of coverage-reliability, thus, depends on our being able to rely upon our social environment to report on events and changes in the world in a timely fashion. There will, of course, be some time lag between an event (say, a political event or a scientific discovery) and its being reported. Nobody can expect information to spread instantaneously. ‘Absolute’ timeliness, however, is not what matters; what is important is timeliness with respect to a subject *H*’s epistemic interests. Goldberg makes a similar point when he demands that

²¹ Indeed, it seems incumbent to mention that the ‘we’ should be understood in the exclusive sense; after all, to this day, many people around the world continue to live in restrictive epistemic environments, with little access to education, a free press, or diverse sources of information.

²² Results of the study, which was conducted by Daniel Cassino and Peter Woolley as part of Fairleigh Dickinson University’s *PublicMind* project, are summarized in a press release (‘Some News Leaves People Knowing Less’, 21 November 2011), available online at <<http://publicmind.fdu.edu/2011/knowless/final.pdf>>, accessed on 10 January 2012.

α (the relied-upon source) must be such that, at the time t at which α is being relied upon by H , it is true that, were there some relevant discovery to be made, α would have made the relevant discovery by t , and would have reported on the matter. (Goldberg 2010: 160-1)

In other words ‘there must be some sort of coordination between the time-related expectations of H , on the one hand, and the abilities of α to make any relevant discoveries, on the other’ (ibid.).

In most cases of our relying on our social environment for coverage, the source α that we would be most likely to encounter in any particular instance (e.g., in the case of new information regarding whether or not p), would not itself be the first to have the relevant knowledge. Newspapers, broadcasters, and popular websites receive much of their information from news agencies, whose reporters in turn rely on spokespeople or informants for knowledge. As consumers of news, we stand at the end of a long chain of communication, with a time lag introduced at every link along the way. Depending on one’s position in society, degree of connectedness, and affiliation with relevant epistemic and social communities, how long it takes before one learns about a new fact may vary considerably.

Consider the case of new scientific discoveries. For example, as I am writing this text, I take myself to know that scientists have not yet discovered the Higgs boson (the mass-giving, symmetry-breaking particle predicted by the Standard Model of elementary particle physics).²³ Why? Because if it had been discovered, I am convinced that, as someone who tries to keep abreast with scientific developments, I would have heard about it by now. However, at the same time, I also know that scientists at CERN in Switzerland are actively engaged in an experimental search for the Higgs boson; I know that scientific consensus formation takes time, that large collaborative research projects may, for logistical reasons, delay the announcement of major breakthroughs, that scientific journals such as *Nature* impose an embargo on discussions with the media before the official publication date, and so forth.²⁴ Perhaps, then, the Higgs boson has already been discovered, except I have not yet heard about it. It is certainly possible. If pressed on any of these points, I might have to concede that perhaps I should not be as certain as I initially was of my knowledge that no Higgs boson has yet been discovered

²³ This passage was written in April 2012.

²⁴ *Nature*’s current embargo policy mandates that ‘[m]aterial submitted to *Nature* must not be discussed with the media, except in the case of accepted contributions, which can be discussed with the media no more than a week before the publication date’. See <<http://www.nature.com/nature/authors/policy/embargo.html>>, accessed on 23 May 2011.

– although, intuitively (and in order to avoid skepticism), I would not want to disavow my knowledge. As John McDowell points out, ‘with respect to masses of what we take ourselves to know, [...] we are quite undisturbed, at least until philosophy breaks out, by the time-lag between [potential] changes in such states of affairs and our hearing about them’ (McDowell 1998: 422).

It is worth noting that in addition to a temporal marker (‘...I would have heard about it *by now*’), the truth-to-testimony conditional also contains a second indexical, ‘I’. Whether or not *I* will be among the first to learn about, say, the discovery of the Higgs boson depends on whether I am one of the CERN researchers involved, a member of the editorial team of *Nature*, or someone who only occasionally flips through the science pages of my weekend newspaper. It is important, then, in order for *my* coverage-supported beliefs to be justified, that relevant information and evidence is not merely available in principle, but has in fact achieved a certain depth of penetration into the social fabric, as it were. Call this additional constraint on the justification of coverage-supported beliefs *epistemic penetration*. Only if relevant information is being diffused beyond the narrow confines of the initial epistemic community (in this case, the scientists at CERN) and has begun to circulate among other groups – including those that routinely function as sources of information *for me* – does it seem reasonable for me to assume that my corresponding coverage-supported beliefs are reliably formed. Once such favourable conditions are in place, however, the truth-to-testimony conditional asserts that we can gain justified belief by inferring that, if the situation was significantly different from the way we think it is, we would have heard about it by now.

The twin constraints of timeliness and epistemic penetration are related in an obvious way, in that it takes time for relevant information to reach a degree of penetration that is sufficient to underwrite the coverage-supported beliefs of an epistemic subject *H*. This raises the possibility of a mismatch between what Goldberg calls ‘the time-related expectations of *H*’ (2010: 161) and the *actual* amount of time required for relevant information to reach *H*. One might worry that there may come a point when incoming testimonial reports may themselves be justifiably dismissed on the ground that, if they were true, one would have heard about them (earlier). The worry is less pronounced in the case of newly formed coverage-supported belief (such as the belief that there has been no major foreign policy change since last week) or coverage-supported sustainment of already held beliefs (e.g., that Higgs bosons have not yet been observed), but it becomes a very real concern in the case of rejecting current testimony to the contrary.

Here is another way of making vivid what is at stake in calibrating a subject's expectation of timeliness against the 'legitimate' time lag that is inevitably associated with the diffusion of new information. Imagine a subject H receives, with a minor delay, a (novel) truthful and reliable report that contradicts a pre-existing belief. Perhaps the report in question is about an unbecoming aspect of the Prime Minister's private life and, for fear of legal action, has not been carried in the mainstream press, but is now spreading as a rumor. How should H respond? By stipulation, the report is truthful and reliable. Let us assume further that, even though H 's source α on this occasion is not her usual source of beliefs about the Prime Minister (after all, most of H 's beliefs about the Prime Minister are of a political, not a private nature), H has in the past successfully relied on information from α for many true, coverage-supported beliefs. (For example, α might be a colleague at work who is generally well-informed about matters of current interest.) Can H legitimately reject the report merely because it fails to meet the timeliness condition?²⁵ Perhaps, but it is worth noting that this leads to a counterintuitive result. After all, if H rejects α 's report that p on the basis that, if p were true, she would have heard it by now, she overlooks the fact that her hearing p from α may very well be the first time that she 'legitimately' encounters the report in question – where what is 'legitimate' depends on the time lag that is objectively appropriate to H 's position in society, degree of separation from the original source of the information, membership in relevant epistemic communities and so forth. In other words, H could not reasonably have expected to hear about p any earlier than in her chance conversation with her colleague α at the watercooler. There is a first time for everything.

5. Rumors: which ones should you trust?

As long as a situation like the one just described – where a chance report from an 'unofficial' source constitutes H 's first legitimate encounter with a relevant piece of information – is a serious possibility, the question of whether rumors are reliable or not cannot be firmly answered in the negative. Instead, one should focus on the contrastive question of why some rumors are more reliable than others.

Playing devil's advocate, one might want to suggest that the situation with rumors – in this respect at least – is not altogether unlike the case of expert testimony. In both cases, the subject matter typically does not lend itself to direct investigation by the hearer H , who lacks either access to the facts or the requisite expertise to investigate the

²⁵ Or, strictly speaking, 'fails to meet her subjective expectations of timeliness'.

claims directly. Often enough, the claims will be novel in character, and coherence with past experiences or background beliefs on the part of *H* will be of limited use.²⁶ To be sure, there are obvious differences between rumors and expert testimony. Experts typically are credited not only with knowledge in the weak sense of (justified) true belief, but also with skill, understanding, and ‘an unusually extensive body of knowledge on both primary and secondary questions in the domain’ (Goldman 2001: 92) – none of which we would expect from the average rumor-monger. Whereas the expert *asserts* his testimony (and, on various accounts of norms of assertion, may be held responsible for his claims), a rumor-monger can always claim to be merely ‘reporting’ what he has heard.

But we are not here concerned with establishing a contrast between the epistemic character of ‘the expert’ and that of ‘the rumor-monger’ – we can safely assume that they are not in the same league.²⁷ Instead, what matters is how *H* should respond to their testimony *on a given occasion*. When viewed from this angle, it is the parallels between rumor and expert testimony that become more pronounced. This observation can be strengthened further by noting, first, that just as there may be conflicting expert testimony, there may be rumors and counter-rumors. Indeed, in the arena of national security, one central goal of strategic communication consists precisely in rebutting rumors (e.g., during wartime). Second, just as rumor will likely be unreliable whenever social conditions are unfavorable to collective ‘filtering’ (see Section 3), so expert testimony can be collectively undermined by prevalent social conditions – for example, when ‘experts for hire’ cater to specific interest groups. Given these parallels, it should come as no surprise that, in assessing rumors and their rebuttals, a hearer can bring to bear criteria and strategies similar to the ones that would be appropriate for the assessment of expert testimony. She might try to engage both sides in argument and search for signs of ‘dialectical superiority’ of one side over the other.²⁸ She might monitor her interlocutor for signs of incompetence, such as overly hesitant delivery – knowing full well that extreme quickness and smoothness may also

²⁶ Take the example of rumors about the private lives of politicians, which typically expose the affairs and secrets of those who, in the past, have appeared to be squeaky-clean. What would be the novelty in reporting the transgressions of known philanderers?

²⁷ Also, highlighting the figure of ‘the rumor-monger’ risks overstating the significance of ‘character’. After all, perfectly normal people pass on rumors – recall the example of *H*’s colleague α – sometimes because they genuinely believe them to be true, sometimes as a way of engaging in small talk. Experts, too, are of course not immune to the passing on – or starting – of rumors, not least when ideological commitments are at play (see fn. 19).

²⁸ The term ‘dialectical superiority’ is Goldman’s (2001: 95).

indicate a desire, on the part of the speaker, to manipulate the hearer. Indeed, one might expect such strategies of assessment to be especially promising in the case of rumors, given that, unlike expert testimony, the content of a rumor typically requires no specialist expertise for it to be understood – thereby allowing for an immediate first check for consistency and coherence with background knowledge.

Why consider trusting (some) rumors in the first place? To see what would be problematic about a policy of outright dismissal, consider what the corresponding (first-person) truth-to-testimony conditional would look like: *If p were true, I would have heard it by now (via official channels of communication)*. The qualification ‘via official channels of communication’ is necessary in order to establish the requisite contrast between rumors and official communication, without which the question of whether to reject a rumor *qua rumor* would not even begin to make sense. However, what if official channels of communication are not open (e.g., in wartime or disaster situations when channels of mass communication are disrupted), or are unreliable (e.g., in countries with heavy censorship)? When official channels of communication cannot be relied upon – either because they are not trustworthy or they are unavailable – rumors may well be the only source of information at hand.

Does this mean that, in the absence of reliable ‘official’ information, we are psychologically compelled to believe rumors, even though they are unreliable? Not quite. What it suggests, rather, is that there exists a trade-off between levels of trust in the coverage-reliability of official sources and (warranted) trust in rumor as a source of information. In information-deprived social environments with little coverage-reliability, rumors may transmit information that could not otherwise be had. This renders rumors neither intrinsically reliable nor essentially unreliable. It does suggest, though, that, in such a setting, dismissing a rumor merely on the basis that *if the rumor were true, one would have heard about it by now*, is not an adequate response. If a rumor is told by a source that one has reason to believe is part of a by and large trustworthy social network, the ‘filter effect’ discussed in Section 3 may well outweigh any argument from the lack of coverage-reliability of official sources of information. On this – admittedly optimistic – view, reliance on rumors can be seen to be as much a factor in the very real phenomenon of epistemic dependence as expert testimony, though one that needs to be managed and calibrated with great care.

References

- Adkins, Karen C. "The Real Dirt: Gossip and Feminist Epistemology." *Social Epistemology* 16, No. 3 (2002): 215-232.
- Allport, Gordon W. and Leo Postman. *The Psychology of Rumour*. New York: Henry Holt & Co., 1947.
- Anscombe, Elizabeth. "Hume and Julius Caesar." *Analysis* 34, No. 1 (1973): 1-7.
- Bell, Daniel. *Work and its Discontents*. Boston: Beacon Press, 1956.
- Ben-Ze'ev, A. "The Vindication of Gossip." In *Good Gossip*, edited by Robert F. Goodman and Aaron Ben-Ze'ev, 11-24, Lawrence: University of Kansas Press, 1994.
- Bergmann, Jörg R. *Discreet Indiscretions: The Social Organization of Gossip*. New York: Aldine de Gruyter, 1993.
- Bordia, Prashant. "Studying Verbal Interaction on the Internet: The Case of Rumor Transmission Research." *Behavior Research Methods* 28, No. 2 (1996): 149-151.
- Buckner, H. Taylor. "A Theory of Rumor Transmission." *Public Opinion Quarterly* 29, No. 1 (1969) 54-70.
- Burge, Tyler. "Individualism and Psychology." *The Philosophical Review* 95, No. 1 (1986): 3-45.
- Caplow, Theodore. "Rumors in War." *Social Forces* 25, No. 3 (1947): 298-302.
- Chakrabarti, Arindam. "On Knowing by Being Told." *Philosophy East and West* 42, No. 3 (1992): 421-439.
- Coady, C.A.J. "Pathologies of Testimony." In *The Epistemology of Testimony*, edited by Jennifer Lackey and Ernest Sosa, 253-271, Oxford: Oxford University Press, 2006.
- Coady, David. "Rumour Has It." *International Journal of Applied Philosophy* 20, No. 1 (2006): 41-53.
- Code, Lorraine. "Gossip, or In Praise of Chaos." In *Rhetorical Spaces: Essays on Gendered Locations*, 144-153, New York: Routledge, 1995.
- DiFonzo, Nicholas. "Ferreting Facts or Fashioning Fallacies? Factors in Rumor Accuracy." *Social and Personality Psychology Compass* 4, No. 11 (2010): 1124-1137.
- DiFonzo, Nicholas and Prashant Bordia. "Corporate Rumor Activity, Belief and Accuracy." *Public Relations Review* 28, No. 1 (2002): 1-19.

- DiFonzo, Nicholas and Prashant Bordia. "Rumor in Organizational Contexts." In *Advances in Social & Organizational Psychology: A Tribute to Ralph Rosnow*, edited by Donald A. Hantula, 249-274, London: Routledge, 2006.
- Davidson, Donald. "The Method of Truth in Metaphysics." In his *Inquiries into Truth and Interpretation*, 199-214, Oxford: Oxford University Press, 1984.
- Eder, Donna and Janet Lynne Enke: "The Structure of Gossip: Opportunities and Constraints on Collective Expression among Adolescents." *American Sociological Review* 56, No. 4 (1991): 494-508.
- Fairleigh Dickinson University. "Some News Leaves People Knowing Less." [Press release by the *PublicMind* project]. 21 November 2011. Retrieved from <<http://publicmind.fdu.edu/2011/knowless/>> on 10 January 2012.
- Fine, Gary. "Rumors and Gossiping." in *Handbook of Discourse Analysis*, Vol. 3, edited by Teun A. van Dijk, 223-237. London: Academic Press, 1985.
- Fragale, Alison R. and Chip Heath. "Evolving Informational Credentials: The (Mis)Attribution of Believable Facts to Credible Sources." *Personality and Social Psychology Bulletin* 30, No. 2 (2004): 225-236.
- Fricker, Miranda. "Rational Authority and Social Power: Towards a Truly Social Epistemology." *Proceedings of the Aristotelian Society* (New Series) 98 (1998): 159-177.
- Goldberg, Sanford. "Testimonial Knowledge Through Unsafe Testimony." *Analysis* 65, No. 4 (2005): 302-311.
- Goldberg, Sanford. *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge: Cambridge University Press, 2007.
- Goldberg, Sanford. *Relying On Others: An Essay in Epistemology*. Oxford: Oxford University Press, 2010.
- Goldman, Alvin. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63, No. 1 (2001): 85-110.
- Goodman, Robert F. and Aaron Ben-Ze'ev, eds. *Good Gossip*. Lawrence: University of Kansas Press, 1994.
- Gottschalk, Louis Moreau. *Notes of a Pianist*. Edited by Jeanne Behrend, Princeton: Princeton University Press, 2006.
- Hardwig, John. "Epistemic Dependence." *The Journal of Philosophy* 82, No. 7 (1985): 335-349.
- Helm, Sabrina, Kerstin Liehr-Gobbers and Christopher Storck, eds. *Reputation Management*. Berlin: Springer, 2011.

- Jones, Ward E. "Rumor, Reproach, and the Norms of Testimony." *Public Affairs Quarterly* 19, No. 3 (2005): 195-212.
- Kant, Immanuel. *The Metaphysics of Morals*. Translated and edited by Mary J. Gregor, Cambridge: Cambridge University Press, 1996.
- Keren, Arnon. "Epistemic Authority, Testimony, and the Transmission of Knowledge." *Episteme: A Journal of Social Epistemology* 4, No. 3 (2007): 368-381.
- Kitcher, Philip. "The Division of Cognitive Labor." *The Journal of Philosophy* 87, No. 1 (1990): 5-22.
- Kuran, Timur and Cass R. Sunstein. "Availability Cascades and Risk Regulation." *Stanford Law Review* 51, No. 4 (1999): 683-768.
- Kurland, Nancy B. and Lisa Hope Pelled. "Passing the Word: Toward a Model of Gossip and Power in the Workplace." *The Academy of Management Review* 25, No. 2 (2000): 428-438.
- Lackey, Jennifer. "Testimonial Knowledge and Transmission." *The Philosophical Quarterly* 49, No. 197 (1999): 471-490.
- Levy, Neil. "Radically Socialized Knowledge and Conspiracy Theories." *Episteme: A Journal of Social Epistemology* 4, No. 2 (2007): 181-192.
- McDowell, John. "Knowledge by Hearsay." In his *Meaning, Knowledge, and Reality*, 414-443, Cambridge (Mass.): Harvard University Press, 1998.
- Miller, Dan E. "Rumor: An Examination of Some Stereotypes." *Symbolic Interaction* 28, No. 4 (2006): 505-519.
- Montaigne, Michel de. *The Complete Essays*. Translated by Donald M. Frame. Palo Alto: Stanford University Press, 1958.
- Noon, Mike and Rick Delbridge. "News From Behind My Hand: Gossip in Organizations." *Organization Studies* 14, No. 1 (1993): 23-36.
- Packer, George. "Knowing the Enemy." *The New Yorker*, December 18, 2006: 60-69.
- Paine, Robert. "What is Gossip About? An Alternative Hypothesis." *Man (New Series)* 2, No. 2 (1967): 278-285.
- Prasad, Jamuna. "The Psychology of Rumour: A Study Relating to the Great Indian Earthquake of 1934." *British Journal of Psychology* 26, No. 1 (1935): 1-15.
- Pritchard, Duncan. "The Epistemology of Testimony." *Philosophical Issues* 14 (2004) 326-348.
- Rosnow, Ralph. "Rumor as Communication: A Contextualist Approach." *Journal of Communication* 38, No. 1 (1988): 12-28.

- Rosnow, Ralph and Gary A. Fine. *Rumor and Gossip: The Social Psychology of Hearsay*. New York: Elsevier, 1976.
- Rudolph, Evan E. "Informal Human Communication Systems in a Large Organization." *Journal of Applied Communication Research* 1, No.1 (1973): 7-23.
- Schindler, Mark. *Rumors in Financial Markets: Insights into Behavioral Finance*. Hoboken: Wiley, 2007.
- Shibutani, Tamotsu. *Improvised News: A Sociological Study of Rumor*. Indianapolis: Bob Merrill, 1966.
- Sunstein, Cass R. and Adrian Vermeule. "Conspiracy theories: causes and cures." *The Journal of Political Philosophy* 17, No. 2 (2009): 202-227.
- Sunstein, Cass R. "'She Said What?' 'He Did That?': Believing False Rumors." (21 November 2008). *Harvard Public Law Working Paper*, No. 08-56. Available from SSRN, <<http://ssrn.com/abstract=1304268>>. Accessed on 13 May 2011.
- Webster, Thomas B.L. "Personification as a Mode of Greek Thought." *Journal of the Warburg and Courtauld Institutes* 17, No. 1/2 (1954): 10-21.
- Weinberg, Sanford B. and Ritch K. Eich. "Fighting Fire with Fire: Establishment of a Rumor Control Center." *Communication Quarterly* 26, No. 3 (1978): 26-31.
- Westacott, Emrys. "The Ethics of Gossiping." *International Journal of Applied Philosophy* 14 No. 1 (2000): 65-90.
- Wrinch, Pamela N. "Science and Politics in the U.S.S.R.: The Genetics Debate." *World Politics* 3 No. 4 (1951): 486-519.